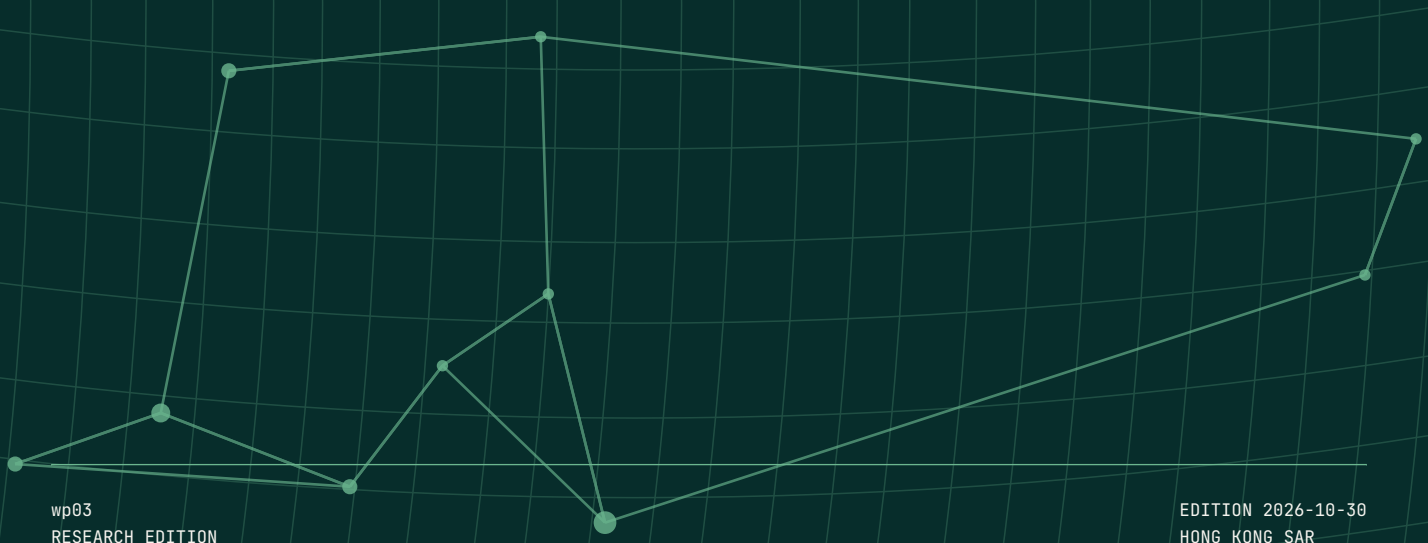
GREEN FUTURE TECH
LIMITED / HONG KONG SAR

GFT RESEARCH / WARP AND SIGNAL

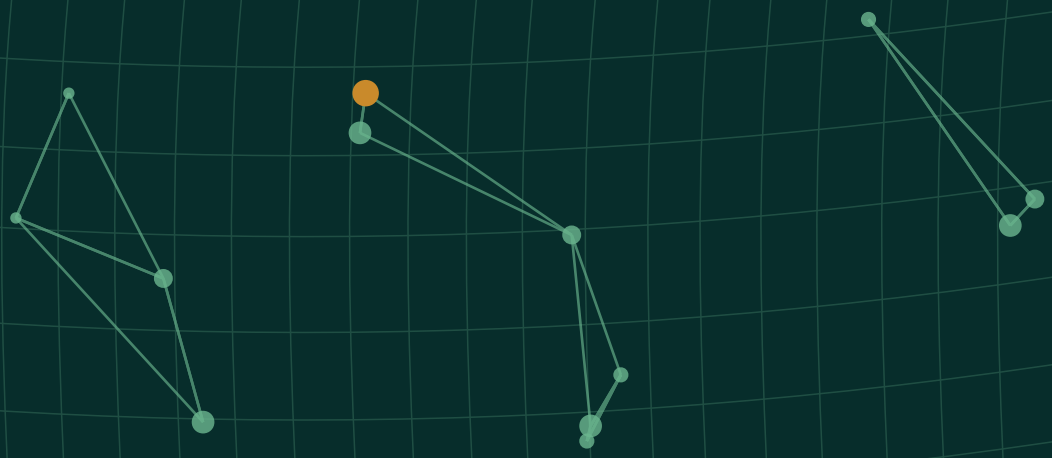
Seeing Every Defect

Computer-vision quality inspection on the garment floor: accuracy, economics and deployment



wp03
RESEARCH EDITION

EDITION 2026-10-30
HONG KONG SAR



THE GFT THESIS

Value is made where it is measured.

GFT engineers sustainable factories that capture their own data from the first shift, so raw-material economies can keep the value they grow.



Read the evidence before acting on the recommendation.

ISSUER	Green Future Tech Limited, Hong Kong SAR
DOCUMENT	wp03 / Seeing Every Defect
EDITION DATE	2026-10-30
AUDIENCE	Garment factory owners and industrial sponsors in Cameroon, CEMAC and Sub-Saharan Africa, Plant managers, quality-assurance leads and industrial engineers designing cut-make-trim lines, Chief digital officers and operational technology architects deploying edge AI on weak electrical grids, Development-finance institutions and technical partners evaluating textile-industrialisation capital plans, Institutional apparel buyers and compliance auditors specifying Acceptable Quality Limits and Digital Product Passports
STATUS	Research edition. Generated locally. Publication and operating-performance acceptance are separate decisions.
PROGRAMME	GFT's Cameroon textile programme: feasibility, performance-management system design and digital strategy. Prepared 31 January 2024; source-refresh annotations retained in 2026.

IMPORTANT NOTICE

This document is not investment, legal or tax advice. Proposed controls and illustrative calculations require project-specific validation. Targets and announced third-party projects are not evidence of delivered operating performance.

The confidential programme PDFs are not reproduced or published. No client identity or engagement code is included.

EVIDENCE REMAINS GRADED

A: official. B: institutional or major research. C: trade or local press and market-research vendors. D: unverified; excluded from use. A grade describes the supplied register and does not certify accuracy.

Inherited publication dates unrecorded: R01, R02, R03, R06, R07, R09, R10, R11, R12, R14, R15, R16, R17, R18, R24, R26, R27, R28, R29, R30, R31, R32. Temporal eligibility is unverified; these are inherited source-text checks, not independent source validation.

The operating case runs from fibre to feedback.

01 / Computer-vision defect detection is the only industrial AI application a greenfield garment factory can deploy profitably on shift one 15

1.1 / Greenfield garment makers face a cold-start data paradox from day one 16

1.2 / Published results favour machine inspection, but they do not create a common baseline 17

1.3 / Vision needs a camera rather than a multi-year sales ledger 19

1.4 / Quality consistency decides market entry before digital sophistication does 22

1.5 / A growing vendor market widens buyer choice without proving maturity 23

1.6 / Warp and Signal turns the first inspection into a management control 25

1.7 / The paper tests vision-first as an engineering and economic hypothesis 27

02 / Human visual inspection fatigues under line speed and glare, allowing up to 30 percent of defects to escape downstream 29

2.1 / Visual vigilance decays within the first thirty minutes of continuous roll watching 30

2.4 / Inter-rater subjectivity turns one physical defect into several data records 32

2.7 / The ASTM D5430 four-point system normalises roll scores but cannot correct missed flaws 34

2.10 / ISO 2859-1 AQL sampling accepts lots after production rather than controlling defects as they form 35

2.13 / A flaw missed at fabric entry can cost ten times more at packing in an explicit cost model 37

2.14 / Manual cycle time becomes a throughput constraint before it becomes a labour-saving case 38

2.15 / The manual baseline must measure its own blindness before vision is judged against it 40

03 / Textile and garment defects span more than 40 distinct classes, each with its own optical signature across warp, weft and seam 42

3.1 / Yarn irregularities need an identity before they become fabric complaints 43

3.2 / Weaving faults need fabric coordinates rather than screen coordinates 45

3.3 / Knitted structures require loop-aware definitions and a qualified presentation 46

3.4 / Wet-processing labels must separate appearance from chemical explanation 48

3.5 / Cutting faults need a pattern reference and a retained bundle identity 49

3.6 / Sewing faults need separate stitch, seam and surface judgements 51

3.7 / The quality log must preserve uncertainty and count the right event 53

04 / Line-scan cameras and multi-angle LED illumination decide 80 percent of inspection reliability before a neural network runs 58

4.1 / Moving fabric webs demand CMOS line-scan cameras synchronised up to 60 metres a minute 59

4.3 / Spatial resolution arithmetic sizes pixel pitch to single-thread diameter across wide cotton rolls 62

4.5 / Multi-angle LED illumination geometries separate holes, stains and raised three-dimensional defects 64

4.10 / Vision-based tactile sensors and contact image bars isolate structural weave relief from complex print patterns 66

4.13 / Positive-pressure air knives, sealed housings and rigid frames protect optics in tropical cotton mills 67

4.17 / Trilinear RGB calibration and inline CIELAB delta-E monitoring catch shade drift before cutting 69

4.19 / Standard optical bills of materials govern GFT's three shop-floor inspection stations 70

05 / Modern single-stage and two-stage neural detectors exceed 90 percent accuracy when kernel geometry matches fabric texture 74

5.1 / Classical spectral and statistical filters still earn their place as preprocessing layers 75

5.3 / Two-stage instance segmentation earns its compute when the boundary changes the grade 77

5.5 / Single-stage detection earns its place through direct, multi-scale prediction 78

5.7 / Geometry-specific modules address strip and micro-defects that standard kernels miss 80

5.9 / Assignment and attention decide whether close defects remain separate 82

5.10 / Public benchmarks are useful only when their protocols remain attached 83

5.11 / GFT's model-selection rules begin with the physical decision 85

06 / Unsupervised anomaly detection solves the cold-start problem when a greenfield factory has zero local defect images 87

6.1 / Why supervised classifiers stall when a new weave, shade or print hits the line 88

6.4 / Patch-distribution and memory-bank methods: PaDiM and PatchCore fit normal texture from defect-free rolls 90

6.7 / Gabor filter banks coupled with convolutional autoencoders and density estimation capture structural regularity 93

6.10 / Reverse knowledge distillation and single-sample motif extraction handle patterned textiles 94

6.13 / Controllable anomaly synthesis: GLASS and feature-space noise bridge unsupervised and supervised detection 95

6.16 / Self-supervised texture regularity loss enables pixel-level segmentation without manual masks 98

6.19 / The GFT cold-start protocol: enrol a new fabric roll in fifteen minutes 100

07 / Extending computer vision from flat fabric webs to sewn garments catches stitch, seam and dimensional faults in 0.45 seconds 107

7.1 / Incoming cloth inspection cannot release the garment created from it 108

7.2 / Cutting-table vision needs a verified coordinate chain before it can protect panels 109

7.3 / Sewing-station vision should follow the visible stitch path and retain human disposition 111

7.4 / Seam-puckering assessment needs controlled presentation and an agreed grading reference 113

7.5 / Finished-garment evidence supports local trials rather than universal accuracy promises 113

7.6 / Dimensional release requires uncertainty small enough for the buyer's tolerance 116

7.7 / Needle engineering should close the loop from repeated symptoms to verified correction 118

08 / Running vision inference at the edge keeps quality control alive when factory internet links drop or grid power fluctuates 120

8.1 / Cloud-streamed video fails under Cameroonian bandwidth and power constraints 121

8.2 / Bayesian network pruning and quantisation fit deep vision models onto embedded accelerators 123

8.4 / Two-stage edge cascades filter nominal fabric before invoking heavy detectors 125

8.5 / Two-tier station topology pairs compact capture nodes with an on-premise factory server 127

8.8 / Power conditioning and lithium-iron-phosphate buffering protect edge nodes on weak grids 129

8.11 / Store-and-forward synchronisation guarantees zero data loss into the Performance-Management System 132

8.12 / Corridor-against-north site trade-offs govern cooling, spares and solar-storage sizing 133

09 / Buying or licensing a pre-trained vision system beats building from scratch on speed, capital risk and Cold-Start accuracy 137

9.1 / Buy-before-build converts research uncertainty into an acceptance test 138

9.2 / Buy the inference engine but own the data-capture spine 139

9.3 / Three procurement tiers allocate cost and responsibility differently 141

9.4 / Vendor bake-offs must measure defect escape and operating burden together 143

9.5 / Contractual guardrails keep a licensed system replaceable 145

9.6 / Five-year cost comparison must include delay and exit, not just licence price 147

9.7 / Factory-tuned weights should replace the baseline only after evidence matures 148

10 / In a low-wage garment economy, computer vision pays back by saving fabric and preventing rework rather than replacing inspectors 151

10.1 / Low nominal wages make labour substitution a weak investment thesis 152

10.2 / Fabric recovery carries the largest addressable value 154

10.3 / Rework and seconds destroy capacity as well as minutes 155

10.4 / Buyer-failure savings remain contingent evidence, not a payback plug 157

10.5 / An illustrative plant pays back only when material recovery is counted 158

10.6 / Published benchmarks guide tests but cannot certify a factory's return 161

10.7 / Sensitivity makes utilisation and recovery quality the investment gates 161

11 / Vision only improves quality when camera bounding boxes drive hourly line huddles and upstream corrective action 164

11.1 / Deliverable 2 gives every detection a management destination 165

11.2 / One Quality Log preserves meaning from spreadsheet to edge inference 167

11.4 / Statistical control turns hourly counts into a disciplined andon 170

11.5 / Management cadences assign a decision to every timescale 172

11.6 / Roll coordinates move corrective action upstream of the needle 173

11.7 / DMAIC converts the Pareto into a tested cause and a held gain 174

11.8 / OEE and incentives must reward prevention without punishing honest detection 175

12 / Computer vision succeeds on the shop floor when designed as human-AI triage that elevates supervisors rather than alienates inspectors 178

12.1 / Why autonomous pass-or-fail without human triage fails in early factory deployment 179

12.2 / Active learning on the floor: every operator confirmation sharpens the local model 181

12.3 / Ergonomics and occupational health: eliminating backlit perch-frame eye strain 182

12.4 / Bridging Cameroon's supervisor gap through visual quality coaching 184

12.5 / Role redesign: from manual sorters to Quality Technicians and Vision Leads 185

12.6 / Curriculum for a four-week vision operations bootcamp 187

12.7 / Aligning people KPIs in the Performance-Management System 189

13 / Every defect frame logged at the edge builds an auditable quality and waste ledger for Digital Product Passports 191

13.1 / Shop-floor quality becomes traceable only when product genealogy survives every transformation 192

13.2 / ESPR readiness means portable evidence, not a speculative passport template 194

13.3 / Pre-consumer waste accounting must reconcile kilogrammes before claiming savings 196

13.4 / A buyer quality certificate should expose its evidence and its limits 198

13.5 / Lightweight MLOps keeps the ledger meaningful as optics and fabrics change 199

13.6 / Schema continuity turns spreadsheets, edge buffers and central records into one ledger 201

13.7 / Measured quality supports circular claims only when its evidentiary boundary stays visible 202

14 / A four-phase crawl-walk-run roadmap deploys computer vision in step with factory ramp and proves payback inside the PMS 204

14.1 / One gated roadmap aligns the factory ramp, the PMS and digital adoption 205

14.2 / Phase 0 proves the measurement system before approving the vision station 206

14.4 / Phase 1 stabilises two inspection boundaries through human triage 208

14.6 / Phase 2 converts verified detections into cutting decisions and local model evidence 209

14.8 / Phase 3 adds critical-seam control and traceability without destabilising the line 211

14.9 / Condition gates and paired controls decide whether the programme advances 212

14.10 / The PMS proves payback only when physical and financial records reconcile 214

Executive summary 13

Evidence register 219

Glossary 237

Term index 240

Method and imprint 241

Every exhibit states a testable claim.

EXHIBIT 01 / dumbbell
Published machine results exceed manual references, but protocols are not paired 18

EXHIBIT 02 / matrix
Vision is the only assessed application close to an inspectable shift-one decision 21

EXHIBIT 03 / area
The market projection endpoints imply 32.35 percent growth, not 33.8 percent 24

EXHIBIT 04 / flow
A shift-one inspection cell links physical disposition to the Quality Log 26

EXHIBIT 05 / line
Continuous inspection loses capture while task rotation interrupts the decline 32

EXHIBIT 06 / table
Roll grading and lot sampling make delayed decisions from incomplete observation 37

EXHIBIT 07 / waterfall
Late detection turns a small fabric fault into accumulated garment loss 38

EXHIBIT 08 / bar
Image processing takes 0.45 seconds against a derived 27-second manual equivalent 39

EXHIBIT 09 / table
Yarn observations require different confirmation views 44

EXHIBIT 10 / table
Directional observations require material-axis references 46

EXHIBIT 11 / table
Loop and alignment candidates require qualified presentation 47

EXHIBIT 12 / table
Colour and stain labels must retain their reference condition 49

EXHIBIT 13 / table
Panel and bundle judgements require different evidence 50

EXHIBIT 14 / table
Stitch integrity and seam appearance need separate decisions 52

EXHIBIT 15 / bar
Selected jacket records distinguish thread, pucker and opening faults 52

EXHIBIT 16 / heatmap
Qualification effort depends on the feature and its reference 54

EXHIBIT 17 / small-multiples
Synthetic profiles demonstrate contrast shapes without proving class identity 55

EXHIBIT 18 / flow
Release follows adjudication and the applicable product rule 56

EXHIBIT 19 / layered-architecture
Optical and mechanical stack of a 60-metre-a-minute fabric roll line-scan inspection rig 60

EXHIBIT 20 / line
Required line-scan acquisition rate rises linearly with fabric web speed and optical pixel pitch 61

EXHIBIT 21 / heatmap
Relative optical contrast of five canonical textile defect families across four illumination geometries 66

EXHIBIT 22 / table
Camera sensor and illumination selection matrix across fabric roll, cutting table and sewing stations 71

EXHIBIT 23 / bar
Pruned segmentation retains high accuracy on three fabric sets but not on the nanofibre set 79

EXHIBIT 24 / range
Geometry-specific modules improve reported mAP, but not under one common protocol 81

EXHIBIT 25 / flow
A testable hybrid keeps spectral context, spatial detail and assignment logic separate 82

EXHIBIT 26 / table
Each benchmark answers a different question and supports a different claim 84

EXHIBIT 27 / bar
Unsupervised and self-supervised models exceed 89 to 99 percent anomaly detection without labelled defect history 92

EXHIBIT 28 / table
Technical comparison of cold-start anomaly detection families for textile inspection 97

EXHIBIT 29 / matrix
Architecture selection across local defect sample volume and background pattern complexity 99

EXHIBIT 30 / flow
Fifteen-minute cold-start enrolment workflow from nominal fabric scan to active anomaly triage 100

EXHIBIT 31 / line
Illustrative twelve-month transition from Phase 1 one-class anomaly triage to Phase 2 fine-tuned multi-class detection 105

EXHIBIT 32 / flow
Inspection decisions should follow the operations that change the workpiece 109

EXHIBIT 33 / bar
Coordinate Attention gives the largest reported gain in this sewing ablation 112

EXHIBIT 34 / hbar
Reported trouser performance varies materially by class 114

EXHIBIT 35 / range
An uncertain measurement near a tolerance boundary should be reviewed 117

EXHIBIT 36 / layered-architecture
GFT five-layer reference architecture keeps vision inference at the edge and surfaces insights in the PMS 122

EXHIBIT 37 / dumbbell
Bayesian network pruning cuts inference time to 17.56 to 37.03 percent on embedded AGX Xavier hardware 124

EXHIBIT 38 / waterfall
Overlapped decoding and two-stage CAE-YOLO cascading lift Jetson Nano throughput from 9.86 to 13.45 FPS 126

EXHIBIT 39 / table
Edge hardware specification, power budget and 4-hour LiFePO4 UPS sizing across inspection stations 130

EXHIBIT 40 / bar
Garoua stations require double the LiFePO4 UPS autonomy of Douala-Dibamba rigs to bridge seasonal grid gaps 134

EXHIBIT 41 / table
Buy reusable models and build only the factory's data-capture spine 140

EXHIBIT 42 / radar
A modular licensed stack balances deployment speed with factory control 142

EXHIBIT 43 / matrix
Technical accuracy cannot compensate for weak data rights or exit terms 145

EXHIBIT 44 / waterfall
Licensing removes research cost but retains a five-year support floor 147

EXHIBIT 45 / flow
Factory-tuned weights earn deployment through a gated comparison 149

EXHIBIT 46 / bar
Low monthly wages cannot carry a camera case based on headcount removal 153

EXHIBIT 47 / dumbbell
A large detection gain does not translate one-for-one into first-pass yield 156

EXHIBIT 48 / waterfall
Fabric recovery dominates the illustrative annual value bridge 160

EXHIBIT 49 / heatmap
Low utilisation defeats payback even when defect prevalence rises 162

EXHIBIT 50 / keyfigures
Garment studies support integrated control rather than isolated detection 166

EXHIBIT 51 / flow
A detection becomes valuable only after action and verification close the loop 168

EXHIBIT 52 / table
One event can inform six KPI families without creating six records 169

EXHIBIT 53 / line
The andon fires on a process signal rather than on every defect 171

EXHIBIT 54 / sankey
Closed-loop routing separates avoidance and repair from final packing 174

EXHIBIT 55 / flow
Human-AI triage reserves judgement for exceptions and corrective action 180

EXHIBIT 56 / stacked
Illustrative triage moves inspector time from scanning to verification and coaching 183

EXHIBIT 57 / matrix
Greater decision authority requires broader demonstrated competence 186

EXHIBIT 58 / table
A four-week bootcamp certifies observed decisions rather than screen familiarity 188

EXHIBIT 59 / layered-architecture
A conserved garment genealogy carries provenance into quality and passport evidence 193

EXHIBIT 60 / matrix
Vision covers process evidence while composition and durability still need external proof 195

EXHIBIT 61 / sankey
Two illustrative fabric inputs reconcile into packed, offcut and rejected mass 197

EXHIBIT 62 / table
A light MLOps cadence preserves evidence before and after every model change 200

EXHIBIT 63 / gantt
Four conditional phases sequence inspection capability with factory and PMS maturity 206

EXHIBIT 64 / table
Each phase releases capital only after operational and economic evidence agrees 212

EXHIBIT 65 / matrix
Six deployment risks require both a technical control and a management response 213

EXHIBIT 66 / keyfigures
External benchmarks guide acceptance targets but do not prove local value 214

EXHIBIT 67 / waterfall
Material avoidance, not inspection labour, carries the illustrative monthly value case 217

Build the factory and its evidence as one system.

A greenfield garment manufacturer faces an immediate operational paradox: commercial survival depends on consistent quality and high capacity utilisation, yet the factory opens with zero historical production data and a thin pool of experienced quality supervisors. In Green Future Tech Limited's Cameroon textile programme (prepared 31 January 2024, sources refreshed 2026), that constraint governs the sequence of digital adoption. [R01] [S0801] While machine-learning demand forecasting and generative design pipelines require accumulated sales history or specialist engineering teams, computer-vision defect detection requires only an industrial camera, controlled illumination and a pre-trained model running at the edge. Published textile-inspection literature demonstrates that manual visual inspection captures only 60 to 70 percent of fabric defects, missing 20 to 30 percent under operator fatigue and line speed, whereas automated computer-vision systems achieve above 90 percent accuracy across more than 40 defect types at speeds up to 60 metres a minute. [R29] This paper sets out the engineering, economic and operational blueprint for deploying computer-vision quality control across an illustrative 3.0 million garment a year cut-make-trim facility in Cameroon. Across 14 chapters in four parts, we examine manual inspection failure, optical hardware, deep-learning and cold-start anomaly architectures, edge inference on weak grids, build-against-buy economics at Cameroon's XAF 43,969 a month minimum wage floor, and closed-loop integration into the Performance-Management System. Every figure is sourced, derived from cited inputs or labelled an illustrative estimate. Nothing here constitutes investment, legal or tax advice.

01 Manual visual inspection on textile rolls and garment tables achieves only 60 to 70 percent accuracy and misses 20 to 30 percent of defects because human inspectors fatigue rapidly under repetitive visual load, whereas computer-vision systems exceed 90 percent accuracy across more than 40 defect categories at web speeds up to 60 metres a minute. [R29] [S01] [S02] [S03] [S08]

02 Computer-vision defect detection is the optimal Tier-1 digital entry point for a greenfield factory because it operates on pre-trained convolutional, transformer or one-class anomaly models that require a live camera feed rather than multi-year historical sales datasets. [R27] [R29] [S04] [S05] [S06]

03 The global artificial-intelligence market for fabric defect detection is expanding from US$523.4 million in 2024 to a projected US$6,522.1 million by 2033 at a 33.8 percent compound annual growth rate, with Asia-Pacific holding about 46 percent in 2024, creating a mature vendor base from which African entrants can license proven systems rather than build bespoke models. [R28] [S01] [S07]

04 Unsupervised patch-distribution and memory-bank algorithms (PaDiM, PatchCore, reverse knowledge distillation and Gabor-autoencoder hybrids) solve the cold-start training problem on new cotton weaves and dyed fabrics by learning nominal texture regularity from defect-free samples alone, reaching up to 99.6 percent image-level AUROC on industrial benchmarks. [S04] [S05] [S06] [S12] [S13]

05 Single-stage object detectors from YOLOv5 through YOLOv11 enhanced with attention modules and adaptive spatial convolutions detect small, elongated warp, weft and sewing defects in real time, improving mean Average Precision (mAP@0.5) by 3.5 to 8.1 percentage points on industrial benchmarks such as Tianchi, AITEX and TILDA. [S01] [S09] [S10] [S11] [S16] [S17]

06 Extending vision from flat fabric inspection frames to garment sewing lines captures broken stitching, open seams, seam slippage, skipped stitches and seam puckering with 86.7 to 97.16 percent classification accuracy, processing frames in 0.45 seconds or less and delivering a 12 to 18 month payback in developing-economy apparel plants. [S17] [S18] [S19] [S20] [S21] [S22]

07 Network pruning, FP16/INT8 quantisation and two-stage autoencoder-plus-YOLO cascades allow vision inference to run entirely on low-cost edge hardware (such as NVIDIA Jetson Nano or AGX Xavier/Orin) at 13.45 frames a second or 17 to 37 percent of unpruned latency, isolating quality control from Cameroon's variable internet links and power interruptions. [S02] [S07] [S14] [S23]

08 In a greenfield Cameroonian cut-make-trim plant sized at an illustrative 3.0 million garments a year with a direct labour floor of XAF 43,969 a month (about US$76, against US$204 in Cambodia), vision inspection pays back primarily by intercepting fabric flaws before cutting and preventing rework, rather than by replacing low-cost inspection labour. [R12] [R14] [R17] [R29] [S19] [S24]

09 Writing edge defect detections directly into the Performance-Management System's Quality Log (Defects Per Hundred Units and Right First Time) connects camera bounding boxes to daily line huddles, weekly plant reviews and supervisor training, helping the venture avoid the 10.83 to 48.9 percent capacity-underutilisation trap seen in legacy domestic plants. [R11] [R12] [R29] [S24] [S25] [S26]

10 Time-stamped defect imagery linked to SODECOTON ginning lots creates an immutable quality and pre-consumer waste record that supports institutional tender verification domestically and prepares the factory for European Union Ecodesign for Sustainable Products Regulation (ESPR) Digital Product Passports. [R01] [R03] [R06] [R07] [S27] [S28] [S30]

01

Computer-vision defect detection is the only industrial AI application a greenfield garment factory can deploy profitably on shift one

When a new factory has no sales history and no legacy data spine, pre-trained computer vision can begin inspection immediately, with payback dependent on verified waste and rework savings.

CHAPTER 01
DECISION TEST

When a new factory has no sales history and no legacy data spine, pre-trained computer vision can begin inspection immediately, with payback dependent on verified waste and rework savings.

Within the industrial AI portfolio assessed in this paper, computer-vision defect detection is the only application that can begin useful work on shift one with a credible route to payback. Its input is present at the factory gate: light reflected or transmitted by fabric. A pre-trained or one-class model can turn that input into an inspection decision before the business has accumulated a sales ledger, a maintenance history or a seasonal demand record. The economic case is equally specific. Vision must prevent flawed material and defective work from consuming further value. It does not earn a return merely because an algorithm produces a confident label.

That recommendation is conditional, not a promise of automatic profitability. A greenfield factory still needs controlled optics, representative local samples, independent human labels, an operator disposition workflow and a measured cost baseline. Published accuracy figures establish technical plausibility across particular datasets. They do not establish performance on Cameroonian cotton, a buyer's quality standard or a factory's cash flow. The vision-first thesis therefore means deploy the measurement capability first, then release capital and autonomy only as local evidence earns it.

1.1

Greenfield garment makers face a cold-start data paradox from day one

A new garment factory must sell consistent quality before it has had time to learn what consistent quality means in its own operating conditions. Fabric lots vary. Operators are learning new routines. Needles, tension settings, lighting and material handling change together. Yet the plant cannot postpone quality control until its databases are full. Every uncorrected escape can become rework, a downgraded garment or a rejected delivery. The first production period is therefore both the least data-rich and one of the most consequential periods to measure.

The scale of Cameroon's processing ambition makes this more than a narrow software-selection problem. The inherited evidence register records seed-cotton production of 314,455 tonnes in 2023 and about 340,000 tonnes in 2024, with a projection above 350,000 tonnes for 2025. These are sector observations and a projection, not garment-factory output or evidence of realised demand. [R01] [R02]

The same register reports that about 95 percent of cotton fibre is exported raw and about 5 percent is processed locally. It also records the national objective of processing 50 percent locally by 2030. The baseline and target are not a forecast that new garment plants will succeed. They indicate why domestic conversion capacity matters and why a slow digital learning curve would be costly. [R06] [R07]

Macro-industrial data reinforce the need for careful execution rather than easy optimism. World Bank observations put manufacturing value added at 13.36 percent of Cameroon's GDP in 2022 and 12.88 percent in 2024. This economy-wide indicator does not diagnose the textile sector or prove that inspection investment will reverse the change. It does prevent a cotton-processing target from being mistaken for an operating plan. [S36]

A greenfield plant's data paradox has two distinct halves. The first is absent history: no local record of which styles sell, which suppliers fail, which machines drift or which operators need coaching. The second is immediate physical evidence: every roll, panel and garment already presents a visible surface that can be measured. Demand forecasting belongs to the first half. Visual inspection belongs to the second. Treating both as if they required the same data foundation would delay a control that can start now.

This distinction changes the meaning of pre-training. A vision model does not need the new factory's past sales, but it is not knowledge-free. Supervised detectors bring features learned from earlier image collections. One-class methods bring pretrained embeddings and then model what nominal material looks like at the local station. PatchCore explicitly addresses the cold-start problem by fitting with nominal, non-defective example images only. [S04] PaDiM likewise uses a pretrained convolutional network for patch embeddings and represents the normal class with multivariate Gaussian distributions. [S05]

The paper's shared operating model of an illustrative 3.0 million garments a year makes the sequencing concrete without pretending that such a plant exists. At launch, its sales table is empty, but its fabric-entry and inspection tables can generate useful records immediately. The correct first dataset is therefore not an artificial sales history. It is a versioned collection of local images, defect dispositions, roll identities and operator confirmations linked to physical production.

GFT's Cameroon textile programme places vision defect detection first because its value-capture boundary sits inside the factory and its minimum input is available with the first material lot. Forecasting and generative design remain valid later applications. They simply depend on different evidence and different capabilities. The recommendation is a sequence for industrial adoption, not a claim that other AI uses lack value.

1.2

Published results favour machine inspection, but they do not create a common baseline

Manual inspection supplies the counterfactual against which vision must be judged. The inherited inspection summary reports 60 to 70 percent manual accuracy and 20 to 30 percent of defects missed. It separately reports computer-vision performance above 90 percent across more than 40 defect types and line-scan inspection at up to 60 metres a minute. These figures define the programme's screening hypothesis. They are not measurements from a single controlled comparison on one fabric, one defect catalogue and one acceptance rule. [R29]

The research literature supports the reason for seeking an alternative, not a universal numerical uplift. A recent review describes manual oversight as subjective, inefficient and inconsistent in high-speed manufacturing. It also identifies dataset limits, domain generalisation and computing constraints as continuing challenges for automated inspection. The same source that supports machine inspection therefore warns against treating a laboratory metric as a factory guarantee. [S01]

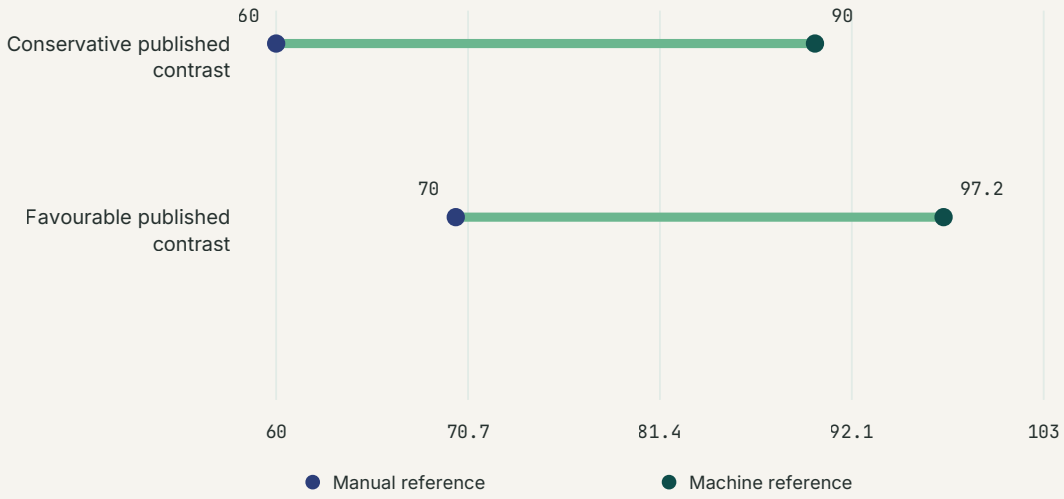
Precision about denominators matters here. Detection accuracy, defect recall, image-classification accuracy and the share of defects missed are not interchangeable. Nor should the reported manual miss share be subtracted mechanically from a machine accuracy reported elsewhere. A valid local baseline must state how many truly defective units were presented, how many were detected, how many good units were rejected and who established the reference label. Without those counts, an impressive percentage can hide a weak acceptance decision.

Reported machine results nevertheless justify a floor trial. Deep-GD, for example, reports 97.16 percent overall classification accuracy across 12 classes using patterned-garment images from the HKBU database. That is a defined study result, not evidence that every fabric construction or sewing defect will reach the same value in a new plant. [S18] The broad inspection summary's references to more than 40 classes and speeds up to 60 metres a minute similarly describe technical reach, not a service-level commitment for an unconfigured station. [R29]

Exhibit 1, p. 18 keeps the source boundaries visible. The lower machine marker represents the inherited summary's above-90-percent threshold conservatively at its boundary. The upper marker is the separate Deep-GD result. The pairings are comparison devices, not before-and-after observations. [R29] [S18]

EXHIBIT 01 / PERCENT ACCURACY, MIXED STUDY PROTOCOLS

Published machine results exceed manual references, but protocols are not paired



SOURCED [R29] [S18] Manual references of 60 and 70 percent and the above-90-percent vision threshold come from the broad inherited summary R29. The lower machine marker plots that threshold at 90, not as an exact mean. The upper machine marker is the separate 97.16 percent, 12-class Deep-GD result in S18. These are not paired trials, identical denominators or local factory results.

The exhibit supports a decision to test, not a decision to believe. A local acceptance trial should present the same independently labelled units to the existing manual process and the candidate vision process. It should report critical-defect recall, false alarms, review time and final disposition by fabric family. It should also preserve borderline cases rather than discard them, because disagreement between inspectors and machines is part of the operating evidence.

WORKED EXAMPLE: CALCULATE INTERCEPTION BEFORE CLAIMING SAVINGS

Take an illustrative cohort of 100,000 garment-equivalents of incoming fabric and assume that 4.0 percent contains a defect that should stop the affected material before cutting. The prevalence is an explicit planning assumption, not a factory observation. The capture references are the broad manual and machine ranges cited above, used here only to demonstrate the calculation. [R29]

The illustrative arithmetic gives 100,000 × 0.04 = 4,000 truly defective garment-equivalents. At 60 to 70 percent capture, manual inspection finds 2,400 to 2,800. At the conservative 90 percent machine boundary, vision finds 3,600. The incremental flagged volume is therefore 800 to 1,200 units in this constructed cohort, before considering false alarms or what happens after a flag. [R29]

Now assume, in this illustrative case, that operators confirm and act on 75 percent of those incremental flags early enough to prevent further processing. The result is 600 to 900 additional garment-equivalents diverted before cutting. This is a disposition outcome, not a detection score. If the flagged material is recut, spliced, claimed from a supplier or discarded, each route has a different cost and must be recorded separately.

No monetary saving should be assigned yet. In this illustrative case, the annual benefit is the count of additional actionable diversions multiplied by the verified net saving for each disposition route, less false-alarm review cost and annual vision operating cost. Payback is installed cost divided by that net annual benefit. Missing fabric prices, escape costs and station quotations are decision inputs, not invitations to insert convenient values.

The worked example explains why a higher detection percentage is not itself profit. Value appears only when detection changes a physical disposition early enough to avoid cost. It also explains why false positives belong in the economic model. An unnecessary review consumes time, and an unnecessary rejection can waste good fabric. A local business case must count both sides rather than treating every machine flag as recovered value.

1.3

Vision needs a camera rather than a multi-year sales ledger

The strongest reason to put vision first is not that it is fashionable or simple. It is that its decision input arrives with production, whereas other applications must infer a local operating pattern from records that do not yet exist. This is a difference in data dependency. It affects what can be validated, what can be automated and what can produce cash-relevant evidence during launch.

Demand forecasting illustrates the contrast. McKinsey's fashion analysis reports potential forecast-error reductions of 20 to 50 percent and inventory-cost reductions of 35 to 40 percent when AI is applied at scale. Those are sector-level potential outcomes, not a greenfield entitlement. [R27] A new factory can obtain a generic forecast from an external model, but it cannot test that forecast against its own realised orders, returns and style substitutions until those events exist.

A forecast also has a longer causal route to operating value. The prediction must change a buying, cutting or scheduling decision, and the consequence must later be observed. Early forecasts may be useful planning scenarios, but labelling them as locally trained demand intelligence would confuse transferred knowledge with factory evidence. The honest starting point is a documented planning assumption, followed by forecast-error measurement as actual orders accumulate.

Inventory optimisation adds further dependencies. It needs meaningful demand distributions, lead times, service rules, minimum order quantities and reliable stock identities. Some of these can be specified by management on the first shift. Others must be learned from supplier and production behaviour. The application can therefore be configured early, but its local optimisation claims should remain provisional until the underlying records stabilise.

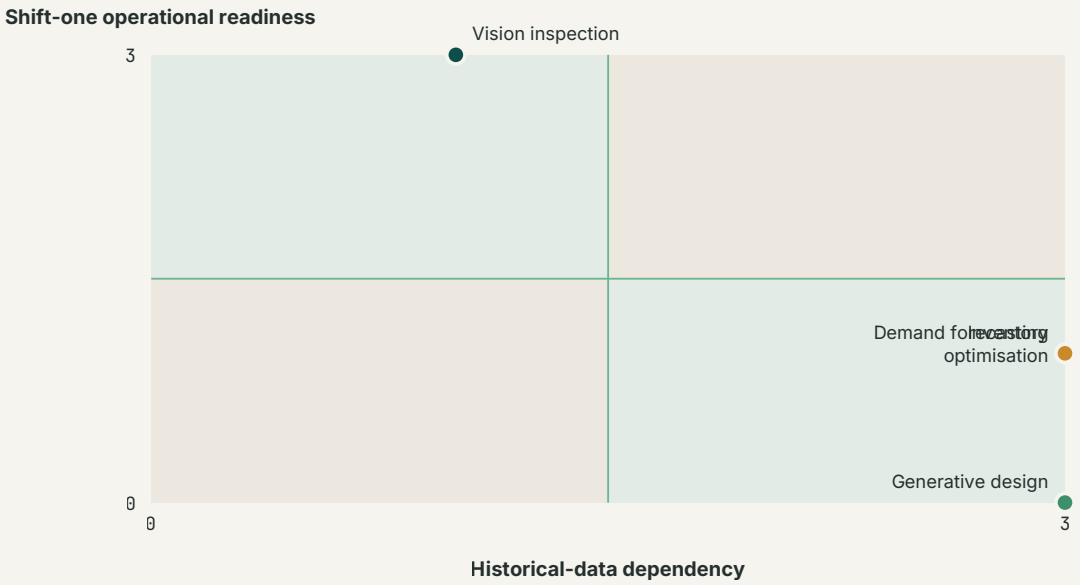
Generative design has a different constraint. The inherited register describes apparel use in 2023 and 2024 as concentrated mainly in image generation across style, silhouette, colour and pattern. [R30] Later work evaluates multiple models against fabric construction, shade, weave and silhouette, which shows expanding technical scope. [R31] [S31] Yet design output still requires product judgement, pattern engineering, material testing and buyer acceptance. It does not inspect the physical garment being made now.

The comparison is not that vision needs no data. It needs the right data at the right boundary. Nominal images can support one-class enrolment, while pretrained features supply a starting representation. [S04] [S05] The factory must still collect examples of its actual lighting, fabric, defects and operator decisions. What vision avoids is the wait for multi-period commercial history, not the discipline of local validation.

Exhibit 2, p. 21 converts that reasoning into an explicit analytical score. Historical-data dependency runs from none to multi-period. Shift-one readiness runs from research activity to an inspectable physical decision. GFT derived the positions from the cited application requirements; they are not percentages or measurements reported by the sources. [R27] [R29] [R30] [R31] [S04] [S05] [S31]

EXHIBIT 02 / GFT ANALYTICAL SCORE, 0 LOW TO 3 HIGH

Vision is the only assessed application close to an inspectable shift-one decision



DERIVED [R27] [R29] [R30] [R31] [S04] [S05] [S31] Derived ordinal coding, not source measurements. Dependency: 0 none, 1 current nominal images, 2 several records, 3 multi-period history. Readiness: 0 exploratory, 1 weakly validated, 2 assisted decision, 3 inspectable physical decision. Vision still needs nominal local frames. Forecasting and inventory need accumulated transactions. Generative design needs specialist product validation.

The matrix also shows where the argument could fail. If a vision station cannot obtain representative nominal frames, if its optics hide the target defect or if no operator can resolve alerts, its practical readiness falls. Conversely, a factory with unusually good inherited order records may make forecasting useful earlier. Greenfield status alone does not determine the score; evidence availability does.

GFT NOTE

Commission vision as assisted inspection before commissioning autonomous acceptance. On the first fabric family, freeze the camera geometry and lighting, assemble an independently labelled local test set, and record the existing manual disposition as the counterfactual. Release each additional fabric family only after critical-defect recall, false-alarm workload and operator agreement meet the buyer-linked acceptance rule. Keep demand forecasting and generative design in measured pilot lanes until their own input records become dependable.

This recommendation preserves the speed advantage without confusing speed with haste. A camera can generate events immediately. The governance around those events determines whether they become production knowledge. The first useful output is not a fully autonomous factory. It is a repeatable disposition with an auditable image and a named human decision where uncertainty remains.

1.4

Quality consistency decides market entry before digital sophistication does

A new garment producer enters a market through acceptable product, delivered repeatedly. Computer vision matters commercially because it can make inspection more consistent and its evidence easier to retrieve. The buyer does not purchase an algorithm. The buyer purchases a garment whose material, construction and dimensions meet an agreed specification. The digital system must remain subordinate to that transaction.

The market base case also needs discipline. The inherited evidence records the termination of Cameroon's AGOA eligibility effective 1 January 2020 and lists the country as ineligible for calendar year 2025. On that evidence, the venture should not make duty-free United States access a foundation for its launch economics. Domestic and regional buyers provide the more defensible planning boundary, subject to current trade advice before any commitment. [R18]

Market-size estimates do not remove the need for product-level proof. One commercial research estimate places the African textile market at US$37.3 billion in 2025 and projects US$49.41 billion by 2030. These are vendor estimates for a broad regional category, not measured demand for a Cameroonian entrant or a substitute for buyer discovery. [R24] Quality consistency helps a factory compete for a specified order; it does not create the order.

Domestic operating history offers a warning without supplying a simple cause. The inherited register records CICAM's market share falling from about 85 percent three decades earlier to about 5 percent. It also records July 2023 utilisation of 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, alongside about CFAF 2 billion of underutilisation cost in 2020. These observations do not prove that inspection failure caused the decline. They show how little room a producer has when capacity and commercial performance are already under pressure. [R11] [R12]

For institutional uniforms and repeatable ready-to-wear, a useful quality record answers practical questions. Which roll supplied the panels? Which defects were found? Which units were repaired, downgraded or released? Which model and threshold produced the flag? Who resolved an ambiguous case? Vision can preserve this chain at the point of work. A manual tally entered later can lose the image, location and sequence that make corrective action possible.

The record still has limits. An image of a visible flaw does not establish colourfastness, seam strength, restricted-substance compliance or every dimensional requirement. A stored pass label is not a buyer certificate unless the underlying inspection method and acceptance rule are agreed. The factory should present vision evidence as proof that defined checks occurred, not as proof of properties the station never measured.

Recent African factory research makes the operational possibility tangible. The inherited paper register reports image processing in 0.45 seconds a sample and a 60-fold throughput improvement on consumer-grade hardware in an Addis Ababa trouser

study. The same entry records an estimated return on investment of 12 to 18 months. These are study-specific results and estimates, not a supplier quotation, a GFT result or a guaranteed payback for Cameroon. [S19]

The wage context clarifies where the return must come from. Cameroon's monthly minimum wage rose to XAF 43,969 from 23 February 2024, recorded as about US$76 in the inherited register. A business case based only on removing inexpensive inspection labour can therefore be fragile. Avoided fabric waste, reduced rework, fewer downgrades and protected delivery performance must carry the economic argument. [R14]

1.5

A growing vendor market widens buyer choice without proving maturity

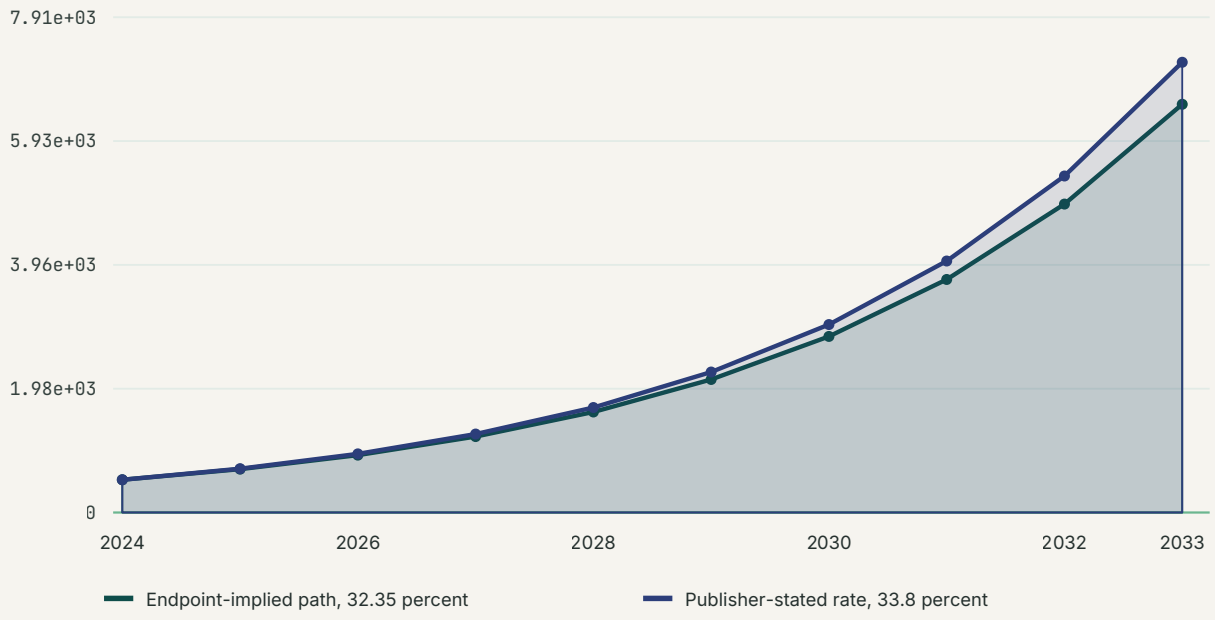
A greenfield entrant should prefer a tested inspection component over a bespoke research programme, but market growth is only a weak proxy for what can be bought and supported. The inherited market entry values the global fabric-defect-detection AI market at US$523.4 million in 2024 and projects US$6,522.1 million by 2033. It also states 33.8 percent compound annual growth and an Asia-Pacific share of about 46 percent in 2024. These are commercial market-research estimates, not audited shipments or a supplier census. [R28]

The published endpoints and growth rate do not reconcile. GFT arithmetic gives (6,522.1 ÷ 523.4)^(1 ÷ 9) − 1 = 32.35 percent over the nine annual intervals from 2024 to 2033. Applying the stated 33.8 percent instead produces US$7,193.5 million in 2033. Both calculations use the inherited source inputs; neither should be silently adjusted to make the register look consistent. [R28]

Exhibit 3, p. 24 therefore shows two geometric paths. The endpoint-implied path connects the published start and end values. The stated-rate path preserves the published growth rate and exposes its different endpoint. These are arithmetic interpolations, not annual market observations.

EXHIBIT 03 / US$ MILLION, GEOMETRIC INTERPOLATION

The market projection endpoints imply 32.35 percent growth, not 33.8 percent



DERIVED [R28] Derived from R28 endpoints of US$523.4 million in 2024 and US$6,522.1 million in 2033. Endpoint-implied series: 523.4 × (1 + 0.3235129431)^n, where n runs from 0 to 9. Stated-rate series: 523.4 × 1.338^n. Rounded to US$0.1 million. Intermediate years are calculated paths, not reported annual observations. The source's stated 33.8 percent reaches US$7,193.5 million in 2033 and therefore conflicts with its endpoint.

The inconsistency does not invalidate the decision to investigate commercial systems. It does invalidate using the headline forecast as evidence of payback or supplier capacity. Procurement should rely on dated quotations, reference-site evidence, exportable model and data terms, local support arrangements and an acceptance test on the factory's material. A regional market share does not establish that any vendor can maintain a station in Douala or Garoua.

Recent edge research shows both the opportunity and the remaining work. One study deployed a two-stage fabric-inspection cascade on an NVIDIA Jetson Nano, but its disjoint benchmark contained 249 images, of which only 20 were defective. At a recall-prioritised threshold, the first stage detected all 20 defective images while recording a 49.3 percent false-positive rate on non-defective images. The authors positioned the system as assisted triage rather than autonomous acceptance. [S07]

The same study reported 13.45 frames a second for its parallel pipeline against 9.86 for a sequential detector-only loop, while attributing 91 percent of the speedup to overlapping image decoding with inference rather than to the cascade itself. It also reported that 85 percent of forwarded frames were false alarms. These details are more useful to a buyer than a generic real-time claim because they reveal where workflow and computing gains actually arise. [S07]

Buyer adoption is consequently a qualification exercise. The factory should acquire the right to run inference locally, retain its images and annotations, export quality events in an open schema and change integrator without losing its operating records. It should buy demonstrated capability, not a promise that market growth has already solved local inspection.

1.6

Warp and Signal turns the first inspection into a management control

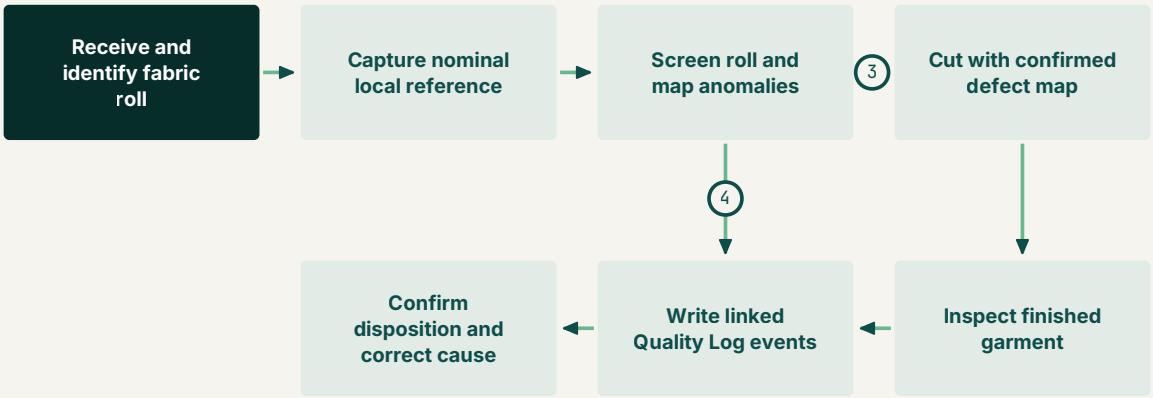
The vision-first thesis becomes operational when a camera event changes work. Production is the warp: rolls become panels, panels become garments, and garments become released orders. Measurement is the weft: identity, image, defect, disposition and corrective action cross that material flow at defined points. A detection without those connections is only an automated observation.

The first control point is material intake. The station verifies the roll identity before interpreting its pixels. Nominal frames establish the local reference where a one-class method is used. The inspection output then needs a route: accept, hold for review, map for cutting avoidance, claim against the supplier or reject. Each route has an owner and a different economic consequence. Without that routing table, higher recall can simply create a larger queue.

Exhibit 4, p. 26 sets out the proposed minimum flow. It is an illustrative GFT architecture, not an installed system or a claim about an actual supplier relationship. Its important feature is not the number of gates. It is that material disposition and the quality record retain the same identity from intake to corrective action.

EXHIBIT 04 / ILLUSTRATIVE PROCESS AND INFORMATION FLOW

A shift-one inspection cell links physical disposition to the Quality Log



3. Accepted with map

4. Flags and holds

ILLUSTRATIVE [R29] [S04] [S05] Illustrative GFT process design informed by the inspection context in R29 and nominal-image cold-start methods in S04 and S05. It does not depict a commissioned GFT installation, a verified supplier relationship or measured factory performance. Human confirmation remains required where the acceptance rule or model confidence does not support automatic disposition.

A useful event record is deliberately plain. It carries the roll or bundle identity, station, timestamp, fabric family, model version, threshold, defect class where known, confidence, image reference and operator disposition. It also records the final physical route. This allows the factory to distinguish a correct detection from a correct business outcome. A true flaw that is still cut into a garment is a failed control even when the model worked.

The Quality Log should compute Defects Per Hundred Units and Right First Time from consistent event definitions. The first measure counts defects rather than only defective garments, so repeated faults on one unit remain visible. The second measures whether units complete without rework. Vision can automate entries, but management must freeze the denominator and disposition rules before comparing periods.

Daily review then becomes specific. A supervisor can inspect the image behind a recurring fault, locate its station or roll, and assign a mechanical, material, method or training action. Upstream review can compare fabric lots using the same defect dictionary. The purpose is not to automate a scorecard. It is to shorten the time between a physical deviation and a verified correction.

This architecture also suits a weak connectivity environment because the essential decision can remain at the station. The edge node can retain its image buffer and quality events while wider systems are unavailable. Later synchronisation should carry

structured records and selected evidence rather than make basic inspection dependent on a remote service. The engineering detail belongs to the edge architecture; the shift-one principle is local continuity of the disposition.

Human review is not a temporary embarrassment in this design. It is the control that handles ambiguous cases and produces trustworthy local labels. Over time, verified examples can support recalibration or fine-tuning. The factory then improves from its own production evidence without pretending that it possessed that evidence before launch.

1.7

The paper tests vision-first as an engineering and economic hypothesis

The argument proceeds through four tests. Data independence asks whether the application can produce a valid local decision without multi-period factory history. Decision proximity asks whether its output acts directly on material or work in progress. Economic mechanism asks which cost is avoided, when it is avoided and who confirms the disposition. Learning curve asks whether each shift creates evidence that improves the next decision. Vision passes these tests conditionally; forecasting and generative design fail the first test at greenfield launch.

These tests also prevent a common sequencing error: purchasing a broad analytics platform before the factory can act on its outputs. A platform may aggregate data efficiently, but aggregation is not an operating decision. The first digital capability should expose a discrepancy that a named role can resolve at a physical boundary. Once that loop is stable, its structured records become dependable inputs for forecasting, optimisation and design rather than another stream of unverified dashboard activity.

The remaining chapters examine the conditions separately. Manual inspection and defect taxonomy establish what must be found. Optics and model architecture determine whether the station can see and classify it. Edge design, procurement and unit economics test whether it can run and pay back. Integration, workforce, traceability and deployment chapters specify how the signal becomes a managed operating capability. None of those later subjects should be used to smuggle an unproven local result into this opening thesis.

The evidence method is correspondingly strict. Shared register entries provide sector context and inherited programme findings. Chapter-local sources support particular research claims. Derived values expose their inputs and operations. Constructed operating examples are labelled illustrative and do not become forecasts through repetition. Where sources use different metrics or disagree, the chapter preserves the difference rather than averaging it away.

Several evidence gaps remain material. Actual product mix, buyer specifications, local defect prevalence, inter-inspector agreement, image quality, review queues, supplier prices, escape costs and operating results are unavailable. The chapter therefore does not claim a realised payback, a validated Cameroonian accuracy or a completed GFT vision deployment. Those questions require local specimens, independent labels, site acceptance tests and dated commercial quotations.

The appropriate investment decision is staged. Approve a bounded inspection cell and its measurement protocol. Test the manual counterfactual and candidate system on the same units. Price each disposition route. Expand only when the recorded net benefit, workforce load and buyer acceptance support expansion. This is how a shift-one application can become profitable without assuming profitability in advance.

Nothing in this chapter constitutes investment, legal or tax advice. Trade treatment, buyer compliance and procurement terms require current specialist confirmation before commitment. The chapter's contribution is a deployment sequence and an evidence standard for that sequence.

Start with the defect you can see, and let each verified disposition earn the next layer of intelligence.

CHAPTER SOURCE NOTES

R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R02 Fibre2Fashion; Business in Cameroon. Publication date unrecorded; inherited source-text check.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S36 World Bank. 2026-10-08.

S04 arXiv / IEEE CVPR. 2021-06-15.

S05 arXiv / ICPR. 2020-11-17.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S01 Sensors (MDPI). 2025-04-03.

S18 International Journal of Electrical and Electronics Research. 2024-02-05.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

R30 Design Journal (Taylor and Francis). Publication date unrecorded; inherited source-text check.

R31 arXiv 2509.04705. Publication date unrecorded; inherited source-text check.

S31 arXiv. 2025-09-04.

R18 Congressional Research Service IF10149; AGOA (Wikipedia). Publication date unrecorded; inherited source-text check.

R24 Mordor Intelligence. Publication date unrecorded; inherited source-text check.

R11 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S19 Textile Research Journal. 2026-09-06.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

R28 DataIntelo. Publication date unrecorded; inherited source-text check.

S07 arXiv. 2026-08-12.

02

Human visual inspection fatigues under line speed and glare, allowing up to 30 percent of defects to escape downstream

Understanding why human eyes miss defects on moving fabric explains why adding more manual inspectors cannot solve a factory's quality problem.

CHAPTER 02
DECISION TEST

Understanding why human eyes miss defects on moving fabric explains why adding more manual inspectors cannot solve a factory's quality problem.

Manual inspection cannot provide a reliable quality floor merely by covering every roll or garment. Coverage describes how much product passes in front of an inspector. Capture describes how many of its defects that inspector actually finds. Textile literature reports manual accuracy of 60 to 75 percent, while the inherited programme register records a 60 to 70 percent range and defect losses of 20 to 30 percent under operating conditions. [S03] [S08] [R29] These are reported ranges from different sources, not a universal law and not exact complements that should be forced into one equation. They nevertheless support one operational conclusion: a manual gate can process all the cloth and still let a material share of its flaws reach cutting, sewing or packing.

The causes sit at four levels. Visual search becomes less reliable when a repetitive web moves quickly, defects are rare and attention has been sustained without a break. Lighting can improve visibility, but excessive brightness or glare can defeat the same task. Human decisions also vary around borderline defects, so two inspectors need not assign the same class or severity. Finally, four-point roll grading and Acceptable Quality Limit sampling convert imperfect observations into decisions with different blind spots. [S0205] [S0204] [S0202] [S0203] Adding inspectors may increase throughput, but it does not remove those mechanisms.

For Green Future Tech Limited's Cameroon textile programme, prepared 31 January 2024 with sources refreshed in 2026, the correct response is therefore not to discard manual measurement. It is to measure its error before computer vision is introduced. The manual baseline must record speed, time on task, lighting, defect class, inspector agreement and downstream escapes. That baseline can then be compared honestly with an automated gate whose inherited benchmark is above 90 percent detection at web speeds up to 60 metres a minute. [R29] The comparison remains a design target to verify on local cloth, not a promised factory result.

2.1

Visual vigilance decays within the first thirty minutes of continuous roll watching

Visual inspection is not one act of looking. The inspection model reviewed by Sandia National Laboratories separates search from decision. The inspector first searches for a possible signal, then compares it with a standard and responds. Most information during search is gathered during fixations, which average about 300 milliseconds in the reviewed literature. [S0205] On a moving fabric web, the next section of cloth arrives whether or not the inspector has completed that search and decision sequence. Pacing therefore removes time from both stages.

The repetitive structure of woven and knitted cloth makes the task harder, not easier. A regular background offers few natural places to stop, while the signal may be a small break in that pattern. The Sandia review reports that vigilance losses tend to be worse for rare events, difficult detection, repetitive tasks and situations without performance feedback. [S0205] Those conditions describe a long run of nominally acceptable fabric interrupted by an occasional fault. Training can improve search strategy and knowledge of the defect catalogue, but it cannot make the incoming signal frequent or stop the web.

2.2 FATIGUE CHANGES THE OBSERVATION, NOT JUST THE OPERATOR'S COMFORT

Textile-specific evidence points in the same direction. A comprehensive fabric-inspection review identifies training burden, ignored small defects, human effort and difficulty maintaining focus for more than ten minutes as drawbacks of manual inspection. [S03] A knitted-fabric study separately describes eye fatigue from continuous attention to an illuminated inspection board and notes that inspection quality changes with working time, experience and personal viewpoint. [S0203] These findings do not justify a single fatigue curve for every mill. They do invalidate the assumption that an inspector's detection rate is constant across a shift.

The wider inspection literature provides a time scale and a warning against false precision. Sandia reports that the decline in vigilance performance is typically complete 20 to 35 minutes into a session, with at least half of the final loss occurring in the initial 15 minutes. It also records inspection studies in which detection deteriorated by up to 40 percent in 30 minutes, while stressing that performance depends on the task and its conditions. [S0205] Those figures come from the broader inspection literature, not a calibration of Cameroonian cotton lines. They support short inspection blocks and local measurement, not a universal claim that every fabric inspector fails at the twentieth minute.

A practical rotation follows from that evidence. Sandia's principal recommendation is to limit continuous inspection to 30 minutes, with a rest break or a change of task. [S0205] A factory can test that pattern without delaying production by rotating inspectors between roll watching, defect recording and exception review. The test should compare capture by time block on the same fabric family and defect mix. A simpler rota is useful only if the measured capture rate improves.

2.3 GLARE AND LINE SPEED NARROW THE USABLE SEARCH WINDOW

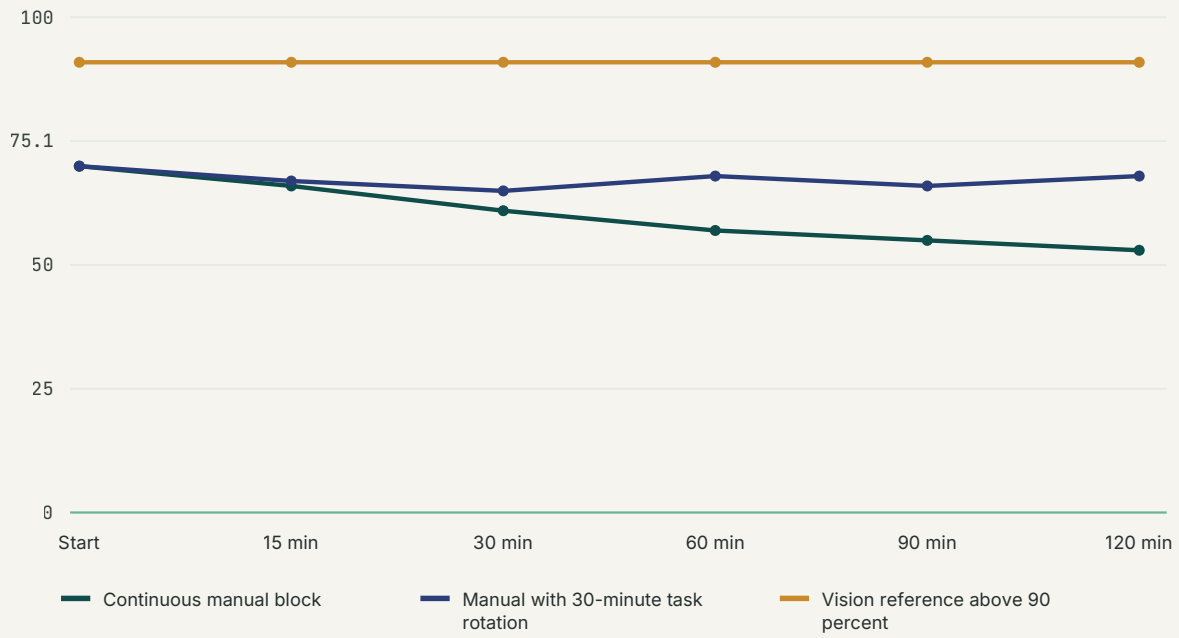
Lighting has a similarly conditional effect. The Sandia review identifies brightness, glare, colour and focus as the main lighting design variables. Increasing illumination can improve vision only up to a point, after which added illumination can create glare. [S0205] Manual fabric inspection commonly requires the operator to scan cloth transported across a brightly illuminated plane. [S0204] More backlight is therefore not automatically better inspection. The engineering question is whether each target defect retains sufficient contrast against its local weave, colour and surface relief.

Speed creates the second constraint. If an inspector has less time per unit area, the probability that a defect was never fixated rises. Sandia specifically identifies overly fast pacing as a cause of omissions because there is insufficient time for thorough inspection. [S0205] This is why a higher nominal roll throughput can produce a lower effective inspection coverage even when every metre passes the perch. The relevant measure is not metres presented per minute. It is the share of target defects found per metre, by class and contrast.

Exhibit 5, p. 32 makes the distinction visible without pretending to offer measured factory curves. The manual values are an illustrative scenario anchored to the reported textile capture range. The break rotation is a test design, not a guaranteed recovery pattern. The vision line marks the inherited lower-bound benchmark above 90 percent and does not represent a locally validated result. [R29] [S0205]

EXHIBIT 05 / ILLUSTRATIVE PERCENT OF TRUE DEFECTS CAPTURED

Continuous inspection loses capture while task rotation interrupts the decline



ILLUSTRATIVE [R29] [S03] [S0205] Illustrative scenario, not observed factory data. Continuous manual capture starts at 70 percent, within the inherited 60 to 70 percent range (R29), then declines over two hours. The rotation series resets after each 30-minute task change to test Sandia's recommended maximum continuous inspection block (S0205). Textile evidence also reports difficulty sustaining focus beyond ten minutes (S03). The 91 percent vision reference represents the inherited above-90-percent benchmark, not a measurement or guaranteed local result (R29).

The exhibit cannot establish the slope at a particular factory. It shows why the baseline protocol must record time on task rather than pooling the whole shift into one percentage. It also shows why automation should be judged by sustained capture under the same optical and production conditions, not by a short demonstration on selected defects.

2.4

Inter-rater subjectivity turns one physical defect into several data records

Finding a flaw and grading it are separate opportunities for variation. Signal detection theory, as summarised in the Sandia review, distinguishes perceptual sensitivity from response criterion. Sensitivity concerns the ability to separate signal from background noise. Response criterion concerns how much evidence an inspector requires before rejecting an item. [S0205] Two people can therefore see the same mark and make different decisions without either being careless.

That distinction matters on a production floor under conflicting pressures. A conservative inspector who rejects only clear defects will record fewer hits and fewer false alarms. A more lenient inspector will catch more borderline defects but also flag more acceptable cloth. The reviewed model predicts this movement between misses and false alarms as the decision threshold changes. [S0205] A single accuracy percentage hides which side of that trade-off a team has chosen.

2.5 BORDERLINE CALLS NEED PHYSICAL REFERENCES RATHER THAN LOUDER INSTRUCTIONS

Textile research explicitly describes manual inspection as subjective and inconsistent in high-speed manufacturing. [S01] Seam puckering provides a particularly clear example. The established grading method described in the research literature remains visual and subjective, while an automated alternative locates the seamline and extracts profile features before classifying the grade. [S21] The issue is not simply whether a wrinkle exists. It is where acceptable surface variation ends and a severity grade begins.

Similar ambiguity arises when a slub is close to a four-point length bracket or a stain is close to the background shade. A measurement taken after stopping and isolating the defect may differ from a judgement made on a moving web. The four-point method itself assigns points by fault length, so a boundary error changes the roll score even when the inspector detected the flaw. [S0203] Detection recall and grading agreement must therefore be measured separately.

A usable manual log needs three reference devices: a defect dictionary with photographs, boundary samples for each severity class, and a measured rule for length or position. The samples should include accepted as well as rejected cases. Otherwise, inspectors learn the centre of each category but not its edges. Periodic paired review then measures agreement on the same physical items, rather than comparing different rolls and calling the difference inter-rater variation.

2.6 NOISY LABELS WEAKEN MANAGEMENT DECISIONS AND LATER MODEL TRAINING

If one shift records a borderline pucker as a major defect, another records it as minor and a third accepts it, the quality log contains three decisions but not three physical truths. The resulting Pareto chart can misdirect the next review towards whichever classification habit is most common. This is a measurement-system problem before it is a production problem. The fabric review identifies inconsistency and subjectivity as manual-inspection limitations, and the seam-puckering study was designed specifically to improve grading consistency. [S01] [S21]

The same noise also affects any later vision model trained from manual labels. A model cannot learn a stable boundary from labels that change with the observer unless those disagreements are retained and adjudicated. GFT's practical requirement is to preserve the image or physical sample reference, the original inspector decision and the reviewer decision. The reviewer should not silently overwrite the first record. Agreement, escalation and final disposition are different fields.

> **MEASURES THAT SEPARATE SEEING, GRADING AND RELEASING**
>
> Use four distinct denominators in the manual baseline. **Detection recall** is defects found divided by defects confirmed by a subsequent reference review. **False-alarm rate** is acceptable items flagged divided by acceptable items inspected. **Grading agreement** is matching severity decisions divided by items graded by both inspectors. **Escape rate** is defective items released divided by defective items presented. These measures answer different questions and should not be averaged into one score. The framework follows the distinction between search, decision, misses and false alarms in the inspection literature. [S0205]

2.7

The ASTM D5430 four-point system normalises roll scores but cannot correct missed flaws

The four-point system gives a fabric buyer and seller a compact way to describe roll quality. The knitted-fabric study used here describes it under ASTM D5430-93 and records its use in quality control, fabric purchasing and buyer-seller agreements. [S0203] Its strength is arithmetic consistency. Its weakness is that the calculation can only process defects that were detected and measured.

2.8 PENALTY BRACKETS CONVERT LENGTH INTO POINTS, NOT INTO A COMPLETE QUALITY JUDGEMENT

The penalty schedule in the opened study assigns one point for a fault of less than 3 inches, two points for 3 to 6 inches, three points for 6 to 9 inches and four points for more than 9 inches. Holes and setups receive four points. [S0203] In millimetres, the corresponding boundaries are 76.2, 152.4 and 228.6 millimetres, derived using 25.4 millimetres per inch. These are length brackets, not a full assessment of colour, handle, dimensions or garment performance.

After the inspector sums the points, the score is normalised to 100 square yards. Using the formula presented in the study, `P_100 = (P_observed x 100 x 36) / (L_yd x W_in)`, where `P_observed` is total observed penalty points, `L_yd` is roll length in linear yards and `W_in` is cuttable width in inches. [S0203] The factor of 36 converts width in inches to yards. The denominator, `L_yd x W_in / 36`, is roll area in square yards.

Acceptance must then be tied to a stated specification. The opened knitted-fabric study uses Grade A for 1 to 20 points per 100 square yards, Grade B for 21 to 28 and not accepted above 28. [S0203] A factory or buyer may use another threshold, but that threshold is a contractual choice, not a universal constant. This distinction matters because the commonly quoted 40-point rule cannot be treated as though it were the score table of every fabric programme.

2.9 MANUAL UNDER-COUNTING CREATES A FALSE-ACCEPTANCE BAND

Let `r_points` be the share of true penalty points that reaches the tally sheet. This is an illustrative modelling variable, not the same measure as the percentage of individual defects detected. A long defect can carry more points than a short one, so defect recall and point recovery need not be equal. Under a proportional under-counting assumption, $P_{100,observed} = r_{points} \times P_{100,true}$. If the specified rejection threshold is T, a truly rejectable roll can be accepted when $P_{100,true} > T$ but $r_{points} \times P_{100,true} \leqslant T$.

In this illustrative model, the false-acceptance band is therefore $T < P_{100,true} \leqslant T / r_{points}$. Lower point recovery widens that band. The model is deliberately simple: it assumes that missed observations reduce the total proportionally and that no false alarms add points. A real audit should record the actual defects missed, their lengths and any points added incorrectly. Those records can replace the assumption with a measured roll-level error distribution.

WORKED EXAMPLE: A ROLL FAILS THE TRUE SPECIFICATION BUT PASSES THE MANUAL TALLY

Take an illustrative roll 120 yards long and 60 inches wide. Its area is `120 x 60 / 36 = 200` square yards. Suppose its confirmed defects carry 96 true penalty points. Its true normalised score is therefore `96 x 100 / 200 = 48.0` points per 100 square yards. The roll fails an illustrative buyer threshold of 40 points.

Now suppose the manual tally captures only 65 percent of the true penalty points in this constructed case. It records `96 x 0.65 = 62.4` points, which normalises to `62.4 x 100 / 200 = 31.2` points per 100 square yards. The roll passes the illustrative 40-point gate. At 80 percent point recovery, the observed score is 38.4 and it still passes. At 90 percent, the observed score is 43.2 and the roll is rejected. These percentages are scenario inputs used to expose the threshold mechanism, not measured performance at a GFT site. The brackets and normalisation formula follow the opened four-point study. [S0203]

The example shows why improving arithmetic cannot rescue incomplete observation. It also shows why a roll close to the threshold deserves a second inspection even when the first score says pass. A camera-based implementation can calibrate pixels to inches and assign four-point scores from detected fault dimensions, as demonstrated in the knitted-fabric study. [S0203] Yet automation still requires validation of bounding-box dimensions and missed defects. A digital score is not automatically a true score.

A second limitation is spatial. A total penalty score says nothing about where the defects lie. If the cutting room receives only a total, it cannot place pattern pieces around a local hole or stain without finding it again. The automated four-point prototype links calibrated image coordinates to fault dimensions and scoring. [S0203] That connection between score and location is operationally more useful than a higher total by itself.

2.10

ISO 2859-1 AQL sampling accepts lots after production rather than controlling defects as they form

Acceptable Quality Limit sampling answers a different question from inline inspection. It estimates whether a lot is acceptable from a sample, under a buyer-selected sampling plan. The garment study used here defines AQL as the maximum percent defective, or defects per hundred units, that can be considered satisfactory as a process average for sampling inspection. [S0202] It does not say that every lot at that percentage will pass, nor that the sample percentage is a permit to ship the same percentage of defective garments.

The distinction between acceptance number and AQL is essential. For code letter K in the study's normal, single-sampling example, the sample size is 125. At AQL 1.5, the acceptance number is 5 and rejection number 6. At AQL 2.5, they are 7 and 8. [S0202] Those integers are sample outcomes. AQL 2.5 does not mean that 2.5 defective items are allowed in every sample, and it does not mean that a lot containing 2.5 percent defectives is certain to be accepted.

2.11 A SMALL SAMPLE HAS A WIDE DECISION BOUNDARY EVEN WHEN IT IS DRAWN CORRECTLY

For an illustrative 12,000-piece lot, a sample of 125 garments covers `125 / 12,000 x 100 = 1.04 percent` of the lot. This is an arithmetic coverage calculation using the sample size from the opened study. [S0202] The remaining garments have not been visually checked by that audit. Random selection is what makes the sample decision meaningful. Selecting convenient cartons from the top of a pallet or from one production period weakens that basis.

Clustered production raises two different problems. First, a true random sample from the whole lot can still miss or under-represent a localised burst because sampling variability remains. Second, an informal sample drawn from only some cartons or time periods can be systematically blind to the burst. The remedy for the second problem is better lot definition and stratified random selection across production time, lines and cartons. The remedy for the first is an appropriate sampling plan and, where the risk requires it, inline screening.

The timing problem remains even with a perfectly random sample. A final audit happens after the sampled garments have been produced and usually after they have been packed. A drifting needle, tension fault or print problem can therefore continue while the lot is being assembled. Acceptance sampling can trigger rejection and containment, but it cannot rewind the process to the first defective piece. The garment AQL study itself presents inspection across finished fabric, cutting, sewing, finishing, pre-final and final stages rather than relying on final inspection alone. [S0202]

2.12 REJECTION AND ESCAPE ARE BOTH COSTLY OUTCOMES OF LATE INFORMATION

If the audit misses the burst, defective garments can ship. If it finds enough defects to reject the lot, the factory must contain, re-inspect or rework product after value has been added. Neither outcome is evidence that sampling has no value. It means AQL is a release control, not a substitute for process control. The buyer's contractual audit and the factory's inline prevention system should coexist.

Exhibit 6, p. 37 places the two manual protocols beside automated screening. It deliberately keeps sourced rules separate from illustrative factory thresholds and derived coverage. The purpose is to show which decision each method can support.

EXHIBIT 06 / PROTOCOL RULES AND DERIVED COVERAGE

Roll grading and lot sampling make delayed decisions from incomplete observation

Method and decision	Rule used here	What remains unresolved
Four-point length brackets	1 point below 3 in; 2 at 3 to 6 in; 3 at 6 to 9 in; 4 above 9 in or for holes	Undetected or mismeasured faults never enter the score
Opened knit-study roll grades	A at 1 to 20; B at 21 to 28; not accepted above 28 points per 100 sq yd	Study-specific acceptance, not every buyer's contract
Illustrative 40-point roll gate	True score 48.0 becomes 31.2 at 65 percent point recovery	False acceptance when observed points fall below threshold
AQL 1.5 and 2.5, code K	Sample 125; accept at 5 or 7; reject at 6 or 8 respectively	Lot release after production, not immediate station correction
Illustrative AQL lot coverage	125 pieces from 12,000 equals 1.04 percent	Unsampled pieces and clustering between sample points
Automated fabric and garment screening	Above 90 percent vision benchmark; up to 60 m/min; 0.45 s/image garment result	Local recall, optics, false alarms and physical handling still need validation

DERIVED [R29] [S19] [S0202] [S0203] Four-point brackets, study acceptance grades and formula come from S0203. The 40-point row is an illustrative buyer specification used in the chapter case, not a universal ASTM threshold. AQL code K sample and acceptance numbers come from S0202. Coverage of 1.04 percent is derived as 125 divided by 12,000 times 100 for an illustrative lot. Vision benchmarks above 90 percent and up to 60 metres a minute come from R29. The 0.45-second garment image-processing result comes from S19 and excludes physical handling.

The table should not be read as a claim that 100 percent automated inspection eliminates error. The inherited register itself reports automated detection above 90 percent, not perfection. [R29] The improvement is continuous screening with structured exception records. AQL can then perform its proper role as an independent lot-release check rather than carrying the whole quality burden.

2.13

A flaw missed at fabric entry can cost ten times more at packing in an explicit cost model

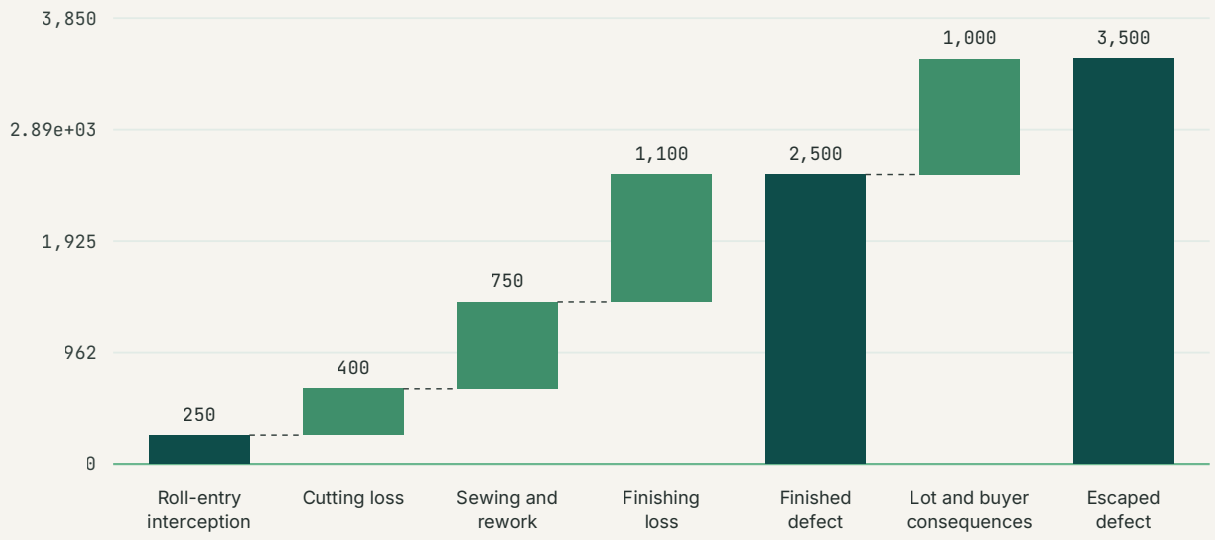
The cost of a defect depends on where it is stopped. The knitted-fabric study explains that a fault passing inspection can be dyed and then rejected after finishing, consuming time, chemicals, dyestuff and utilities without a recoverable product. [S0203] The same logic continues through cutting and garment assembly. Every completed operation adds material, machine time, labour or logistics to the flawed item. Early detection avoids those later inputs.

The financial comparison between error types is also asymmetric. The fabric-defect study explicitly treats undetected defects, or false negatives, as more costly than falsely rejecting acceptable fabric, and builds false-negative reduction into its proposed system. [S08] A false positive creates recheck work and possible unnecessary rejection. A false negative can create a finished defective garment or a customer escape. This does not make false alarms free. It means the inspection threshold should reflect both costs rather than maximising reported accuracy alone.

Exhibit 7, p. 38 converts that mechanism into a transparent illustrative case. It starts with an assumed XAF 250 cost to isolate and splice around a fault at roll entry. Cutting, sewing and finishing then add assumed losses until the finished defect costs XAF 2,500, ten times the entry cost. A further assumed XAF 1,000 for lot-level handling and commercial consequences brings the escaped case to XAF 3,500. These are scenario values, not observed Cameroon costs or a general tenfold law.

EXHIBIT 07 / ILLUSTRATIVE CFA FRANCS (XAF) LOST PER DEFECT

Late detection turns a small fabric fault into accumulated garment loss



ILLUSTRATIVE [S08] [S0203] Illustrative model, not a cost benchmark. Assumed entry interception cost is XAF 250. Assumed additions are XAF 400 after cutting, XAF 750 after sewing and XAF 1,100 after finishing, producing XAF 2,500 or ten times entry cost. A further illustrative XAF 1,000 lot-handling and commercial allocation produces XAF 3,500. S0203 supports the mechanism of unrecoverable processing losses when faults pass early inspection. S08 supports the higher cost assigned to false negatives. No GFT operating cost is claimed.

The model should be replaced with local accounts before any investment decision. Fabric consumed, replacement-panel labour, unpicking time, restitching, pressing energy, downgrade value, reinspection and buyer charges each need a measured basis. Some defects are repairable, while others destroy the garment. Some buyer consequences fall on a whole lot rather than one item. Allocating that lot total to one escaped defect requires an explicit denominator. The exhibit is a method prompt, not a payback forecast.

2.14

Manual cycle time becomes a throughput constraint before it becomes a labour-saving case

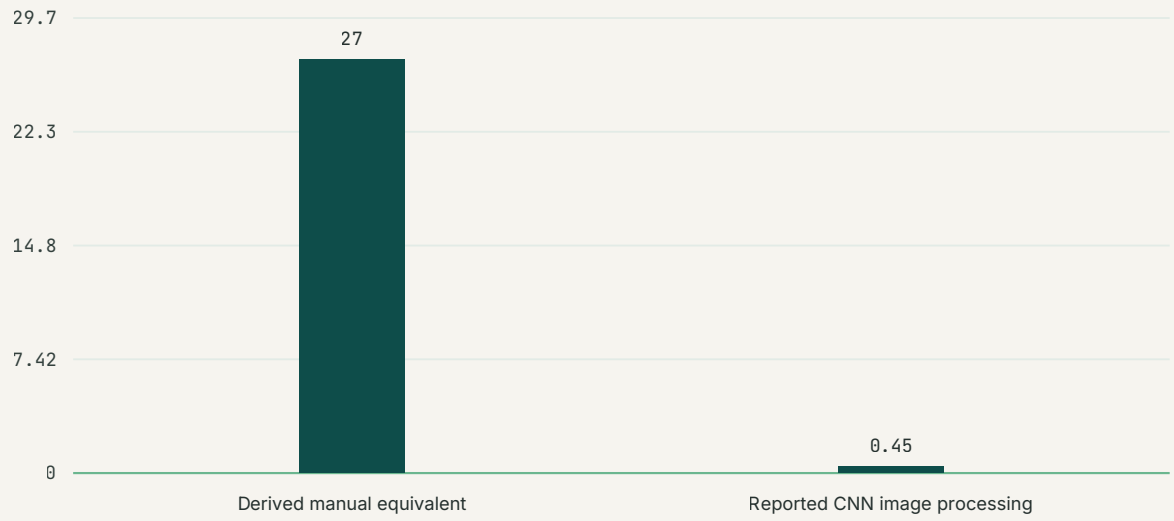
Manual inspection consumes time even when it finds the right defects. The fabric review lists time consumption among its main drawbacks, while the Ethiopian garment study describes manual visual inspection as subjective, fatigue-prone and a limit on throughput. [S03] [S19] A queue at final inspection also delays information. By the time the tally reaches the sewing supervisor, the line may have repeated the same fault.

The Ethiopian study provides a useful but bounded benchmark. Its VGGNet-19 and artificial-neural-network system used a dataset of 1,000 trouser images collected in Addis Ababa garment facilities under factory lighting. It achieved 86.7 percent overall classification accuracy and processed each image in 0.45 seconds on consumer-grade hardware, reported as a 60-fold improvement over manual inspection throughput. [S19] The result concerns image classification, not complete physical handling, measurement or buyer acceptance.

Exhibit 8, p. 39 derives the manual equivalent from that reported ratio: `0.45 x 60 = 27.0 seconds`. It is a comparison of processing time, not proof that a complete automated station handles sixty times as many garments. [S19]

EXHIBIT 08 / SECONDS PER GARMENT SAMPLE

Image processing takes 0.45 seconds against a derived 27-second manual equivalent



DERIVED [S19] Derived from the opened Ethiopian garment study (S19): reported automated processing is 0.45 seconds per sample and 60-fold manual-equivalent throughput. Manual equivalent is 0.45 times 60 equals 27.0 seconds. The comparison excludes garment placement, turning, dimensional checks, exception handling and packing, so it is not a complete station cycle-time claim.

A line-balancing example shows the denominator. In the paper's illustrative model of 3.0 million garments a year over 300 working days, daily output is 10,000 garments. At 27.0 seconds each, manual inspection requires 270,000 inspector-seconds a day. One eight-hour shift contains 28,800 seconds, so the theoretical requirement is `270,000 / 28,800 = 9.375` full-time equivalents before breaks, fatigue, handling variation or re-inspection. All inputs in this paragraph are illustrative planning assumptions, not a GFT staffing plan.

In the same illustrative planning case, 0.45 seconds per image requires 4,500 seconds of inference for the daily volume, or 75 minutes of compute time across the available devices. That arithmetic does not justify a one-person inspection station. Garments must still be presented consistently, exceptions reviewed and defective pieces removed. The correct design use of **Exhibit 8**, p. 39 is to size image-processing capacity and identify where manual queues form. Station labour requires a separate measured handling study.

2.15

The manual baseline must measure its own blindness before vision is judged against it

A before-and-after comparison is valid only if both periods use compatible definitions. GFT's manual baseline should retain roll and bundle identifiers, fabric family, inspection speed, lighting condition, time on task, defect class, severity, inspector decision and downstream disposition. These are recommended data fields for the Cameroon programme, not claims about an already operating factory. The objective is to distinguish changes in production quality from changes in measurement sensitivity.

The reference review also needs honest limits. A slower second inspection by a senior reviewer is a practical reference, not infallible ground truth. Where possible, disputed items should receive a panel decision or a measurement using a ruler, calibrated image or colour reference. The manual detection estimate should then be reported by defect class. A single pooled recall can hide poor performance on exactly the small or low-contrast defects that matter most.

GFT NOTE: RUN THE BASELINE AS A MEASUREMENT-SYSTEM AUDIT

Before accepting the vision benchmark, sample across fabric families, production lines and time blocks. Pair the ordinary inspector's decisions with a slower reference review, retain disagreement rather than overwriting it, and record downstream escapes. Test a 30-minute task rotation against the existing rota, following the vigilance intervention reviewed by Sandia. [S0205] Report detection, false alarms, grading agreement and escape rate separately. Use the same defect dictionary and denominators in the manual and vision periods so the comparison remains auditable.

A HIGHER DEFECT COUNT AFTER AUTOMATION CAN MEAN BETTER DETECTION

A camera that finds defects previously missed will increase recorded Defects Per Hundred Units even if physical production has not changed. Managers should not interpret that initial rise as a sudden process failure or penalise lines for becoming more visible. Compare confirmed downstream escapes, rework and Right First Time alongside detections. The manual reference audit supplies the context needed to separate improved capture from a real deterioration in production. [S03] [S08]

The baseline is therefore both a quality measure and a warning about the measure itself. Manual logs remain useful for management, provided their uncertainty is visible. Vision should then be introduced as a measured change in search coverage and

decision consistency, with humans still reviewing exceptions and ambiguous grades. The next chapter defines the defect classes against which that comparison must be run.

Do not ask tired eyes to prove what the line never gave them time to see.

CHAPTER SOURCE NOTES

S03 Mathematical Problems in Engineering (Wiley). 2020-11-16.

S08 International Journal for Research in Applied Science and Engineering Technology. 2023-05-31.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S0205 Sandia National Laboratories, United States Department of Energy. 2012-10-01.

S0204 TEKSTIL VE KONFEKSIYON. 2019-03-26.

S0202 Journal of Textile Science & Fashion Technology. 2019-05-16.

S0203 Heliyon (Elsevier). 2024-08-24.

S01 Sensors (MDPI). 2025-04-03.

S21 Textile Research Journal (SAGE Publications). 2016-10-14.

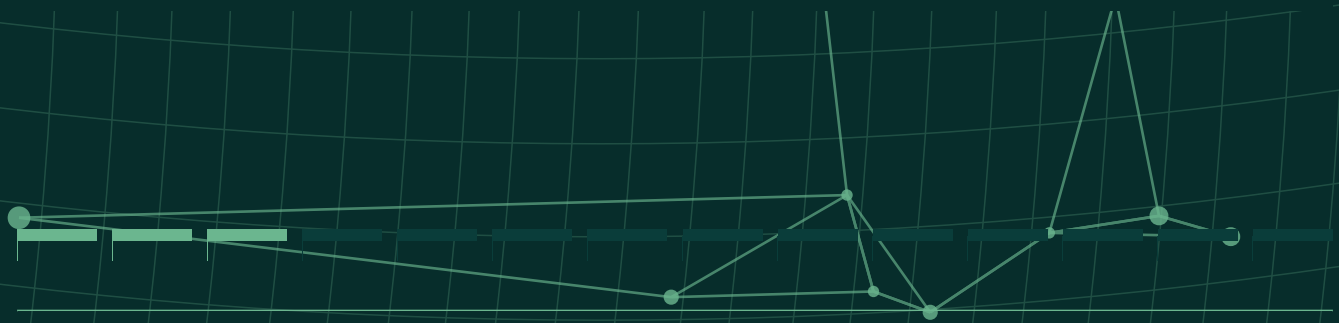
S19 Textile Research Journal (SAGE Publications). 2026-09-06.

03

THE QUALITY GAP AND THE CASE FOR VISION FIRST / DEFECT TAXONOMY

Textile and garment defects span more than 40 distinct classes, each with its own optical signature across warp, weft and seam

A useful inspection dictionary describes what the camera can see, what the inspector must confirm and which production decision follows.



CHAPTER 03
DECISION TEST

A useful inspection dictionary describes what the camera can see, what the inspector must confirm and which production decision follows.

A factory should commission its inspection system against a defect dictionary, rather than an undifferentiated promise to detect faulty cloth. This chapter proposes an illustrative catalogue of 42 operational classes across yarn preparation, weaving, knitting, wet processing, cutting and sewing. The grouping is GFT's design proposal, not a universal classification or evidence that any detector recognises every class. Optical signatures are inspection hypotheses to test on the intended fabric. Several classes can look alike, and the same class can change appearance with fabric, lighting and presentation. Keep observed appearance, suspected process origin, severity and release decision as separate fields.

Published work provides bounded examples, not a complete factory specification. Deep-GD includes a defect-free category in its reported class list. That nominal category is a reference condition, not another defect. The airbag seam study addresses particular stitch configurations, while the warp-knitting study addresses broken-end inspection. Their vocabularies help define the inspection task; their results do not establish interchangeable coverage across products. [S18] [S22] [S15]

The catalogue below is deliberately operational. Each code should become a controlled definition, an image collection, a disagreement rule and an escalation route before automatic release is considered. The proposed visual checks and routing rules throughout this chapter are GFT engineering recommendations. They are not measured detection results. The chapter establishes what must be distinguished; later chapters address optics, models, garment presentation and corrective-action workflows.

3.1

Yarn irregularities need an identity before they become fabric complaints

Start the dictionary upstream of the camera's installation point. A mark first seen on incoming cloth need not have originated in weaving. Conversely, an irregular yarn appearance is not sufficient evidence that a spinner supplied defective yarn. Record the observation as an appearance class until a retained sample or process investigation supports attribution. This distinction is especially useful when a factory buys finished cloth rather than controlling yarn preparation itself. Its quality team can document a reproducible complaint without claiming access to the supplier's machinery records.

For the illustrative yarn family, retain separate entries for nep-like specks, local slubs, extended thick places, thin places, excessive hairiness, foreign-fibre inclusions and prominent knots. Their proposed boundaries distinguish compact particles, short bulges, sustained width changes, filament-like fringes and attached junctions. These are working annotation distinctions. They must be reconciled with the purchased yarn and fabric specification before production use. Deep-GD's use of hairiness and lumpy appearances supports distinguishing texture abnormalities from stains, but does not validate this complete family. [S18]

The proposed nep-like class should request a close view of an attached compact fleck and a wider view of its surrounding texture. Keep loose surface debris in a separate review state until attachment is checked. For the slub class, request the visible

bulge and enough neighbouring yarn to identify its direction. For thick and thin places, request a longer contextual crop than for an isolated particle. This keeps the annotation tied to a spatial extent rather than to an operator's impression that the cloth looks uneven.

Treat hairiness as a comparison against an approved normal sample from the same construction and finishing condition. A silhouette against a contrasting background may be a useful commissioning view, but qualification must establish whether it predicts the buyer's assessment. Likewise, the proposed foreign-fibre code should describe a visible inclusion without claiming chemical identification. A camera flag can request investigation; it should not certify fibre composition from appearance alone. For a knot, record whether the junction protrudes, affects an intended visible panel or falls within an agreed acceptable appearance.

The illustrative dictionary in **Exhibit 9**, p. 44 assigns stable yarn-family codes and candidate optical cues. It is a commissioning worksheet, not a prevalence table. An entry's presence means that the inspection designer must decide whether it applies to the product. It does not mean the factory has observed that defect, or that the proposed view can reliably separate it from every lookalike.

EXHIBIT 09 / PROPOSED OPERATIONAL CLASSES

Yarn observations require different confirmation views

Proposed class	Candidate cue	Confirmation
Y01 Nep-like speck	Compact attached fleck	Check attachment
Y02 Local slub	Short local bulge	Compare nearby yarn
Y03 Thick place	Sustained broad segment	Check contextual length
Y04 Thin place	Sustained narrow segment	Compare construction
Y05 Hairiness	Peripheral fibre fringe	Compare approved finish
Y06 Foreign inclusion	Distinct embedded feature	Confirm material identity
Y07 Prominent knot	Local junction or nub	Apply agreed sample rule

ILLUSTRATIVE Illustrative GFT dictionary. Seven entries per family, six families, 42 entries in total. Candidate cues and checks are design hypotheses requiring local validation, not measured signatures, universal terminology or fixed acceptance rules. Codes are class labels, not source identifiers.

For procurement, require the supplier and factory to agree how short local bulges are distinguished from sustained thick places. Store examples on both sides of that boundary. If the distinction cannot be applied consistently, keep a broader confirmed irregularity code and retain the finer annotation as provisional. Forcing certainty at intake creates an apparently precise supplier scorecard whose differences may reflect inspection language rather than material quality.

GFT NOTE

Approve a physical sample pack alongside the dictionary. Include accepted texture, confirmed defects, borderline examples and examples that require another test. Make the quality lead responsible for definition changes. Purchasing should use confirmed observations and agreed acceptance rules in supplier discussions, while proposed causes remain visibly provisional.

3.2

Weaving faults need fabric coordinates rather than screen coordinates

A directional definition must travel with the cloth. The ASTM terminology preview distinguishes a warp-direction void from a filling-direction discontinuity. Those material directions remain meaningful when an image is rotated. A screen label such as horizontal does not. Record warp and weft orientation in the image metadata and retain a reference showing how the fabric was presented. The preview supports these terminology distinctions; it is not a complete acceptance specification. [S0301]

The illustrative weaving family separates broken ends, missing or broken picks, double ends, double picks, reed-spacing marks, unintended floats and snarled yarn inclusions. These are proposed operational labels. The candidate distinctions are absent material, additional material, spacing irregularity, exposed unbound length and a local tangled shape. Deep-GD uses double-end, snarl and miss labels, but a factory must inspect its actual label definitions before mapping those names into the local dictionary. [S18]

For a broken-end candidate, ask whether the narrow feature follows the recorded warp direction and remains distinguishable from an intentional stripe. For a missing-pick candidate, ask the corresponding question across the filling direction. A dark line alone should not decide either class. The proposed decision record should describe continuity, neighbouring yarn structure and the inspection view that exposed the suspected gap. If those observations remain inconclusive, preserve a line-anomaly label and refer the sample for confirmation.

Double-end and double-pick candidates need a different reference question: does the local structure contain additional material relative to an approved repeat? Reed-spacing marks should ask about persistent spacing rather than local yarn thickness. Floats should request enough context to distinguish a missed interlacing appearance from the intended weave. Snarls should request a local shape description before the investigator attributes the event to any particular machine condition. These are proposed checks, rather than guaranteed optical discriminators.

The worksheet in **Exhibit 10**, p. 46 keeps these questions separate. It should be reviewed with the supplier's construction terminology. A detector may initially issue a broader directional anomaly flag while the quality lead assigns a confirmed class. That staged approach retains useful localisation without pretending that an early training set already resolves every structural distinction.

EXHIBIT 10 / PROPOSED OPERATIONAL CLASSES

Directional observations require material-axis references

Proposed class	Candidate cue	Confirmation
W01 Broken end	Warp-aligned gap	Check yarn continuity
W02 Missing/broken pick	Weft-aligned gap	Check filling continuity
W03 Double end	Extra warp-like line	Compare intended repeat
W04 Double pick	Extra weft-like line	Compare intended repeat
W05 Reed-spacing mark	Persistent spacing line	Check repeat geometry
W06 Unintended float	Exposed unbound length	Compare intended weave
W07 Snarled inclusion	Tangled local shape	Confirm structure locally

ILLUSTRATIVE Illustrative GFT dictionary. Seven entries per family, six families, 42 entries in total. Candidate cues and checks are design hypotheses requiring local validation, not measured signatures, universal terminology or fixed acceptance rules. Codes are class labels, not source identifiers.

Define the relationship between a physical defect and an image crop before annotating long lines. A continuous feature may cross several crops or camera fields. Treat those crops as evidence attached to a candidate event, then use the verified material position to decide whether they belong together. A disagreement about event continuity should remain visible in the record. Do not quietly convert a crop count into a count of defective garments or a length of rejected cloth.

For model qualification, require directional cases in both their normal presentation and plausible rotated presentations. This is a proposed evaluation condition, not an instruction to make the model invariant to every rotation. Material orientation may carry useful information. Preserve it explicitly so that a change in cloth presentation can be distinguished from a change in the process. The dictionary should help the optical and model teams state that requirement without selecting their hardware or architecture for them.

3.3

Knitted structures require loop-aware definitions and a qualified presentation

A knitting dictionary should begin with the intended structure rather than borrow woven labels indiscriminately. The smart-sensor study reports a warp-knitting application aimed at broken-end defects. Its reported 98 percent detection rate follows six months of operation on that application. This is evidence for that inspected task, not evidence that all knit faults or garment styles have the same detectability. [S15]

The illustrative knit family uses dropped-loop candidates, needle-line candidates, ladders, holes, unwanted tuck loops, sinker-line candidates and spirality or course skew. These labels separate a local interruption, a persistent line, an extending opening, an area opening, an unexpected loop appearance, a repeating surface mark and a

global alignment condition. The boundaries are proposed for annotation and require confirmation by a knitting specialist. In particular, a hole records an observed opening; it does not, by itself, identify the mechanism that created it.

For dropped-loop and ladder candidates, request both a detailed crop and a contextual image showing how far the affected structure continues. For a needle-line candidate, retain a normal comparison from the same fabric family. For unwanted tucks, require an approved pattern reference so an intentional feature is not labelled defective. Sinker-line candidates should remain provisional until the visual mark is connected to a confirmed structural or process observation. Naming a machine component in a class must not turn an image prediction into a maintenance diagnosis.

For spirality or course skew, the proposed boundary is larger than a local defect crop. Record the presentation state, reference axis and measurement method. Avoid assigning a precise alignment failure from a loosely positioned garment photograph. If the assessment requires conditioning, relaxation or a specified test procedure, mark the camera observation as a request for that assessment. The intended inspection gate should determine which presentation belongs in the dictionary.

The proposed classification in **Exhibit 11**, p. 47 is therefore accompanied by a presentation requirement. A useful commissioning sample should show the accepted state, the suspect state and any manipulation needed for confirmation. Keep an opened hole, a slack fold and a dark printed feature available as deliberate lookalikes. This gives reviewers a practical way to record what the installed view can and cannot decide.

EXHIBIT 11 / PROPOSED OPERATIONAL CLASSES

Loop and alignment candidates require qualified presentation

Proposed class	Candidate cue	Confirmation
K01 Dropped loop	Local loop interruption	Inspect surrounding loops
K02 Needle line	Persistent wale-like line	Compare normal texture
K03 Ladder	Extending loop opening	Retain full extent
K04 Hole	Area of open material	Confirm under set tension
K05 Unwanted tuck	Unexpected loop motif	Check pattern reference
K06 Sinker-line candidate	Repeated surface mark	Confirm structural origin
K07 Spirality/course skew	Global axis departure	Apply agreed test method

ILLUSTRATIVE Illustrative GFT dictionary. Seven entries per family, six families, 42 entries in total. Candidate cues and checks are design hypotheses requiring local validation, not measured signatures, universal terminology or fixed acceptance rules. Codes are class labels, not source identifiers.

Create an observability status for every inspection event. Use it to distinguish an adequate view from an obscured, folded or otherwise unsuitable view. An uninspectable area should not silently inherit an acceptable label. Route it for another presenta-

tion or an explicit manual check. Record the reason so that repeated presentation failures can be addressed independently of model errors. This is a proposed operating rule, and it should be tested with the people who will handle the material.

The class list should remain stable when presentation improves. What changes is the evidence available for assigning a class, not necessarily the underlying definition. Retaining that distinction allows the quality lead to compare historical records while the inspection team changes fixtures, handling instructions or camera recipes. If improved presentation exposes a distinction that was previously impossible to adjudicate, introduce the refined code through a documented dictionary revision rather than rewriting old records without explanation.

3.4

Wet-processing labels must separate appearance from chemical explanation

The proposed wet-processing family contains dye spots, oil-like stains, soil-like stains, water-mark candidates, shade drift, weak or faded print and print misregistration. Deep-GD explicitly distinguishes soil stain, oil stain and dye spot. Its vocabulary is useful evidence that stain classes can be separated in a research dataset. It does not demonstrate that an ordinary colour image proves the substance responsible for any new factory mark. [S18]

For that reason, prefer oil-like and soil-like while the event is awaiting confirmation. Store the visual class and the confirmed substance, if established, in different fields. A maintenance investigation may later connect a stain to a machine or handling step. That conclusion should be added with its evidence and reviewer, rather than embedded in the original detection as if the camera observed the cause. The quality log then remains useful even when the first explanation is rejected.

For dye spots and water-mark candidates, the proposed annotation should capture the region and its surrounding colour context. For shade drift, request a comparison against the approved reference under the qualified inspection recipe. Preserve exposure settings and any reference-capture procedure used for that decision. Do not give a numerical colour-difference threshold authority simply because the software can calculate it. The acceptance threshold and measurement method belong to the product agreement and require their own approval.

For weak print, define the expected motif and reference condition before asking an inspector to grade missing coverage or reduced density. For misregistration, define which pattern features must align and the permissible relationship between them. The proposed distinction is between an appearance or coverage failure and a positional failure. Keep both labels available when they coexist, while recording the material as a single affected unit for unit-based yield calculations.

The worksheet in **Exhibit 12**, p. 49 reflects these different decision boundaries. It intentionally excludes inferred chemical composition, fastness certification and bleaching adequacy from the camera's automatic conclusions. A visible mark may be

reason to request a test. It is not, without appropriate evidence, the result of that test. Keep specialist laboratory assessments linked to the event without implying that the image alone supplied them.

EXHIBIT 12 / PROPOSED OPERATIONAL CLASSES

Colour and stain labels must retain their reference condition

Proposed class	Candidate cue	Confirmation
D01 Dye spot	Local colour concentration	Compare colour reference
D02 Oil-like stain	Local tonal or sheen patch	Confirm substance separately
D03 Soil-like stain	Local mottled region	Confirm substance separately
D04 Water-mark candidate	Diffuse ring or boundary	Check retained sample
D05 Shade drift	Broad colour departure	Use qualified colour recipe
D06 Weak/faded print	Low motif coverage	Compare approved artwork
D07 Print misregistration	Displaced motif features	Check alignment reference

ILLUSTRATIVE Illustrative GFT dictionary. Seven entries per family, six families, 42 entries in total. Candidate cues and checks are design hypotheses requiring local validation, not measured signatures, universal terminology or fixed acceptance rules. Codes are class labels, not source identifiers.

A proposed commissioning exercise should include repeated captures of an unchanged approved sample as well as captures of suspect material. Use these to ask whether a changed alert follows the material or the inspection conditions. Document the result as part of recipe qualification. If the system cannot distinguish those alternatives reliably enough for the agreed decision, route the class to manual confirmation and state that limitation in its coverage record.

Keep local prevalence estimates separate from external case studies. A print shop, a dye house and a sewing line have different opportunities to observe and introduce faults. Combining their category frequencies into a synthetic factory Pareto would give a misleading denominator. The appropriate local distribution should come from the factory's own confirmed events, with its inspection boundary, product mix and coverage stated. External evidence can suggest classes to investigate; it should not fill missing local counts.

3.5

Cutting faults need a pattern reference and a retained bundle identity

For the illustrative cutting family, use panel-outline distortion, missing edge material, misplaced or excessive notches, drilling-position errors, frayed or ragged cut edges, fused plies and reversed nap or directional motif. These are proposed operational categories. Their candidate checks compare an observed panel or bundle with the approved pattern, orientation and handling requirements. They should not be represented as a classifier validated by a fabric-surface research dataset.

The training material on cutting distinguishes outline accuracy, edge condition, material faults within cut parts and shade differences between bundled pieces. It also treats fusion as a cutting concern. This supports keeping dimensional and handling checks alongside surface inspection, rather than assuming that a clean fabric image proves an acceptable cut component. [S0302]

For panel-outline distortion, request the approved pattern revision and a qualified presentation before measuring deviation. For missing edge material, identify the expected boundary and whether the relevant edge is actually visible. An apparent omission caused by an overlapping panel should remain a presentation failure until separated. For notches, record the expected position and permitted geometry. For drilling, request the intended location and whether the mark is supposed to penetrate the inspected stack or only the component presented.

For frayed edges, retain a detailed crop alongside the panel identity. For fused plies, require a handling confirmation or another qualified observation if the visible top face cannot resolve separation. For reversed nap or directional motifs, compare the component with the approved orientation rule and its paired components. These are proposed inspection tasks. No universal dimension or tolerance is supplied here, because the dictionary should carry the reference to the actual product requirement.

The worksheet in **Exhibit 13**, p. 50 also helps define the inspection unit. A top panel, a bundle and a complete garment are different units. The proposed record must say which was inspected and which material the decision covers. An alert from a visible outer component should not be treated as evidence that every concealed ply was checked. Explicit coverage makes a narrow but reliable inspection more useful than an ambiguous whole-bundle pass.

EXHIBIT 13 / PROPOSED OPERATIONAL CLASSES

Panel and bundle judgements require different evidence

Proposed class	Candidate cue	Confirmation
C01 Outline distortion	Contour departure	Use approved pattern
C02 Missing edge material	Incomplete contour	Expose full panel edge
C03 Notch error	Wrong location or depth	Compare pattern geometry
C04 Drilling error	Wrong mark position	Confirm intended penetration
C05 Frayed/ragged edge	Irregular cut boundary	Inspect detailed edge
C06 Fused plies	Failure of separation	Require handling check
C07 Reversed nap/motif	Directional mismatch	Compare paired components

ILLUSTRATIVE Illustrative GFT dictionary. Seven entries per family, six families, 42 entries in total. Candidate cues and checks are design hypotheses requiring local validation, not measured signatures, universal terminology or fixed acceptance rules. Codes are class labels, not source identifiers.

Retain the connection between roll position and cut component when it is available. A pre-existing hole discovered after cutting should remain a material observation linked to the earlier roll, rather than being assigned to cutting merely because that is where it

was noticed. Conversely, a panel-outline failure should not become a supplier material defect without evidence. Store discovery station separately from suspected origin and from confirmed origin. The same principle applies to stains first found during sewing or finishing.

For rework, keep the original event and add the attempted repair and reinspection result. A replaced panel should retain its relationship to the garment or bundle so the quality lead can identify what actually passed. Closing an event should mean the prescribed action was completed and verified, not merely that someone moved the component away from the inspection station. This proposed record structure supports the later corrective-action chapter without prescribing a particular factory software implementation.

3.6

Sewing faults need separate stitch, seam and surface judgements

The illustrative sewing family separates broken stitching, skipped stitches, open seams, seam slippage, untrimmed thread, needle damage and seam puckering. The airbag paper describes skipped, missed and superimposed seam defects for specific stitch configurations. Pan and colleagues describe image-based seamline localisation and profile analysis for puckering assessment. These studies support different inspection tasks; neither establishes universal acceptance limits for the proposed sewing family. [S22] [S21]

For broken stitching, the proposed question is whether the required thread path has a confirmed interruption. For skipped stitches, compare the observed formation with the expected stitch sequence. For an open seam, inspect the intended join rather than inferring closure from nearby decorative stitching. Seam slippage should request the agreed confirmation procedure and reference condition. If that procedure is not represented in the camera view, retain a suspect opening class until the required check is made.

For untrimmed thread, record the intended visible zone and the applicable finishing rule. For needle damage, distinguish a visible damaged area from a normal penetration appearance through agreed samples and, where necessary, closer examination. For puckering, define the evaluation presentation and reference grades before collecting training labels. An image class should not imply that the inspector has proved a particular tension, feed or shrinkage mechanism. Store any such diagnosis after the relevant investigation.

The proposed worksheet in **Exhibit 14**, p. 52 makes these boundaries inspectable. Keep construction failures, surface appearance and presentation failures separate even if they share a review queue. A garment may have both an interrupted stitch path and an uneven seam appearance. Permit multiple confirmed observations while preventing them from creating duplicate records of the same affected unit.

EXHIBIT 14 / PROPOSED OPERATIONAL CLASSES

Stitch integrity and seam appearance need separate decisions

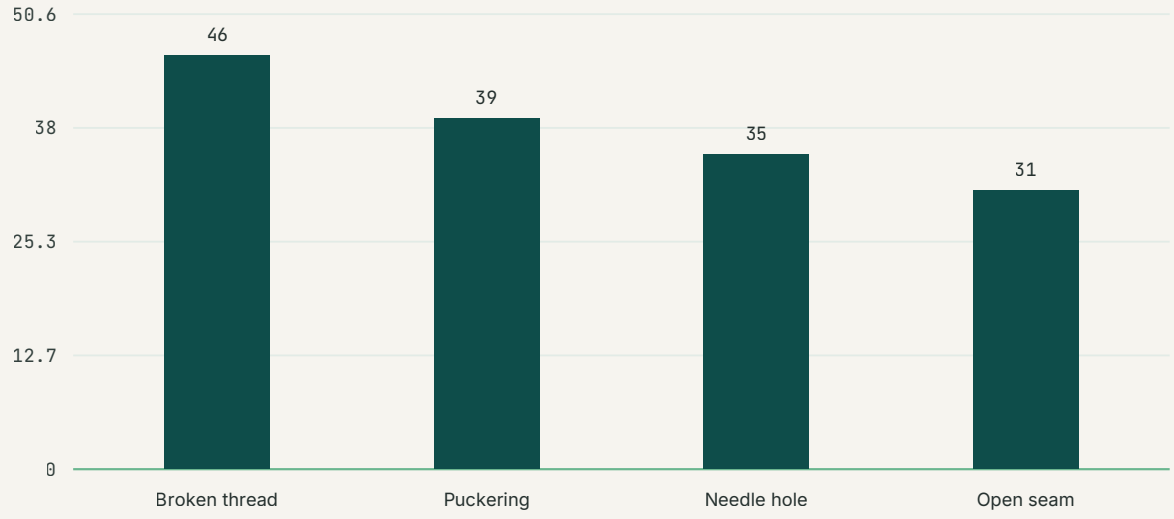
Proposed class	Candidate cue	Confirmation
S01 Broken stitching	Interrupted thread path	Confirm required stitch path
S02 Skipped stitch	Missing expected formation	Compare stitch sequence
S03 Open seam	Unjoined intended join	Inspect intended seam
S04 Seam slippage	Suspect opening at join	Apply confirmation procedure
S05 Untrimmed thread	Free projecting tail	Apply zone finishing rule
S06 Needle damage	Suspect penetration damage	Compare approved samples
S07 Seam puckering	Uneven seam appearance	Use agreed presentation

ILLUSTRATIVE Illustrative GFT dictionary. Seven entries per family, six families, 42 entries in total. Candidate cues and checks are design hypotheses requiring local validation, not measured signatures, universal terminology or fixed acceptance rules. Codes are class labels, not source identifiers.

A useful sourced comparison is available from the jacket study. Summing its daily entries gives 46 broken-thread records, 39 puckering records, 35 needle-hole records and 31 open-seam records. **Exhibit 15**, p. 52 shows only those selected categories, not a complete Pareto. The counts describe that study, not a Cameroon baseline, and should be treated as category records rather than independently verified unique defective garments. [S25]

EXHIBIT 15 / CATEGORY RECORDS OVER THE STUDY OBSERVATION PERIOD

Selected jacket records distinguish thread, pucker and opening faults



DERIVED [S25] Derived from S25 Table 3, fifteen daily values per category, printed page 239. Dashes treated as zero. Inputs: {"Broken thread": [2, 4, 2, 3, 0, 1, 3, 3, 3, 2, 1, 4, 6, 6, 6], "Puckering": [2, 3, 3, 1, 4, 4, 2, 2, 1, 3, 2, 3, 3, 3, 3], "Needle hole": [2, 3, 0, 2, 1, 0, 0, 3, 2, 6, 5, 2, 2, 3, 4], "Open seam": [2, 3, 2, 2, 3, 2, 1, 3, 2, 3, 0, 0, 3, 2, 3]}. Sum each row: 46, 39, 35, 31. Selected categories only; not a complete Pareto or unique-garment count. External study, not a GFT result.

The practical lesson is to preserve the class boundaries before ranking them. A combined sewing-defect code cannot distinguish the next investigation for an interrupted thread path from the next investigation for seam appearance. However, the relative frequency in an external jacket study should not decide the priority of a new uniform or knitwear line. Start with the product's requirements, then build a confirmed local distribution with documented inspection coverage.

Do not equate visible severity with risk without the product context. A small observed opening at a load-bearing join may require a different disposition from a conspicuous cosmetic mark in a concealed zone. The quality lead should approve the severity rule against the intended use and buyer specification. Keep the class stable while allowing severity and disposition to vary by zone and product. A defect dictionary is a language for decisions, not a substitute for the product's acceptance agreement.

3.7

The quality log must preserve uncertainty and count the right event

Implement the proposed dictionary as a versioned reference used by both the inspection application and the Performance-Management System. Each class should have its stable code, preferred name, accepted synonyms, inclusion criteria, exclusions, required view, measurement boundary and reference images. Keep the production-stage family as an organising field. Do not make it compulsory evidence of origin. Every event should retain the dictionary version under which it was first labelled.

For the proposed event record, retain material or garment identity, discovery station, timestamp, image reference, material coordinates where available, presentation state, candidate classes, confirmed class, reviewer, severity, disposition and repair status. Record model version and confidence separately from the confirmed judgement. Confidence expresses the model's output under its configured interpretation; it should not be relabelled as a physical probability of product failure without validation for that purpose.

Separate unknown appearance from uninspectable presentation. Unknown appearance means usable evidence exists but the approved dictionary does not yet resolve it. Uninspectable means the evidence is insufficient for the required check. Give both a review route. Keep confirmed acceptable examples as carefully as confirmed defects so the training record retains the intended boundary. Borderline cases should keep their adjudication notes, particularly when a new product or finish changes what counts as normal.

Exhibit 16, p. 54 is an illustrative prioritisation aid for qualification. Its ordinal scores are assumed effort categories, not measured algorithm accuracy. It asks reviewers to consider a feature against different backgrounds and with different evidence needs. A higher score requests a stronger qualification plan; it does not establish that the class is intrinsically harder under every camera or product configuration.

EXHIBIT 16 / ILLUSTRATIVE ORDINAL EFFORT SCORE

Qualification effort depends on the feature and its reference

	Plain cloth	Twill	Printed motif	Presented garment
Point anomaly	1	2	3	3
Directional line	1	2	3	3
Colour region	2	2	3	3
Seam appearance	3	3	3	2

Cell scale: 1 to 3

ILLUSTRATIVE GFT planning assumptions: 1 = simple reference task, 2 = added discrimination required, 3 = strong reference/presentation dependence. Scores are assumed priorities, not empirical difficulty, probability, accuracy or interval-scale measurements. Do not average them.

The synthetic profiles in **Exhibit 17**, p. 55 serve a narrower purpose. They show possible local contrast shapes for training discussions. They are constructed arrays, not captured images, measured grey levels or Fourier spectra. Their equal sampling positions have no assigned physical pitch. Use them to explain why a crop, a profile and a defect identity are different objects, then replace them with qualified local examples when commissioning.

EXHIBIT 17 / ILLUSTRATIVE SIGNED CONTRAST, ARBITRARY UNITS

Synthetic profiles demonstrate contrast shapes without proving class identity



ILLUSTRATIVE GFT teaching arrays, not measurements or Fourier spectra. Nine equally spaced positions with no physical pitch. Warp gap and slub are transverse cuts, stain and hole are centre cuts, stitch and pucker profiles follow the proposed seam direction. No common physical scale or universal class signature is claimed.

WORKED EXAMPLE: OBSERVATIONS AND AFFECTED GARMENTS HAVE DIFFERENT DENOMINATORS

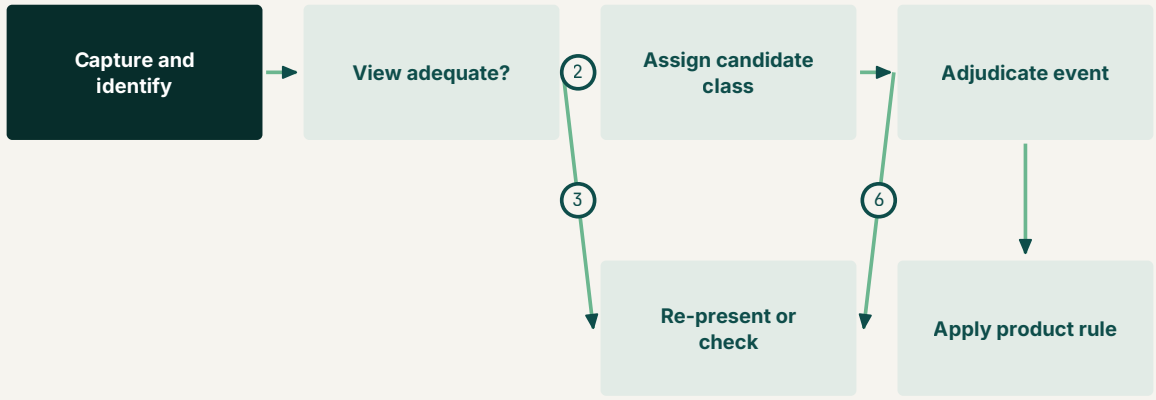
In an illustrative exercise, a reviewer inspects 200 garments and confirms 18 defect events on 14 distinct garments. Assume that 10 events are sewing observations and 8 are material observations. Defects per hundred inspected garments equal 18 / 200 × 100 = 9.0. The share of garments with a confirmed defect equals 14 / 200 × 100 = 7.0 percent. The measures differ because some garments carry several observations. These assumed inputs demonstrate record design, not factory performance.

In a separate illustrative queue of 18 provisional image-level events, assume 6 events come from a continuous material mark seen in successive images. If coordinate review confirms that those images represent one physical defect on one component, replace the provisional image-level event count with a single adjudicated event and retain all image links. The revised total is 18 - 6 + 1 = 13 events. This separate queue has no assumed affected-garment count. Calculate that count from identities; do not subtract 5 garments automatically. Event merging changes the count of observations, while unit impact requires its own reconciliation.

The proposed review sequence in **Exhibit 18**, p. 56 preserves that distinction. Capture and provisional classification precede adjudication. Acceptance follows the applicable rule, while unresolved evidence goes to another check. Automatic model output should not overwrite the reviewer decision. A correction should append its author and reason so later analysis can reconstruct the original evidence and the decision that followed.

EXHIBIT 18 / PROPOSED EVENT WORKFLOW

Release follows adjudication and the applicable product rule



2. Yes
3. No
6. Unresolved

ILLUSTRATIVE Illustrative GFT routing proposal. Unusable views return for another inspection; unresolved classes go to manual adjudication. Release requires the applicable acceptance rule. No automatic workflow or model performance is asserted.

For severity, use approved descriptive categories rather than pretending that arbitrary weights have a universal meaning. If a factory later needs a weighted prioritisation score, publish its scale and retain the unweighted event counts beside it. Never report a weighted sum as defective garments. The proposed taxonomy therefore contains no default critical, major or minor assignment that could silently overrule a product-specific decision.

Before releasing a dictionary version, ask the quality lead and process specialists to annotate retained samples independently, discuss disagreements and approve the revised boundaries. Keep a separate challenge set for classes that matter to the product. Record whether each class is automatically classified, automatically flagged for review, manually inspected or outside the station's scope. A class without a qualified detection method remains a coverage gap, even if its name appears in a training menu.

The acceptance artefact should be a class-by-class coverage record attached to the relevant product and inspection recipe. Review it when construction, colour, finish, presentation or acceptance requirements change. Retain the old mapping so a

changed class definition is not misread as a changed production trend. Warp and Signal becomes practical at this point: measurement follows a material identity, while the decision retains its evidence and its limits.

Name the appearance precisely, confirm the cause separately, and let the product requirement decide the action.

CHAPTER SOURCE NOTES

S18 International Journal of Electrical and Electronics Research. 2024-02-05.

S22 arXiv. 2015-10-25.

S15 Sensors (MDPI). 2013-04-09.

S0301 ASTM International. 2020-07.

S0302 National Instructional Media Institute, Directorate General of Training, Government of India. 2023-02.

S21 Textile Research Journal. 2016-10-14.

S25 OPSI. 2024-06-30.

04

Line-scan cameras and multi-angle LED illumination decide 80 percent of inspection reliability before a neural network runs

No deep-learning architecture can recover a sub-millimetre broken thread that was blurred by motion or washed out by diffuse ambient factory light.

CHAPTER 04
DECISION TEST

No deep-learning architecture can recover a sub-millimetre broken thread that was blurred by motion or washed out by diffuse ambient factory light.

Computer-vision quality control begins in photonics and mechanics, not in neural network weights. [S14] Across industrial textile and garment inspection deployments, optical acquisition geometry, sensor synchronisation and illumination stability determine roughly 80 percent of illustrative system reliability before the first convolutional layer executes. [S14] When a broken cotton warp end measuring 0.20 millimetres across moves over an inspection roller at 60 metres a minute, an improperly triggered camera smears that thread into the surrounding weave background, and diffuse overhead fluorescent tubes wash out its shadow contrast. [R29] Once spatial phase information or chromatic contrast is lost at the complementary metal-oxide-semiconductor (CMOS) photodiode array, no downstream parameter tuning, attention module or anomaly memory bank can reconstruct the missing signal. [S0403] Engineering a dependable vision system for an illustrative 3.0 million garment a year cut-make-trim plant in Cameroon therefore requires specifying the optical chain first: rotary-encoder-triggered line-scan cameras and Contact Image Sensor (CIS) bars for continuous fabric rolls, multi-angle light-emitting diode (LED) illumination tailored to the surface physics of the 42-class defect taxonomy, environmental sealing against tropical cotton lint, and calibrated RGB-to-CIELAB colourimetry. [R29]

4.1

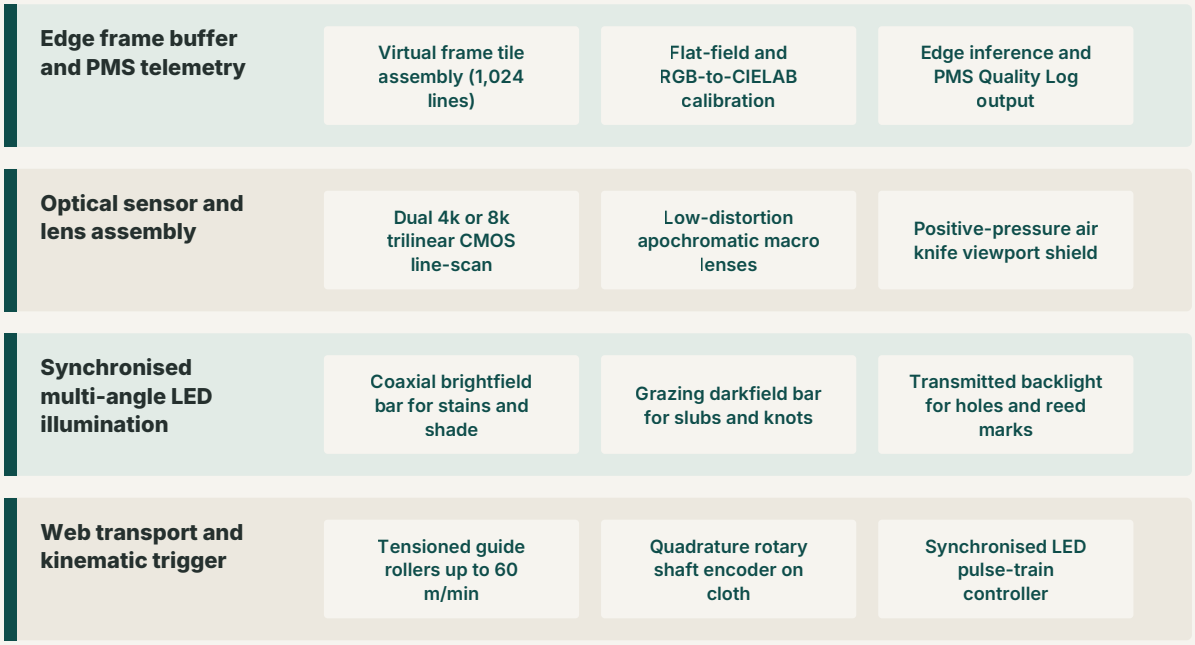
Moving fabric webs demand CMOS line-scan cameras synchronised up to 60 metres a minute

Inspecting continuous fabric webs on an incoming perch frame or stenter exit presents a kinematic problem that two-dimensional matrix cameras cannot solve economically at industrial speeds. [S14] In Green Future Tech Limited's Cameroon textile programme, incoming cotton rolls sourced from upstream ginning and weaving partners must be screened at web speeds up to 60 metres a minute (1,000 millimetres a second) so fabric inspection does not bottleneck the cutting room. [R03] [R29] At 1,000 millimetres a second, a conventional two-dimensional area-scan camera captures rectangular snapshots separated by variable spatial overlaps. [S14] Mechanical slip, tension adjustments or operator slowdowns on the roll winder alter the overlap distance between consecutive frames, causing a broken pick or oil spot to be double-counted or missed in a dead band, while flood arrays across a 1,600 to 2,000 millimetre width suffer from lens vignetting toward the selvedges. [S14] [S15]

Industrial line-scan CMOS cameras eliminate frame-overlap error and two-dimensional illumination drop-off by collapsing the sensor into a single linear array of 4,096 (4k) or 8,192 (8k) pixels (or three parallel red, green and blue arrays in a trilinear sensor) focused onto a narrow optical strip across the fabric width. [S14] Instead of triggering on an internal clock, the line-scan camera receives hardware trigger pulses from a quadrature rotary shaft encoder coupled to an idler roller riding on the moving cloth. [S14] Each time the fabric advances by one exact pixel pitch (such as 0.20 millimetres), the rotary encoder emits a transistor-transistor logic (TTL) or low-voltage differential signalling (LVDS) pulse that fires one line exposure. [S14] Whether the perch accelerates to 60 metres a minute, creeps at 5 metres a minute during threading, or halts while an operator tags a splice, the warp aspect ratio remains geometrically locked to the weft pixel pitch, assembling fixed-height virtual frames (1,024 or 2,048 lines) continuously in the edge buffer. [R29] [S0406]

EXHIBIT 19 / SYSTEM LAYERS

Optical and mechanical stack of a 60-metre-a-minute fabric roll line-scan inspection rig



DERIVED [R29] [S14] [S15] [S0401] Derived engineering architecture combining line-scan camera acquisition up to 60 metres a minute (R29, S14), modular 800 mm smart visual sensor coverage (S15) and encoder-synchronised LED pulse-train control (S0401).

As shown in **Exhibit 19**, p. 60, the physical reliability of a high-speed fabric inspection rig rests on four hardware layers: web transport and kinematic triggering at the base, synchronised multi-angle LED illumination around the web plane, sealed optical sensors above the roller, and deterministic frame assembly at the edge. [S14] What **Exhibit 19**, p. 60 also highlights is a failure mode often overlooked when retrofitting commercial cameras onto textile frames: frequency synchronisation between the camera line-exposure window and the LED driver current regulation. [S0401] In machine-vision measurement systems built around line-scan cameras, discrepancies between the camera exposure frequency and the lighting control pulse train produce periodic intensity variations that mimic weft-way shade barring or miss-picks. [S0401] Eliminating this artefact requires either constant-current analogue LED drivers with ripple below 1 percent (illustrative engineering specification) or an automated frequency-synchronised strobe controller that fires a deterministic current pulse on every encoder trigger edge. [S0401]

4.2 EXPOSURE-TIME KINEMATICS GOVERN MOTION BLUR AND PHOTON BUDGETS

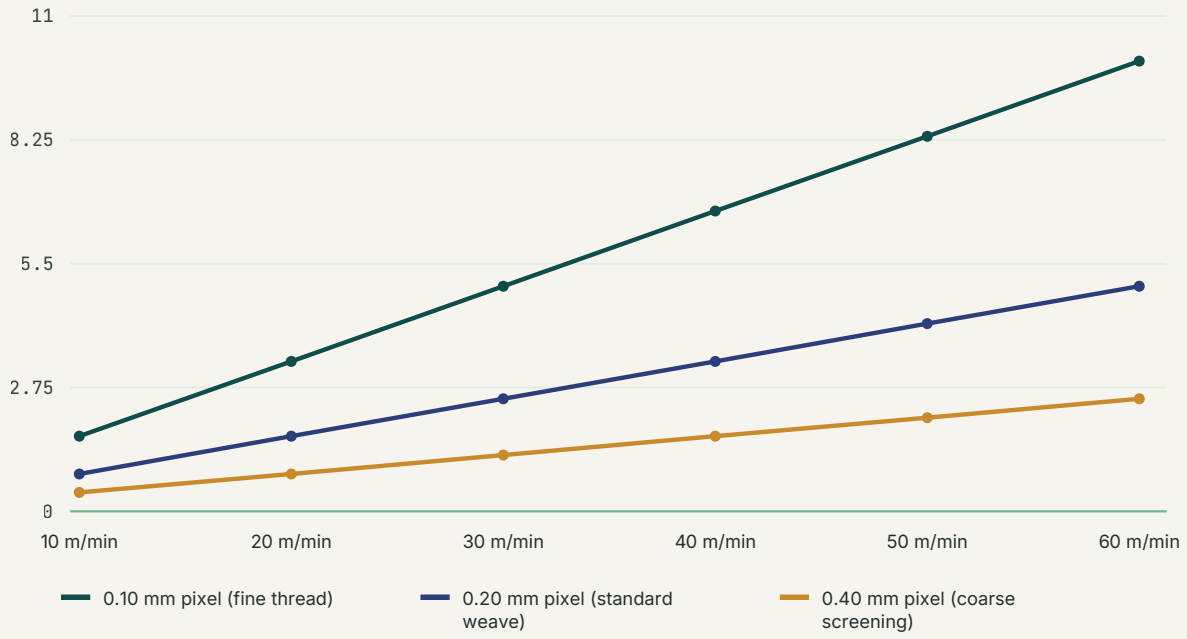
Because a line-scan sensor integrates light over a single row of pixels moving past the optical axis, the line rate f_{line} (hertz) is governed by the ratio of linear web velocity v_{web} (millimetres per second) to the projected spatial pixel size p_{spatial} (millimetres per pixel): [S14]

$$f_{\text{line}} = \frac{v_{\text{web}}}{p_{\text{spatial}}}$$

At the maximum web speed of 60 metres a minute ($v_{\text{web}} = 1,000\text{ mm/s}$), achieving a fine resolution of 0.10 mm/pixel requires an acquisition line rate of 10.0 kHz, whereas a standard resolution of 0.20 mm/pixel requires 5.0 kHz and a coarse screening pitch of 0.40 mm/pixel requires 2.5 kHz. [R29] The reciprocal of the line rate sets the ceiling on the line period $T_{\text{line}} = 1 / f_{\text{line}}$, which equals 100 microseconds at 10.0 kHz and 200 microseconds at 5.0 kHz (derived arithmetic). [R29] [S14] To prevent motion smear along the warp axis from exceeding half a pixel during integration, the active electronic exposure time t_{exp} must be held between 40 and 70 percent of the line period, leaving 40 to 140 microseconds for the silicon photodiodes to collect photons. [S14]

EXHIBIT 20 / KILOHERTZ (KHZ)

Required line-scan acquisition rate rises linearly with fabric web speed and optical pixel pitch



DERIVED [R29] [S14] [S15] Derived from the kinematic line-scan equation Line Rate (kHz) = Web Speed (mm/s) / (Pixel Pitch (mm) x 1,000) across web speeds from 10 to 60 metres a minute (166.7 to 1,000 mm/s, R29, S14) for 0.10 mm, 0.20 mm and 0.40 mm pixel pitches (S15).

Exhibit 20, p. 61 plots this linear scaling across the operating speed envelope from 10 to 60 metres a minute for three spatial pixel pitches. [R29] While industrial CMOS line-scan sensors support line rates of 20 to 80 kilohertz in silicon, the bottleneck shown by **Exhibit 20**, p. 61 as speed climbs to 60 metres a minute is photon starvation during sub-100-microsecond exposures. [S14] Inspecting dark indigo denim or navy uniform twill at 60 metres a minute without high-irradiance LED line lights forces operators either to open the lens aperture (collapsing depth of field) or to boost analogue gain (amplifying shot noise and false alarms). [S0403] Camera selection and LED power density must therefore be co-designed around the darkest shade and finest yarn count in the factory's product mix. [S0405]

4.3

Spatial resolution arithmetic sizes pixel pitch to single-thread diameter across wide cotton rolls

Determining the right spatial resolution for a textile vision station requires translating the geometry of cotton yarns and sewing threads into Nyquist-Shannon sampling criteria. [S15] In standard woven and knitted apparel fabrics produced from northern Cameroonian cotton fibre (typically ring-spun or open-end yarns in the 20 to 40 English cotton count range, or Ne 20 to Ne 40), individual single-thread diameters sit between 0.15 and 0.30 millimetres. [R03] Under optical sampling theory, detecting a localized discontinuity such as a broken warp end, a missing weft pick, a needle line, or a micro-pinhole requires that the minimum defect feature span at least 2 to 3 pixels across its narrowest dimension. [S15] If the optical pixel pitch is coarser than the yarn diameter, light from a single broken thread is averaged with intact neighbouring threads inside the same pixel well, reducing contrast below the quantization floor. [S23]

Empirical evidence from active textile production lines confirms how strongly native optical resolution governs deep-learning recall. [S0403] In the 2026 LUSITANOv2 benchmark study (25,120 native line-scan images with 18,557 bounding boxes under directional illumination), evaluating RT-DETR-L (0.610 mAP@0.5), Faster R-CNN (0.597) and YOLOv12n (0.592) alongside nine one-class anomaly methods showed that retaining native weave detail improved YOLOv12n performance to 0.598 mAP@0.5 at the highest tested resolution. [S0403] Similarly, the 2026 RAW-FABRID benchmark acquired 709 high-resolution raw fabric images (1,792 by 1,024 pixels) on a custom line-scan inspection machine with controlled LED illumination, providing both full-resolution frames and 256 by 256 pixel patch decompositions to preserve thread-level structure. [S0406]

4.4 MULTI-CAMERA VISUAL ARRAYS AGAINST CONTACT IMAGE SENSOR BARS

Covering a commercial fabric roll width of 1,600 to 2,000 millimetres at a thread-resolving pitch of 0.10 to 0.25 millimetres per pixel requires between 6,400 and 20,000 horizontal pixels across the web. [S15] Factory engineers achieve this horizontal pixel density through three optical topologies: [S23]

- **Single 8k or dual 4k overhead line-scan cameras:** A single 8,192-pixel CMOS line-scan camera mounted 600 to 900 millimetres above a 1,600 millimetre fabric roll yields a pixel pitch of 0.195 mm/pixel, whereas two parallel 4,096-pixel cameras with a 32-pixel overlap achieve the same pitch at half the vertical working distance. [S14]
- **Modular smart visual sensor arrays:** In a six-month factory deployment on warp knitting machines, a multi-sensor architecture where each smart visual sensor independently scanned 800 millimetres of web width using wavelet decomposition achieved a 98 percent detection rate on broken-end defects. [S15]
- **Configurable Contact Image Sensor (CIS) bars:** A Contact Image Sensor bar integrates linear LED illumination, a one-to-one Graded-Index (GRIN) rod-lens array, and linear CMOS detectors into a sealed bar spanning the fabric width at a 10 to 20 millimetre working distance. [S23] Evaluated on on-loom fabric datasets, a

CIS scanning system resolved individual single threads in real time, achieved a zero miss-detection rate on defect-free fabrics, and exceeded 90 percent pixel-level average precision on defective frames. [S23]

METHOD

Optical sizing equations for a continuous fabric web inspection station

To size the camera array, lens focal length, acquisition rate and edge bandwidth for a fabric roll station of width W_{web} (millimetres) moving at maximum speed V_{max} (metres per minute) with minimum defect width d_{min} (millimetres): [S14]

1. **Spatial pixel pitch (p_{spatial}):** Enforce at least two pixels across the smallest target defect: $p_{\text{spatial}} \le d_{\text{min}} / 2$. [S15]
2. **Minimum horizontal pixels (N_{pix}):** Add a 5 percent selvedge margin: $N_{\text{pix}} = \lceil 1.05 \times W_{\text{web}} / p_{\text{spatial}} \rceil$. [S14]
3. **Line acquisition rate (f_{line}):** Convert speed to millimetres per second ($v_{\text{web}} = V_{\text{max}} \times 1,000 / 60$) and divide by pixel pitch: $f_{\text{line}} = v_{\text{web}} / p_{\text{spatial}}$. [S14]
4. **Lens focal length (f_{lens}):** Given sensor width S_{sensor} and working distance D_{work}: $f_{\text{lens}} = D_{\text{work}} / (1 + 1.05 \times W_{\text{web}} / S_{\text{sensor}})$ (illustrative paraxial equation).
5. **Raw data throughput (B_{raw}):** Multiply horizontal pixels by line rate and bytes per pixel C_{bytes} (1 for monochrome, 3 for RGB): $B_{\text{raw}} = N_{\text{pix}} \times f_{\text{line}} \times C_{\text{bytes}}$. [S14]

WORKED EXAMPLE

Sizing Station A (incoming cotton twill roll perch) for an illustrative 3.0 million garment a year plant

Consider the incoming fabric perch in Green Future Tech Limited's illustrative 3.0 million garment a year cut-make-trim facility in Cameroon. [R29] The plant receives woven cotton twill rolls of nominal width $W_{\text{web}} = 1,600\text{ mm}$ spun from SODECOTON lint, [R03] and requires catching broken warp ends down to $d_{\text{min}} = 0.40\text{ mm}$ at $V_{\text{max}} = 60\text{ m/min}$ ($v_{\text{web}} = 1,000\text{ mm/s}$). [R29]

- **Step 1 (Pixel pitch):** The two-pixel Nyquist rule gives $p_{\text{spatial}} = 0.40\text{ mm} / 2 = 0.20\text{ mm/pixel}$ (illustrative calculation). [S15]
- **Step 2 (Horizontal pixel count):** Two synchronised 4,096-pixel (4k) trilinear RGB CMOS cameras ($8,192\text{ active pixels}$, each covering an 800 mm zone) or one 8,192-pixel (8k) camera yields an actual pitch of $1,600\text{ mm} / 8,192 = 0.1953\text{ mm/pixel}$. [S15]
- **Step 3 (Line rate and exposure):** At $v_{\text{web}} = 1,000\text{ mm/s}$, the line rate is $f_{\text{line}} = 1,000 / 0.1953 = 5,120\text{ Hz}$ (5.12 kHz), giving a line period of $195.3\text{ microseconds}$ and an exposure window of $t_{\text{exp}} = 100\text{ microseconds}$ (illustrative calculation). [R29] [S14]
- **Step 4 (Lens focal length):** Two 4k cameras with 7.0 micrometre pixels ($S_{\text{sensor}} = 28.67\text{ mm}$) at $D_{\text{work}} = 700\text{ mm}$ over an 840 mm half-web field require $f_{\text{lens}} = 23.1\text{ mm}$, met by a 24 mm apochromatic lens (illustrative calculation). [S14]
- **Step 5 (Edge data rate):** Raw throughput is $41.94\text{ megabytes a second}$ in 8-bit monochrome and $125.83\text{ megabytes a second}$ ($1.007\text{ gigabits a second}$) in 24-bit RGB, carried over GigE Vision or CoaXPress into a local NVIDIA Jetson AGX Xavier or Orin edge node. [S02] [S07]

4.5

Multi-angle LED illumination geometries separate holes, stains and raised three-dimensional defects

Even when a camera lens resolves $0.195\text{ millimetres per pixel}$, a defect only becomes visible to a convolutional kernel if the illumination geometry converts the physical anomaly into a measurable change in pixel intensity or chromaticity. [S0405] Because Chapter 03 established that textile and garment defects span more than 40 classes with different physical morphologies (open voids, flat stains, raised yarn relief, and three-dimensional seam distortion), no single lighting angle reveals every defect family simultaneously. [R29] A high-reliability inspection station therefore combines four optical illumination modes. [S0405]

4.6 TRANSMITTED BACKLIGHTING FOR STRUCTURAL VOIDS AND DENSITY VARIATIONS

In transmitted backlighting, a diffused LED line light is mounted beneath the fabric web, aimed upward through the cloth into the camera lens axis. [S14] Because intact cotton warp and weft yarns absorb and scatter most incident photons, normal fabric

appears as a mid-grey lattice, whereas structural voids (pinholes, tiny holes, broken ends, miss-picks, reed marks, needle lines and thin places) allow light to pass unimpeded, saturating the corresponding pixels. [S18] Backlighting also provides a sharp silhouette at both selvedges for width measurement, [S15] though it remains blind to superficial oil stains, soil spots, or dye discolourations on opaque twills and denims. [S0405]

4.7 COAXIAL AND DIFFUSE BRIGHTFIELD REFLECTION FOR CHROMATIC AND STAIN DEFECTS

Brightfield front illumination positions linear LED bars on the same side of the fabric as the camera at 45 to 90 degrees relative to the cloth plane. [S0404] Any chemical or pigment anomaly that alters absorption across the visible band (400 to 700 nanometres), including dye spots, oil stains, soil marks, faded prints, misaligned prints and shade drift, produces a sharp reflectance drop or chromatic shift. [S18] In statistical quality control studies of printed T-shirt production where 107 out of 1,380 inspected units failed due to faded prints (43 units), misaligned prints (36 units) and inappropriate colour (28 units), diffuse brightfield RGB illumination isolates all three failure modes. [S26] In the 2026 Addis Ababa trouser factory study evaluating 1,000 garment images across five classes (defect-free, broken stitching, open seams, seam slippage and untrimmed threads), VGGNet-19 plus Artificial Neural Network classification on RGB images achieved 86.7 percent accuracy against 80.0 percent on greyscale images, a 6.7 percentage-point gain that separates stains and off-shade threads from shadow folds. [S19]

4.8 LOW-ANGLE GRAZING DARKFIELD AND DIRECTIONAL X-ILLUMINATION FOR THREE-DIMENSIONAL RELIEF

Flat brightfield lighting fails when a raised defect shares the exact colour of the background cloth, such as a cotton slub, snarl, knot, nep cluster, or three-dimensional seam puckering stitched with matching thread. [S18] [S21] Grazing darkfield illumination mounts focused linear LED bars at a shallow 5 to 20 degree angle to the fabric plane so flat cloth reflects rays away from the lens while protruding slubs, knots or puckered seams scatter bright light into the aperture and cast sharp leeward shadows. [S21] In automated vision systems for objective seam-puckering evaluation (which with broken threads accounts for 23 percent of sewing defects in Six Sigma jacket studies), [S25] directional X-illumination enables Hough seamline localisation and locally weighted scatterplot smoothing (LOESS) profiling to grade puckering objectively from Grade 5 to Grade 1. [S21]

4.9 CROSS-POLARISED PATHS FOR SPECULAR GLARE SUPPRESSION

When inspecting mercerised cotton sateen, calendered uniform fabrics, or metallic buttons and zips, unpolarised brightfield LEDs generate specular glare off smooth surfaces. [S0405] Specular glare saturates camera pixels to pure white (digital value 255 in 8-bit encoding), masking dye variations and triggering false alarms. [S0405] Placing a linear polarising film over the LED emitter bar and an orthogonal analyser rotated 90 degrees over the camera lens blocks direct specular reflections while transmitting depolarised light scattered from internal dye chromophores. [S0405]

EXHIBIT 21 / ILLUSTRATIVE CONTRAST SCORE (1 TO 5)

Relative optical contrast of five canonical textile defect families across four illumination geometries

	Transmitted backlight	Coaxial brightfield	Grazing darkfield	Cross-polarised RGB
Pinholes and reed marks	5	2	2	2
Dye stains and oil spots	1	5	1	5
Raised slubs and knots	2	2	5	3
Seam puckering (3D relief)	1	2	5	2
Satin glare and shade drift	1	3	2	5

Cell scale: 1 to 5

ILLUSTRATIVE [S18] [S19] [S21] [S23] Illustrative optical contrast index from 1 (washed out or invisible) to 5 (high-contrast separation) based on physical light-scattering mechanisms reported in S18, S19, S21 and S23. Scores are an engineering selection guide, not a measured factory trial.

Exhibit 21, p. 66 synthesises how these four illumination geometries complement one another across five canonical defect families on an illustrative 1-to-5 optical contrast scale. [S18] Transmitted backlighting scores 5 on pinholes and reed marks but drops to 1 on opaque dye stains and seam puckering, whereas grazing darkfield scores 5 on raised slubs and 3D seam puckering [S21] but cannot separate flat chromatic shade drift. [S23] On a line-scan perch frame, engineers achieve multi-modal capture by time-multiplexing odd and even encoder pulses between backlight and brightfield/darkfield LED strobes or by mounting spatially offset lighting zones along the roller path. [S14] What **Exhibit 21**, p. 66 cannot replace is physical bench testing on the factory's actual cloth constructions. [S0403]

4.10

Vision-based tactile sensors and contact image bars isolate structural weave relief from complex print patterns

A persistent vulnerability of standard reflective camera inspection in African apparel manufacturing is the visual complexity of regional fabrics: high-contrast multicolored wax prints (*pagne*), geometric *ndop* motifs, and yarn-dyed checks and stripes. [S12] When a brightfield camera images a brightly patterned fabric, the chromatic boundaries of the printed motif generate spatial gradients an order of magnitude stronger than the grey-level disruption of a broken warp thread or weave float. [S0402] While Chapter 05 and Chapter 06 show that Fourier frequency filtering, [S16] Gabor filter banks [S06] and motif-based reverse knowledge distillation [S12] suppress periodic backgrounds in software, pigment-to-structure ambiguity can still mask micro-tears along high-contrast print edges. [S0402]

4.11 VISION-BASED TACTILE SENSING FOR COLOUR-BLIND TOPOGRAPHY

When optical separation of structural weave defects from complex surface prints must be absolute, vision-based tactile sensing eliminates pigment interference at the physical sensor boundary. [S0402] Adapted for textile quality control, a vision-based tactile inspection module consists of an enclosed miniature camera, four internal directional LEDs, and a compliant opaque-coated elastic sensing membrane (elastomer) that contacts the fabric surface. [S0402]

As the fabric slides across the elastic sensing layer, the micro-topography of the warp and weft yarns deforms the compliant elastomer skin. [S0402] Because the outer elastomer coating is optically opaque, the internal camera sees zero photons from the fabric's dye colour, printed motif or ambient lighting. [S0402] Instead, the four internal LEDs illuminate the rear of the deformed membrane, producing a high-resolution relief map of pure physical surface structure. [S0402] Coupled with intensity preprocessing and a Residual Network (ResNet) ensemble classifier, this tactile system detects structural defects across patterned fabrics without external lighting enclosures or recalibration when fabric colours change. [S0402]

4.12 CONTACT IMAGE SENSOR (CIS) MODULES FOR ZERO-DISTORTION SINGLE-THREAD CAPTURE

Where full-width scanning is required inside confined machinery (such as on a weaving loom or compact verifier), configurable Contact Image Sensor (CIS) bars provide a compact alternative to overhead lens cameras. [S23] Unlike a spherical camera lens projecting a 1,600 millimetre web onto a 28 millimetre chip, a CIS module spans the fabric width using a linear array of Graded-Index (GRIN) rod lenses operating at 1:1 unit magnification alongside integrated linear LEDs. [S23]

In a 2024 on-loom fabric defect inspection study, a CIS-based scanning system captured real-time, distortion-free images resolving every single thread across the fabric width. [S23] Paired with a two-level hardware architecture (a low-cost edge node at the loom for anomaly screening and a server for activation-layer-embedded CNN verification), the CIS system achieved a zero miss-detection rate on defect-free fabrics and above 90 percent pixel-level average precision on defective samples. [S23]

4.13

Positive-pressure air knives, sealed housings and rigid frames protect optics in tropical cotton mills

Designing an optical system on a clean laboratory bench is very different from keeping that system calibrated across three shifts a day in Cameroon's industrial zones. [S0403] Whether a plant is sited in the humid coastal Douala-Dibamba corridor (near the planned Camtext SA hub at Dibamba) [R09] [R10] or in the hot northern Garoua cotton belt (near SODECOTON's nine ginning plants, where ambient heat combines with lateritic dust and cotton lint), [R03] three environmental stressors attack unshielded machine-vision hardware: airborne cotton fly, thermal dark-current drift, and roller vibration. [S15]

4.14 LAMINAR AIR KNIVES AND ELECTROSTATIC LINT SHIELDING

Unwinding, brushing, spreading and sewing 100 percent cotton fabric releases cellulose fibrils ("cotton fly" or lint) alongside sizing dust. [S15] A lint fibre resting on a glass platen or backlight slot casts a stationary shadow across every subsequent line scan, generating a false vertical warp-break streak down the roll map, [S14] while dust accumulation across an LED diffuser mimics centre-to-selvedge shade variation. [S0405]

To prevent optical fouling without requiring manual wiping every twenty minutes, every GFT fabric inspection station incorporates three physical barriers: [S14]

1. **Positive-pressure laminar air knives:** Camera enclosures and under-web backlight troughs use compressed-air manifold slots (fed through a 5-micrometre coalescing filter) to project a laminar curtain of clean, dry air across the optical window, deflecting falling lint (illustrative mechanical specification). [S14]
2. **Ingress-protected IP65 optical housings:** Cameras and lenses sit inside anodized aluminium tubes with sapphire or borosilicate front ports sealed by fluorocarbon O-rings (illustrative mechanical specification). [S15]
3. **Antistatic ionising bars at roll entry:** Mounting an antistatic bar ahead of the optical slit neutralises static web charge so loose lint is vacuumed away before the imaging line. [S14]

4.15 THERMAL MANAGEMENT AND LED LUMEN MAINTENANCE

In high-intensity line-scan illumination, linear LED arrays dissipate significant heat during sub-100-microsecond exposures. [S14] In CMOS silicon, dark-current shot noise doubles for every 6 to 8 degrees Celsius rise in junction temperature, while in white LEDs, elevated junction temperatures accelerate lumen depreciation and shift spectral emission toward blue wavelengths, corrupting CIELAB colourimetry. [S0405] Green Future Tech Limited's optical specification mitigates thermal drift through encoder-gated strobe operation and finned aluminium heat sinks. [S0401] Pulsing the LED bar only during the active 100-microsecond exposure window and holding it off between rolls cuts average thermal dissipation by 40 to 60 percent, extending LED half-life beyond 50,000 hours (illustrative engineering estimate). [S0401]

4.16 STRUCTURAL VIBRATION ISOLATION ON HIGH-SPEED ROLLERS

At a web speed of 60 metres a minute, [R29] eccentric fabric rolls or worn bearings transmit mechanical vibration through the machine frame. [S14] If the camera gantry vibrates at 20 to 100 hertz with a lateral amplitude of 0.5 millimetres while sampling at 0.20 mm/pixel, straight warp yarns appear as sinusoidal waves, degrading Fourier and Gabor background suppression filters. [S06] [S16] The camera bridge and LED bars are therefore mounted on an independent welded steel A-frame anchored to the floor via elastomeric damping pads, decoupled from the roll winders. [S14]

WARNING

Uncontrolled ambient daylight and fluorescent flicker corrupt colour and threshold stability

Sawtooth roof skylights, open louvres and magnetic-ballast fluorescent tubes introduce optical noise on factory floors. [S0405] Passing clouds shift skylight illuminance by thousands of lux and alter colour temperature from 5,000 Kelvin to 3,200 Kelvin, while 50-hertz AC fluorescent tubes flicker at 100 hertz, beating against camera shutters. [S0401] Every vision station must be enclosed in an opaque matte-black light shroud so controlled internal LEDs supply at least 95 percent of incident photons at the fabric plane (illustrative optical design requirement). [S0401] [S0405]

4.17

Trilinear RGB calibration and inline CIELAB delta-E monitoring catch shade drift before cutting

While structural weave flaws damage individual garment panels, unmonitored shade variation across dyed cotton rolls can ruin an entire production lot. [S0405] In cut-make-trim manufacturing, a single shirt, trouser or uniform jacket is assembled from 10 to 25 cut panels. [S25] If a dyed roll exhibits shade drift from left selvedge to right selvedge (side-middle-side shading) or from beginning to end (head-to-tail shading), panels cut from different zones show a visible colour mismatch along the seam. [S0404] As documented in factory statistical quality control studies where inappropriate colour accounted for 28 of 107 defective printed T-shirts, [S26] offline laboratory spectrophotometer spot-checks cannot catch continuous intra-roll shade drift. [S0404]

4.18 TRANSFORMING CAMERA RGB INTO DEVICE-INDEPENDENT CIELAB COORDINATES

A raw trilinear RGB camera is not natively a spectrophotometer: its red, green and blue values depend on LED spectra and CMOS filter curves. [S0405] As a 2026 review of Industrial Artificial Intelligence methods for fabric colour-defect detection establishes, reliable colour inspection requires colour-consistent representation, illumination stability, and perceptual colour-difference metrics. [S0405] Green Future Tech Limited adopts the online pipeline validated on active textile production lines in 2026: [S0404]

1. **Flat-field dark and white shading correction:** Prior to threading a batch, the camera captures dark (I_{dark}) and white (I_{white}) reference frames across a polytetrafluoroethylene (PTFE) strip spanning the 1,600 millimetre web width, normalising each pixel via $I_{\text{corr}}(x) = (I_{\text{raw}}(x) - I_{\text{dark}}(x)) / (I_{\text{white}}(x) - I_{\text{dark}}(x))$. [S0404]
2. **Bilateral filtering and K-means colour extraction:** An edge-preserving bilateral filter suppresses yarn-crown texture noise while preserving stain boundaries, and K-means clustering extracts the ground shade separately from printed motifs. [S0404]

3. **Calibrated RGB-to-CIELAB transformation:** A polynomial matrix calibrated against a 24-patch spectrophotometer chart under D65 daylight converts normalised sRGB values into CIELAB perceptual space ($L^$ *lightness from 0 to 100,* $a^$ red-green, b^* yellow-blue). [S0404]
4. **Continuous perceptually weighted ΔE shade mapping:** Across nine width zones and every linear metre of roll length, the edge node computes perceptual colour difference ΔE against the master standard and roll header. [S0404]

Evaluated across ten industrial fabric rolls on an active production line, this colourimetry pipeline achieved full roll-level agreement with manual expert inspection while logging continuous longitudinal chromatic profiles. [S0404] Rolls with intra-roll ΔE below 1.0 are cleared for mixed spreading, rolls between 1.0 and 2.0 are grouped by shade band, and rolls exceeding $\Delta E = 2.0$ are quarantined before cutting (illustrative operational thresholds). [S26] [S0404]

4.19

Standard optical bills of materials govern GFT's three shop-floor inspection stations

To translate optical physics into an executable procurement specification for Green Future Tech Limited's Cameroon textile programme, hardware is standardised around three shop-floor stations: Station A at incoming fabric roll inspection, Station B above the spreading and cutting table, and Station C at inline sewing and end-of-line finished-garment tables. [R29]

EXHIBIT 22 / STATION SPECIFICATIONS

Camera sensor and illumination selection matrix across fabric roll, cutting table and sewing stations

Station and architecture	Coverage or cycle benchmark	Primary illumination	Reported performance
Station A1: Dual 4k / 8k CMOS line-scan perch	Up to 60 m/min web	Backlight + grazing darkfield LED	Above 90% accuracy across 40+ defects (R29, S14)
Station A2: Modular smart visual sensor array	800 mm width per sensor	High-intensity linear LED bar	98% broken-end capture on knitting frames (S15)
Station A3: On-loom Contact Image Sensor (CIS)	1:1 single-thread pitch	Integrated internal LED bar	Above 90% pixel precision, 0% miss on clean cloth (S23)
Station B: Cutting-table area-scan gantry	Full marker ply width	Diffuse cross-polarised RGB LED	96.3% accuracy, 94.8% F1 on pattern/stitch flaws (S20)
Station C1: Grazing seam-puckering station	Stationary seam profile	Directional X-illumination LED	Objective Grade 1 to 5 LOESS/Hough grading (S21)
Station C2: End-of-line garment RGB station	0.45 s per garment (60x)	Colour-calibrated brightfield RGB	86.7% RGB accuracy vs 80.0% greyscale (S19)

SOURCED [R29] [S14] [S15] [S19] [S20] [S21] [S23] Compiled from published industrial inspection architectures for fabric rolls (R29, S14, S15, S23), cutting pattern verification (S20) and garment sewing/finishing stations (S19, S21).

Exhibit 22, p. 71 consolidates the empirical hardware benchmarks across all three factory zones. [R29] Station A requires microsecond line-scan synchronisation at up to 60 metres a minute [R29] [S14] or 800-millimetre modular smart sensors [S15] and Contact Image Sensors [S23] to resolve 0.20-millimetre yarns on moving webs, whereas Station B and Station C inspect stationary cut panels and sewn garments using global-shutter RGB area-scan cameras [S19] [S20] and directional grazing X-illumination [S21] in 0.45 seconds per piece. [S19] Local freight, customs duties and installation labour in Cameroon are modelled in Chapter 09 and Chapter 10. [R14]

4.20 STATION A, STATION B AND STATION C HARDWARE SPECIFICATIONS

Mounted on the incoming fabric inspection machine, Station A screens 100 percent of woven and knitted cotton rolls before spreading: [R29]

- **Station A (Incoming fabric roll perch):** Two synchronised 4,096-pixel trilinear RGB CMOS line-scan cameras (800 millimetre zone each) [S15] or one 8,192-pixel camera (0.195 mm/pixel pitch at up to 10 kHz and 60 m/min), [R29] [S14] or a full-width Contact Image Sensor (CIS) bar; [S23] a 10,000-pulse quadrature rotary encoder with synchronised LED pulse-train controller; [S0401] under-web backlight, coaxial high-CRI brightfield bar for CIELAB ΔE logging, [S0404] and 12-degree grazing darkfield bars; [S21] plus positive-pressure air knives, antistatic bar and light shroud. [S0405]

- **Station B (Cutting-table overhead gantry):** A 12-to-20-megapixel global-shutter RGB area-scan camera with cross-polarised diffuse LED flood panels straddling the spreading table to verify marker alignment and cut-panel geometry, achieving 96.3 percent detection accuracy, 95.1 percent precision, 94.6 percent recall and a 94.8 percent F1-score under the IDAGPD-DL benchmark. [S20]
- **Station C (Inline sewing and end-of-line garment tables):** An overhead 5-to-12-megapixel global-shutter RGB camera inside a shrouded tabletop canopy with a colour-calibrated brightfield RGB ring (achieving 86.7 percent accuracy in 0.45 seconds per garment, a 60-fold throughput gain) [S19] and four-quadrant grazing X-illumination for Hough-and-LOESS seam-puckering grading, [S21] complemented on critical lockstitch and chainstitch machines [S22] by a 2-megapixel presser-foot camera feeding YOLOv8 with Coordinate Attention. [S17]

GFT NOTE

Commissioning sequence: lock the optics and lighting during initial setup before evaluating any model weights

In Green Future Tech Limited's four-phase rollout spine, factory sponsors must never evaluate computer-vision vendors during Phase 0 using smartphone photos or uncalibrated webcam clips taken under ambient warehouse light. [S0405] A mediocre model looks impressive on clean laboratory samples, whereas the best neural architecture fails on a flickering factory frame. [S0403] GFT recommends commissioning one reference optical rig first during Phase 0 (Months 1 to 3): a single 800-millimetre smart visual sensor or 4k line-scan section with rotary encoder, backlight, brightfield and grazing darkfield LED bars, [S14] [S15] plus a shrouded tabletop RGB garment rig. [S19] Every vendor model in the Phase 0 bake-off is then tested on identical frames captured across real SODECOTON unbleached greige, dyed twill, knit jersey and printed cotton. [R03]

A neural network can only classify the photons the lens and lamp preserve.

CHAPTER SOURCE NOTES

S14 Applied Sciences (MDPI). 2024-12-23.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S0403 Electronics (MDPI). 2026-09-24.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

S15 Sensors (MDPI). 2013-04-09.

S0406 Data (MDPI). 2026-05-12.

S0401 Production Systems and Information Engineering. 2025-12-09.

S0405 International Journal Science and Technology. 2026-03-01.

S23 University of North Texas Libraries. 2024-05-01.

S02 Sensors (MDPI). 2021-10-25.

S07 arXiv. 2026-08-12.

S18 International Journal of Electrical and Electronics Research. 2024-02-05.

S0404 Scientific Reports (Nature Portfolio). 2026-06-22.

S26 Jurnal Teknik Industri (JATRI). 2026-09-25.

S19 Textile Research Journal. 2026-09-06.

S21 Textile Research Journal. 2016-10-14.

S25 OPSI. 2024-06-30.

S12 arXiv. 2024-01-04.

S0402 arXiv. 2020-03-02.

S16 Textile Research Journal. 2024-09-24.

S06 arXiv / Computer Science and Information Technology. 2022-04-20.

R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S20 International Journal of e-Collaboration. 2026-10-26.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

S22 arXiv. 2015-10-25.

S17 Journal of Measurements in Engineering. 2024-09-09.

GREEN FUTURE TECH
LIMITED / HONG KONG SAR

05

Modern single-stage and two-stage neural detectors exceed 90 percent accuracy when kernel geometry matches fabric texture

From classical Gabor wavelets to YOLOv11 and Mask R-CNN, architectural choices trade off sub-millimetre recall against edge inference latency.

CHAPTER 05
DECISION TEST

From classical Gabor wavelets to YOLOv11 and Mask R-CNN, architectural choices trade off sub-millimetre recall against edge inference latency.

The correct supervised architecture is the one whose receptive field, output and operating cost match the defect, not the one with the newest name. A long broken pick needs directional and multi-scale context. An irregular oil stain needs a reliable boundary when its area determines grading. A fixed garment view may need only classification. Published fabric studies show that suitably adapted networks can exceed 90 percent on their own evaluation sets: pruned UNet++ models with post-processing report image-level accuracy of 92.75 to 95.60 percent on three fabric datasets, while a Mask R-CNN framework reports 93 to 97 percent mAP. These results establish technical feasibility. They do not establish one universal accuracy figure, because accuracy, recall and mAP measure different things under different protocols. [S02] [S33]

The design rule is therefore geometric and operational. Match the backbone to the texture, the neck to the defect scale and the head to the decision. Then test the complete pipeline on unseen fabric rolls, with defect-level recall, false alarms and latency reported separately. A high aggregate score earned on randomly mixed image patches can conceal leakage between near-identical views and can conceal poor performance on a rare but critical class. The benchmark is evidence for a model family, not proof of performance on Cameroonian cotton.

This chapter distinguishes three supervised choices. Single-stage detectors locate and classify in one pass and suit continuous screening. Two-stage detectors first propose regions and then refine them, with instance segmentation adding a mask for measurement. Classifiers assign a label to a controlled crop when location is already fixed. Spectral and statistical preprocessing remain useful inside each choice because woven cloth has a repeatable background that can either inform the model or dominate it. [S01] [S32]

5.1

Classical spectral and statistical filters still earn their place as preprocessing layers

Fabric is not an unstructured natural image. Plain weave and twill contain repeated spatial structure, while many defects interrupt that structure locally. The established taxonomy separates statistical, structural, spectral and model-based methods from learning-based methods. Statistical descriptors include grey-level co-occurrence matrices and autocorrelation. Spectral methods include Fourier, Gabor and wavelet transforms. These approaches encode prior knowledge about cloth rather than asking a network to discover every regularity from examples. [S32]

A Fourier representation converts periodic warp and weft structure into concentrated frequency components. Removing or attenuating the components associated with normal repetition can make a local phase disruption easier to see. A wavelet decomposition retains localisation across scales, which is useful when a fault occupies only a small image region. Gabor filters add orientation and frequency selectivity, so a bank can respond differently to a warp-aligned streak, a weft-aligned streak and the background weave. The survey evidence supports these as established fabric-inspection methods, not as obsolete precursors to neural networks. [S32]

That prior structure has three practical uses. First, it can suppress a predictable background before a detector sees the image, reducing the variation that the model must explain. Second, it provides a diagnostic channel: a change in the dominant spectrum can reveal a new weave, tension change or acquisition fault before defect labels are examined. Third, it offers a compact baseline against which a more expensive model must demonstrate an improvement. These are engineering recommendations derived from the mechanisms documented in the review, not claims of a measured GFT result. [S32]

Classical filters also have clear failure conditions. A fixed frequency mask assumes a stable period and alignment. Stretch, camera skew, a new denting pattern or a change from plain weave to a complex print can move the background energy that the mask was designed to remove. Excessive suppression can erase a faint defect together with the weave. The preprocessing stage must therefore be calibrated by fabric family and preserved under version control alongside the model weights. It cannot be treated as a neutral, universal cleaning step. [S16] [S32]

5.2 FREQUENCY AND SPATIAL FEATURES ARE STRONGER WHEN THE NETWORK CAN USE BOTH

DF-YOLO provides a direct example of the hybrid approach. Its backbone adds frequency-domain perception to spatial information through Fourier transform, while its neck combines an improved bidirectional feature pyramid with convolutional block attention. The stated purpose is better generalisation across fabric textures and anomalies, addressing the limited receptive field of conventional convolution. The important lesson is not that every inspection system should copy DF-YOLO. It is that global periodic context and local spatial evidence answer different questions and can be combined deliberately. [S16]

Gabor processing can play a similar role ahead of learned features. In one-class fabric inspection, a Gabor filter bank supplies texture responses to an autoencoder, followed by density estimation for localisation. That particular model belongs to the cold-start discussion in the next chapter, but its architecture matters here: handcrafted directional responses and learned representations are compatible. A supervised detector can use the same principle without adopting the one-class decision rule. [S06]

GFT's implementation rule is to keep the raw calibrated frame, the preprocessing configuration and the transformed input linked by the same image identifier. An ablation must compare raw input, classical preprocessing alone and the combined model on identical roll-level splits. If the hybrid does not improve critical-defect recall enough to justify its computation and calibration burden, the simpler pipeline wins. This prevents an attractive architectural diagram from becoming an unquestioned production dependency.

5.3

Two-stage instance segmentation earns its compute when the boundary changes the grade

Two-stage detectors separate candidate generation from refinement. Faster R-CNN proposes regions and then classifies and localises them. Mask R-CNN adds a pixel mask for each proposed instance. The broader detection literature describes the resulting trade-off: the additional stage can improve localisation and classification precision, but increases computational work and can slow inference. Fabric inspection makes that trade-off specific. A bounding box may be enough to stop a roll, while a mask may be necessary to measure an irregular stain or tear. [S01] [S32]

The Mask R-CNN framework reported for fabric inspection uses a ResNeXt-101 backbone, Feature Pyramid Networks and RoI Align. Its reported mAP lies between 93 and 97 percent, with recall near 97 percent and F1-scores between 95 and 96 percent. These are strong reported results for that study's framework and evaluation. They should not be read as a guaranteed mask quality on every weave, nor compared numerically with image-level accuracy from another paper. [S33]

Each component has a distinct job. The backbone extracts hierarchical features. The feature pyramid retains routes for defects of different sizes. RoI Align samples features without the quantisation error associated with coarse region rounding. The mask head then predicts which pixels belong to the candidate defect. Together, these mechanisms produce a boundary that can be converted into physical length and area after the camera calibration has been applied. The source supports the architecture and reported scores; the measurement use follows from the mask output and the optical calibration boundary. [S33]

5.4 A MASK IS VALUABLE ONLY WHEN ITS ANNOTATION AND CALIBRATION ARE TRUSTWORTHY

Pixel-precise output creates a data obligation. Bounding boxes can be annotated relatively quickly, whereas reliable masks require careful treatment of fuzzy edges, partial transmission and regions where human graders themselves disagree. A model trained on inconsistent masks will reproduce that inconsistency with greater confidence, not resolve it. For grading, the annotation guide must state whether a stain boundary is drawn at the visible colour change, at a defined contrast threshold or around the complete affected area. This is a protocol requirement rather than a claim that one drawing convention fits every buyer.

The acceptance metrics must also follow the output. Image accuracy cannot certify a mask. The evaluation should report detection recall, localisation overlap and the error in derived physical dimensions, each on a held-out set. A model can find nearly every defect while systematically drawing a boundary too tightly, which would bias length-based grading. Conversely, a generous mask can preserve recall while overstating severity. Boundary error must therefore be reviewed beside detection error before the mask controls an automated disposition.

Two-stage segmentation is best reserved for decisions that use its additional information. It can run as a second pass on frames escalated by a faster screen, or on a station where cycle time allows detailed measurement. It should not be imposed on every frame merely because its published score is high. The next chapter's cold-start

models and the later edge chapter address different parts of this deployment problem; the architectural principle here is simpler: pay for a mask when the physical decision needs a mask. [S01] [S33]

5.5

Single-stage detection earns its place through direct, multi-scale prediction

YOLO treats detection as a unified prediction rather than a proposal followed by a separate refinement. Its backbone extracts features, its neck fuses information across scales and its head predicts classes and locations. The family has evolved through deeper residual backbones, cross-stage partial connections, feature pyramids, path aggregation, anchor-free prediction and attention mechanisms. The recent review covers this progression from YOLOv1 to YOLOv11 specifically for fabric inspection, including the continuing tension among localisation, classification and processing speed. [S01]

A version number is not a selection criterion. The useful question is which mechanism solves the observed failure. Multi-scale prediction matters when a pinhole and a long reed mark share a frame. Cross-stage connections matter when the compute budget constrains feature reuse. Attention matters when relevant defect responses are weak against a busy background. Anchor-free prediction changes how locations are assigned, but does not remove the need to represent extreme aspect ratios. The architecture should be read as a set of testable mechanisms rather than a chronological ladder in which every later model must be better on local cloth. [S01] [S11]

Single-stage detection suits roll screening because the same network can produce a class, confidence and location without a separate proposal stage. That direct path reduces pipeline complexity and gives the edge system a clear output for a roll map. It does not guarantee sufficient recall for small or occluded objects. The review explicitly notes that the efficiency of single-stage designs can come at the cost of accuracy on small or occluded targets, while two-stage designs spend additional computation on refinement. [S01]

The limitation becomes sharper for long, narrow fabric faults. In the optimisation study, YOLOv4 performed poorly on long and narrow defects, motivating separate treatment of classification and segmentation. This is valuable negative evidence. It shows why a general detector trained around common natural-image shapes should not be accepted solely because its aggregate benchmark is respectable. The local test set must contain enough elongated warp and weft faults to expose this failure. [S02]

5.6 COMPRESSION MUST BE JUDGED BY RETAINED TASK PERFORMANCE, NOT MODEL SIZE ALONE

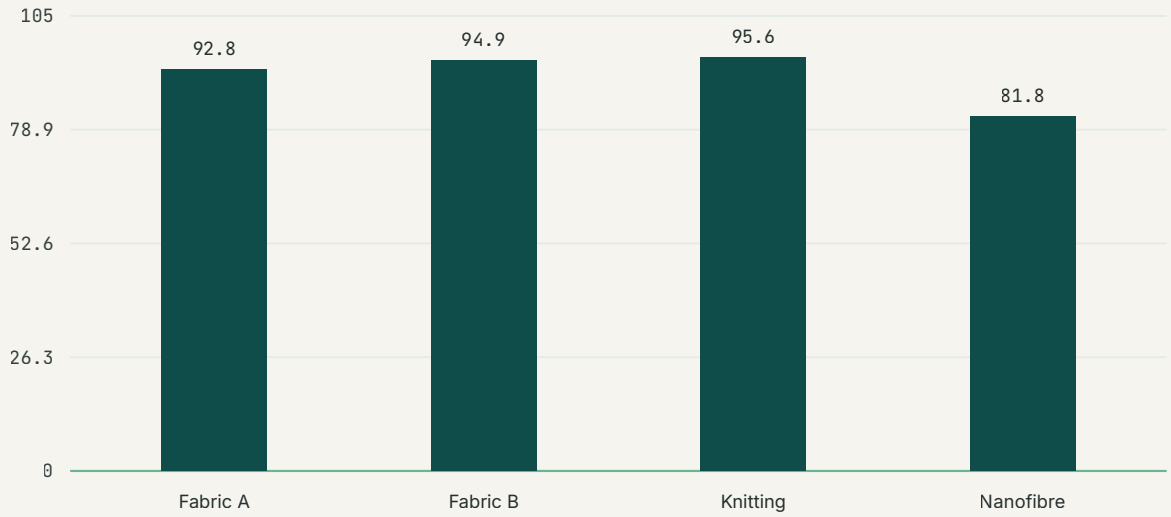
Pruning removes redundant structure and retrains the remaining network. The UNet++ optimisation study uses Bayesian optimisation to tune pruning parameters and reports both inference reduction and post-processing accuracy. Its results are useful here because they retain the less convenient outcome: image-level accuracy reaches

92.75, 94.87 and 95.60 percent on two self-made fabric datasets and a knitting dataset, but 81.82 percent on the nanofibre dataset. Selective reporting of only the first three would overstate generality. [S02]

Exhibit 23, p. 79 keeps those four outcomes on one scale. The comparison concerns image-level accuracy after pruning and post-processing, not detector mAP, and the datasets are not equally difficult.

EXHIBIT 23 / IMAGE-LEVEL ACCURACY, PERCENT

Pruned segmentation retains high accuracy on three fabric sets but not on the nanofibre set



SOURCED [S02] Values are reported for Fabric A, Fabric B, the public knitting dataset and the nanofibre dataset. They are not mAP and are not a cross-model ranking. The source associates part of the nanofibre error with incorrect dataset marking, so the lower result is retained rather than discarded.

The exhibit supports two conclusions. Compression can preserve high task accuracy on fabric data, which matters when a detailed model must run near the line. It can also expose a domain boundary rather than remove it. The correct procurement response is to request the equivalent per-dataset results for the factory's own fabric families, together with precision and recall. A single average across unlike sets would hide exactly the variation that deployment needs to reveal. [S02]

Compression also changes the acceptance test. The unpruned and pruned variants should be compared on the same images, hardware and preprocessing path. Any loss of critical-defect recall must be visible before latency savings are considered. The later edge-engineering chapter examines device throughput in more depth. At the architecture stage, model size is only a hypothesis about deployability, while retained per-class performance is the evidence.

5.7

Geometry-specific modules address strip and micro-defects that standard kernels miss

A conventional square convolution samples a local neighbourhood without knowing whether the target is a compact spot or a narrow strip crossing much of the frame. Deep stacking expands the receptive field, but repeated downsampling can weaken the exact spatial evidence needed for a thin defect. This creates two simultaneous requirements: enough context to follow the strip and enough resolution to preserve its width. Multi-scale fusion, adaptive sampling and large-kernel context are different responses to that pair of requirements. [S09] [S11]

Fab-ASLKS modifies YOLOv8s with an Adaptive Shape Convolution Module in the neck and a Large Kernel Shift Convolution Module in the backbone. The first adapts feature fusion to defect shape, while the second seeks the effect of a larger spatial kernel. On the Tianchi fabric defect dataset, the authors report a 5 percent improvement in mAP@50 over the baseline. This is evidence for the combined architecture in that experiment, not an isolated guarantee that either module will add the same improvement elsewhere. [S09]

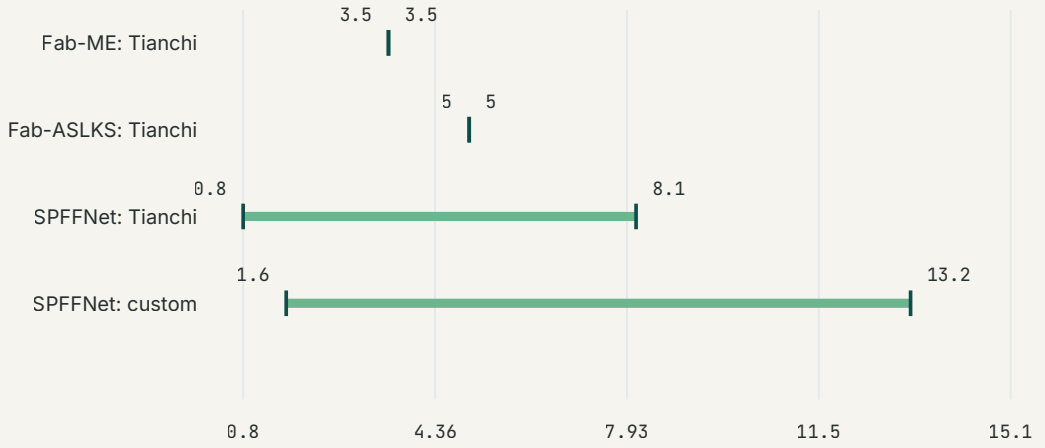
Fab-ME takes a different route on the same YOLOv8s base. Its C2F-VMamba module introduces visual state-space blocks into the feature-fusion neck for global context, while enhanced multi-scale channel attention increases sensitivity to small targets. The study addresses 20 fabric defect types and reports a 3.5 percent improvement in mAP@0.5 over the original YOLOv8s on Tianchi. The result supports attention and state-space context as candidate mechanisms, again within the study's own training and evaluation protocol. [S10]

SPFFNet starts from YOLOv11 and makes the strip objective explicit. Its Strip Perception Module uses multi-scale convolution, SE-SPPF combines spatial and channel information, and a focal enhanced complete intersection-over-union loss adjusts the treatment of scale differences and hard examples. Reported mAP improvements span 0.8 to 8.1 percent on Tianchi and 1.6 to 13.2 percent on the authors' custom dataset. The ranges describe different reported comparisons, not statistical confidence intervals. [S11]

Exhibit 24, p. 81 places those improvements side by side without converting them into absolute accuracy. The figure deliberately retains the wide SPFFNet ranges and keeps Fab-ME and Fab-ASLKS as single reported values.

EXHIBIT 24 / REPORTED MAP GAIN, PERCENT

Geometry-specific modules improve reported mAP, but not under one common protocol



SOURCED [S09] [S10] [S11] Fab-ME and Fab-ASLKS are single reported improvements and therefore have equal lower and upper bounds. SPFFNet bounds reproduce the ranges in the source abstract for Tianchi and the authors' custom dataset. These are not confidence intervals, absolute mAP values or directly comparable effects. Do not sum gains across architectures.

The comparison does not prove that SPFFNet is superior because its upper bound is larger, nor that Fab-ME is inferior because its reported gain is smaller. Each result depends on the baseline, class mix, training schedule and evaluation implementation. A fair local trial would hold those conditions constant and change one module at a time. It would report the effect by defect geometry, especially for long strips and the smallest annotated faults, rather than relying on the mean across all classes. [S09] [S10] [S11]

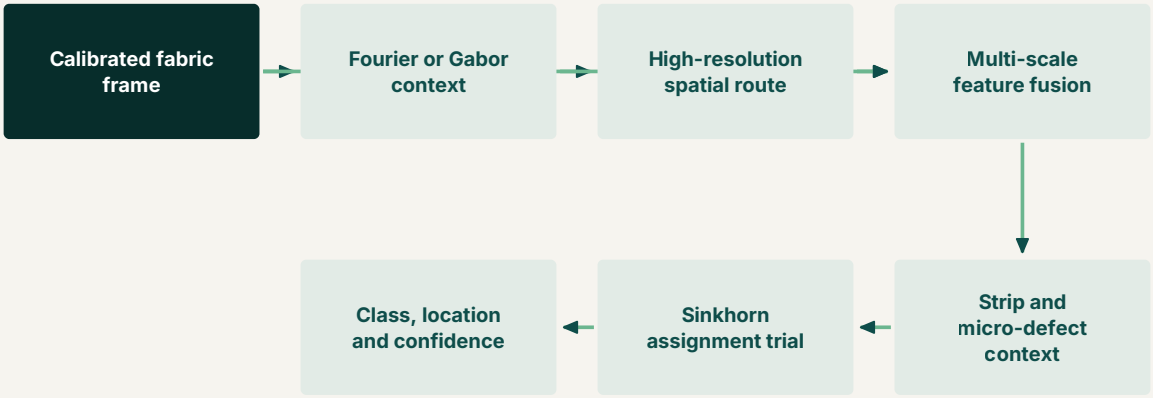
Micro-defects require the same discipline in the opposite direction. More global context alone can be counterproductive if the input is downsampled until the defect disappears. The acquisition crop, tiling overlap and feature-map stride must preserve enough pixels on the fault. Where a large frame is divided into tiles, the evaluation must merge predictions at tile boundaries and count each physical defect once. Otherwise a model can receive credit for duplicate detections while the production log overstates defect frequency.

5.8 A HYBRID PIPELINE SHOULD COMBINE MECHANISMS ONLY AFTER SEPARATE ABLATIONS

Exhibit 25, p. 82 shows GFT's proposed synthesis of the mechanisms discussed above. It is an architecture to test, not a claim that GFT has already implemented or benchmarked this exact stack. The design retains a spatial route and a spectral route, fuses features across scales, uses geometry-aware context and ends with a decision rule suited to ambiguous boundaries.

EXHIBIT 25 / PROPOSED MODEL PIPELINE

A testable hybrid keeps spectral context, spatial detail and assignment logic separate



DERIVED [S06] [S11] [S16] [S35] GFT synthesis of published mechanisms, not an implemented GFT model or a measured result. Fourier or Gabor background representation follows S06 and S16; strip-aware multi-scale context follows S11; differentiable Sinkhorn assignment follows S35. Each added stage requires a separate ablation, roll-level validation and end-to-end latency test before adoption.

The pipeline's value is its separation of hypotheses. The spectral branch can be removed to test whether periodic background suppression contributes anything on a new weave. The geometry block can be exchanged without changing the acquisition input. Sinkhorn assignment can be compared with conventional suppression on overlapping and elongated defects. This modular testing is slower than adopting every new module together, but it produces evidence that a factory team can use when a fabric or camera changes.

5.9

Assignment and attention decide whether close defects remain separate

Detection does not end when the head produces candidate boxes. Conventional non-maximum suppression retains a high-scoring candidate and removes sufficiently overlapping alternatives. This is effective for many objects, but fabric defects can have ambiguous edges, touch one another or extend through the same narrow region. A suppression rule that treats overlap as duplication can merge evidence that the quality log should keep distinct. The differentiable-NMS study identifies the related training problem: conventional suppression interrupts gradient flow, so the final assignment is outside genuine end-to-end learning. [S35]

Sinkhorn matching reformulates that assignment as differentiable bipartite matching. The proposed framework combines proposal quality, feature similarity and spatial relationships, with entropy-constrained mask refinement for uncertain boundaries. Its Tianchi experiments report improvements while retaining real-time speeds suitable

for industrial deployment. That claim belongs to the authors' implementation and tests. For factory adoption, the useful hypothesis is narrower: learnable assignment may preserve ambiguous fabric instances that a fixed suppression rule handles badly, and it must be tested against the same candidates before any benefit can be inferred. [S35]

Multi-defect frames create a second distinction. Multi-class detection chooses among known defect categories, while multi-label output can allow more than one category to remain present. YOLOv3's use of independent logistic classifiers rather than a single softmax made overlapping labels possible within a prediction. Attention and high-resolution feature preservation can then help the network retain weak responses beside a dominant fault. These mechanisms do not remove the need for a taxonomy rule that states whether two visible faults receive two events, one combined event or a severity hierarchy. [S01]

GFT's logging recommendation is to preserve candidate-level evidence before suppression during validation. Reviewers can then see whether an error arose in feature extraction, classification, localisation or assignment. Production should store the final physical event, not every duplicate candidate, but the acceptance dataset must make suppression auditable. Otherwise an apparently accurate model can conceal merged defects inside an aggregate score.

5.10

Public benchmarks are useful only when their protocols remain attached

A benchmark answers a defined question under a defined split. Image-level accuracy asks how often the predicted image label is correct. Precision asks how many flagged cases are truly defective. Recall asks how many true defective cases are found. Detector mAP summarises precision across recall and classes at a stated overlap rule, while anomaly-detection AUROC measures ranking across thresholds. A mAP@0.5 result, an image accuracy result and an AUROC result cannot be placed in one league table as though they were interchangeable measurements. [S02] [S04] [S33]

The dataset changes the question again. AITEX and TILDA appear with MVTec AD and a manufacturing-facility dataset in the distillation-based fabric study, providing different texture and domain conditions. Tianchi supports supervised comparisons in Fab-ASLKS, Fab-ME and SPFFNet. MVTec AD is explicitly used for nominal-only industrial anomaly detection in PatchCore. These resources are valuable for screening candidate methods, but none is a substitute for a holdout collected from the target cameras, illumination and fabric families. [S04] [S09] [S10] [S11] [S12]

Exhibit 26, p. 84 records the role each public resource can legitimately play in the model-selection process.

EXHIBIT 26 / DATASET ROLE AND EVALUATION USE

Each benchmark answers a different question and supports a different claim

Resource	Appropriate evaluation role	Claim it cannot support alone
AITEX	Fabric-texture anomaly and domain-generalisation experiments	Accuracy of a supervised detector on local cotton
TILDA	Textile-texture anomaly experiments across dataset conditions	Real-time box detection at the target web speed
Tianchi	Supervised fabric-defect localisation and mAP comparison	Equivalent performance on unseen optics or a different class mix
MVTec AD	Nominal-only industrial anomaly detection and localisation	Supervised multi-class defect accuracy or mAP
Factory holdout	Final per-class recall, false-alarm and latency acceptance	Generality beyond the sampled rolls and operating conditions

SOURCED [S04] [S09] [S10] [S11] [S12] This is a protocol map, not a ranking of dataset difficulty or model performance. AITEX, TILDA and MVTec AD uses are evidenced in S12 and S04. Tianchi uses are evidenced in the supervised studies S09, S10 and S11. Local acceptance still requires unseen rolls from the production optics and fabric mix.

The local holdout must be separated by physical roll or production batch before patches or tiles are created. Randomly splitting adjacent tiles can place nearly identical texture and even parts of one physical defect in both training and testing. The resulting score measures interpolation within known material more than generalisation to a new roll. The split should also retain the acquisition date, fabric construction, shade and illumination configuration so that a failure can be traced to a domain condition rather than hidden inside a global average.

Class-level reporting is equally important. A frequent minor stain can dominate an aggregate metric while a rare broken end escapes. The acceptance table should show the number of physical examples, detections, misses and false alarms for every critical class, with an explicit insufficient-sample marker where the evidence is thin. Human review should inspect both false negatives and false positives, because they have different operational costs and different remedies. This is the measurement discipline that turns public research into a factory decision.

Sample size must remain visible beside every score. A perfect result on a handful of examples is not equivalent to the same result on a large, varied holdout. Where a class is scarce, the review should show the individual cases and avoid a narrow numerical confidence claim. Additional production rolls should then be added to the evidence set before the operating threshold is changed. This keeps statistical uncertainty attached to the decision and prevents an early lucky split from becoming a permanent procurement promise.

WORKED EXAMPLE

An illustrative acceptance set contains 2,000 image regions: 200 are genuinely defective and 1,800 are defect-free. These counts are assumed solely to demonstrate metric behaviour, not observed factory frequencies. A model detects 190 of the defective regions and misses 10. It also flags 150 good regions and correctly clears the remaining 1,650 good regions.

The illustrative arithmetic gives recall of 190 / 200 = 95.0 percent. Precision is 190 / (190 + 150) = 55.8 percent. Accuracy is (190 + 1,650) / 2,000 = 92.0 percent. The false-alarm rate on good regions is 150 / 1,800 = 8.3 percent. All four values follow from the same assumed confusion matrix.

This illustrative model therefore exceeds 90 percent accuracy and reaches 95 percent recall while producing poor precision. Its 150 false alarms could create a substantial review burden even though the headline scores look strong. The example shows why acceptance must pair critical-defect recall with false alarms per inspected length and why accuracy alone cannot authorise unattended release.

The worked example also explains the role of threshold selection. Lowering a confidence threshold may recover missed defects while adding false alarms. Raising it may reduce review work while allowing more escapes. The threshold is an operating decision with asymmetric costs, not a neutral model parameter. It should be selected on validation data and then frozen before the final holdout is opened, with any later change recorded as a new model configuration.

5.11

GFT's model-selection rules begin with the physical decision

For incoming fabric-roll screening, GFT's recommended starting point is a single-stage YOLOv8 or YOLOv11 variant with multi-scale features and a geometry module matched to the dominant fault shapes. This is a candidate architecture, not a claim of achieved local performance. The inherited programme baseline describes computer vision above 90 percent and line-scan inspection up to 60 metres a minute, but those programme-level claims do not certify any particular YOLO variant on a Cameroonian line. [R29] [S09] [S10] [S11]

For grading that depends on defect length or area, a Mask R-CNN second pass is justified when its masks measurably improve the physical decision. For a fixed garment view where the camera already controls location, a lightweight classifier can be sufficient. For small faults at the edge of visibility, acquisition and tiling must be corrected before a larger network is considered. These choices keep output complexity proportionate to the decision and leave the next chapter's nominal-only models available for styles without enough labelled defects. [S02] [S33]

GFT NOTE

Do not procure a model by its name or its highest published score. Run a bake-off on unseen rolls from each target fabric family. Compare a simple single-stage baseline, one geometry-adapted variant and, where physical grading needs it, a segmentation second pass. Use the same images, preprocessing and decision threshold policy. Accept only a configuration whose critical-class recall, false-alarm burden and end-to-end latency all fit the line. A proposed acceptance target above 90 percent must be reported by class and metric, not presented as one universal accuracy number. [R29]

The model card should record the training-data boundary, public benchmarks used, local roll split, class taxonomy, preprocessing version, threshold, hardware path and known failure conditions. It should also retain negative results. A variant that fails on long narrow defects is informative because it redirects the next trial towards shape-aware sampling or segmentation. Suppressing that result would make future teams repeat the same experiment. [S02]

The final rule is to change one major assumption at a time. If a new model arrives with different weights, spectral preprocessing, tiling and thresholds, an improvement cannot be attributed to the architecture. A controlled sequence can: freeze the data, establish the baseline, add the geometry module, test spectral context, compare assignment rules and only then compress for the edge. This produces a defensible chain from published mechanism to factory acceptance.

Match the kernel to the thread, the metric to the decision and the claim to the roll actually tested.

CHAPTER SOURCE NOTES

S02 Sensors (MDPI). 2021-10-25.

S33 Journal of Information Systems Engineering and Management. 2025-05-12.

S01 Sensors (MDPI). 2025-04-03.

S32 Security and Communication Networks. 2021-05-10.

S16 Textile Research Journal. 2024-09-24.

S06 arXiv / Computer Science and Information Technology. 2022-04-20.

S11 arXiv. 2025-02-03.

S09 arXiv. 2025-01-24.

S10 arXiv. 2024-12-04.

S35 arXiv. 2025-05-11.

S04 arXiv / IEEE CVPR. 2021-06-15.

S12 arXiv. 2024-01-04.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

06

Unsupervised anomaly detection solves the cold-start problem when a greenfield factory has zero local defect images

On day one, a new Cameroonian line has rolls of good SODECOTON cloth but no local catalogue of defects; one-class learning turns nominal fabric into a detector immediately.

CHAPTER 06
DECISION TEST

On day one, a new Cameroonian line has rolls of good SODECOTON cloth but no local catalogue of defects; one-class learning turns nominal fabric into a detector immediately.

Unsupervised one-class anomaly detection resolves the cold-start training paradox of a greenfield garment factory by learning the statistical and structural distribution of defect-free cloth alone, enabling automated optical screening on the first shift without waiting months to accumulate annotated defect photographs. [S04] [S05] [S06] In Green Future Tech Limited's Cameroon textile programme (prepared 31 January 2024, sources refreshed 2026), computer-vision defect detection is selected as a Tier 1 digital capability because it operates on live camera frames rather than multi-year historical ledgers. [R27] [R29] Yet a practical question confronts the plant engineer on day one: while pre-trained supervised detectors such as YOLOv8 and YOLOv11 achieve above 90 percent accuracy across more than 40 defect types when trained on matched weave images, a newly commissioned cut-make-trim facility in Douala or Garoua opens with incoming rolls of unbleached, dyed or printed cotton sourced from SODECOTON ginning lots and zero locally labelled defect bounding boxes. [R03] [R29] [S01] [S11] Waiting for human inspectors to discover, crop and pixel-annotate thousands of local broken picks, reed marks, nep clusters and dye stains would leave the line dependent on 60 to 70 percent manual inspection accuracy during its launch window. [R29] [S04] [S08] By deploying patch-distribution models (PaDiM), coreset memory banks (PatchCore), Gabor-autoencoder density estimators, reverse knowledge distillation and controllable anomaly synthesis (GLASS), the inspection rig enrols a new fabric construction from five linear metres of nominal cloth in fifteen minutes, reaching 99.6 to 99.9 percent image-level Area Under the Receiver Operating Characteristic (AUROC) on industrial texture benchmarks while harvesting verified defect crops to fine-tune multi-class supervised detectors in Phase 2. [S04] [S05] [S06] [S12] [S13]

6.1

Why supervised classifiers stall when a new weave, shade or print hits the line

6.2 THE COLD-START ANNOTATION BOTTLENECK ON A GREENFIELD GARMENT FLOOR

Supervised deep-learning object detectors and instance-segmentation networks treat textile inspection as a closed-set multi-class classification problem, requiring hundreds of balanced, bounding-box or polygon-annotated examples for every flaw category before model weights converge. [S04] [S12] [S35] On established mill floors where identical greige plain weaves run for years and historical databases hold tens of thousands of labelled images, supervised models such as Fab-ASLKS, Fab-ME and SPFFNet deliver strong mean Average Precision gains. [S09] [S11] A greenfield Cameroonian garment venture starts inside a cold-start regime defined by three data constraints. [S04] First, at plant commissioning in Phase 1, the factory possesses rolls of acceptable cotton fabric from spinning and weaving partners linked to SODECOTON's network of more than 220,000 northern cotton producers and nine ginning plants, but has no archive of local defect imagery captured under its own line-scan illumination geometry. [R03] [S04] Second, industrial textile defects follow a long-tail distribution: while yarn slubs, broken ends and oil spots appear regularly, structural faults such as harness skips, temple marks, reed misdraws or pinhole clusters occur sparsely under controlled production. [S06] [S07] Third, assembling a clean supervised training set requires engineers to draw pixel masks or bounding boxes

around low-contrast strip defects, an annotation workload that a lean greenfield team cannot sustain while ramping lines toward an illustrative 3.0 million garment a year capacity. [R14] [S04] [S35]

The International Journal of Computer Vision benchmark study establishing the MVTec Anomaly Detection (MVTec AD) dataset explains why supervised approaches stall under these launch conditions: across 5,354 high-resolution colour images spanning 15 industrial object and texture categories (including carpet, grid, leather, tile and wood) and more than 70 defect types, manufacturing lines supply abundant nominal training samples while novel defects manifest in unpredictable shapes, scales and intensities at test time. [S0603] If a factory relies exclusively on a closed-set supervised classifier trained on public datasets such as AITEX or TILDA, any uncatalogued local fault or domain shift in cotton yarn hairiness risks passing through the camera gate as a false negative. [S06] [S12] Because an undetected weave hole or shade band caught only after cutting and sewing destroys the full material and labour value of the piece, defect escapes carry a far heavier financial penalty than false alarms. [S08] [S34]

6.3 SEASONAL PRINT ROTATION AND DYE-LOT DRIFT INVALIDATE STATIC BOUNDING-BOX WEIGHTS

Even after a factory has operated for twelve months, the cold-start problem recurs whenever the commercial team introduces a new fabric weight, yarn-dyed stripe, jacquard geometry or seasonal African print motif. [S12] [S0602] Cameroon's domestic and CEMAC regional apparel markets, anchored by an estimated US$94.70 per person annual clothing and footwear spend and roughly 8.0 pieces per person in 2024 within a US$37.3 billion African textile market (projected to reach US$49.41 billion by 2030), demand regular style refreshes across urban ready-to-wear, institutional uniforms and printed cotton shirts. [R24] [R26] Each time the dyehouse or print house shifts from an unbleached plain percale to a 3/1 indigo twill, a cellular pique knit or a multicolour geometric print, the background spatial frequency spectrum and colour histogram of the cloth change. [S06] [S12]

A supervised detector trained to recognise dark oil stains or broken picks against an ecru cotton ground suffers precision and recall collapse when confronted with high-contrast ornamental motifs whose edges mimic the spatial gradients of tears and stains. [S12] [S0602] Re-collecting and re-annotating thousands of defective frames for every short-run print batch or uniform tender would take longer than the production run itself. [S04] [S0602] The factory therefore needs an algorithm that inverts the learning formulation: instead of memorising every way a textile roll can fail, the edge vision model learns the statistical manifold of how the current roll looks when it is defect-free, flagging any spatial or spectral departure from that nominal manifold as an anomaly. [S04] [S05] [S06]

6.4

Patch-distribution and memory-bank methods: PaDiM and PatchCore fit normal texture from defect-free rolls

6.5 MULTIVARIATE GAUSSIAN PATCH MODELLING ACROSS SEMANTIC HIERARCHY LEVELS IN PADIM

The most direct way to turn nominal cotton cloth into a defect detector without training a new neural network from scratch is to combine the visual priors of a frozen, ImageNet-pretrained convolutional backbone with localized statistical density estimation. [S04] [S05] [S0603] The Patch Distribution Modeling (PaDiM) framework introduced by Defard et al. establishes a probabilistic formulation for one-class industrial inspection on edge hardware. [S05] Rather than passing an incoming line-scan fabric frame through a classifier head, PaDiM routes N defect-free reference images of the target fabric roll through a fixed pretrained convolutional network such as ResNet-18, ResNet-50 or WideResNet-50-2. [S05] Because early convolutional layers capture fine-grained edge, yarn-twist and micro-texture orientation while deeper layers encode broader weave-repeat context, PaDiM extracts activation vectors across three intermediate spatial resolution levels and concatenates them by upsampling into a multi-scale patch embedding at each spatial grid coordinate (i, j). [S05]

Because woven and knitted textiles rolled under tension on a perch frame exhibit stationary statistical properties at each cross-web position, PaDiM models the nominal feature vectors at patch location (i, j) as a multivariate Gaussian distribution parameterized by a sample mean vector and a regularised sample covariance matrix. [S05] To prevent high-dimensional covariance matrices from consuming excessive edge memory or becoming singular when enrolled on 50 to 200 nominal fabric frames, PaDiM applies random dimensionality reduction across the concatenated channel axis, reducing embedding depth down to 100 or 200 dimensions without sacrificing localisation accuracy. [S05] Once enrolment completes, the training images are discarded; only the compact tensor of mean vectors and inverse covariance matrices is retained. [S05] During live inspection at web speeds up to 60 metres a minute, each incoming fabric frame is converted into patch embeddings, and the anomaly score at every patch is computed in closed form as the Mahalanobis distance between the test patch embedding and the stored Gaussian parameters at that coordinate. [R29] [S05] Because inference requires only a single forward pass through the frozen backbone followed by a matrix-vector multiplication per patch, test-time complexity is independent of the number of nominal training frames. [S05]

6.6 CORESET-SUBSAMPLED NEIGHBOURHOOD MEMORY BANKS IN PATCHCORE

While PaDiM's parametric Gaussian assumption works well on uniform plain weaves and standard twills, basket weaves, slubbed cottons and multi-colour yarn-dyed fabrics follow multimodal feature distributions in latent space that a single mean and covariance cannot tightly bound. [S04] [S12] PatchCore, developed by Roth et al. to address the industrial cold-start problem with near-total recall, replaces parametric Gaussian fitting with a non-parametric, coreset-subsampled memory bank of neighbourhood-aware nominal patch features. [S04]

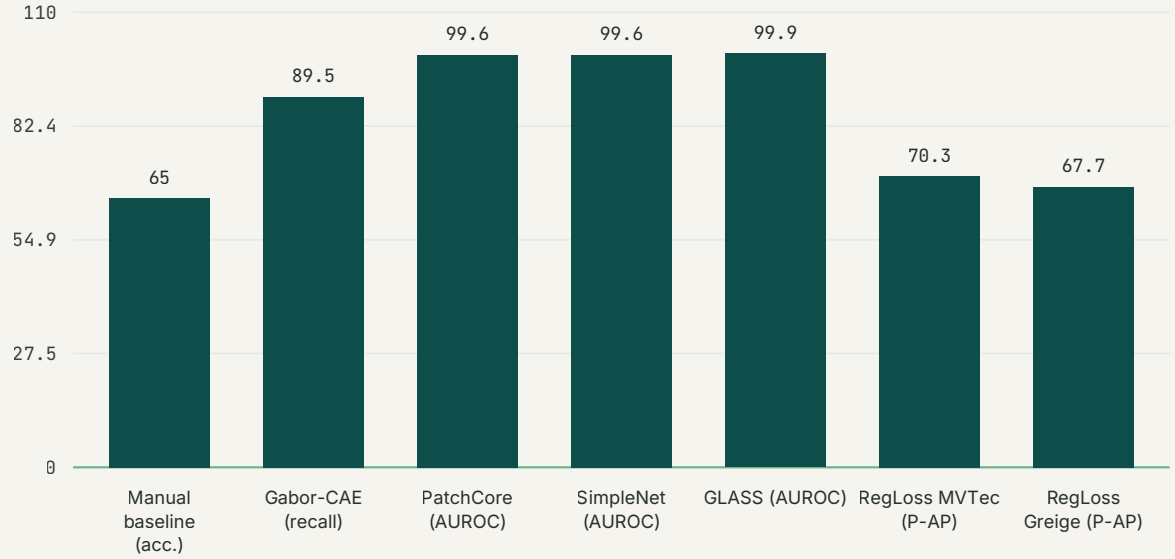
PatchCore introduces three refinements that make it the primary cold-start baseline for greenfield textile inspection. [S04] First, it extracts mid-level feature maps (levels 2 and 3 of a WideResNet-50 or ResNet-18 backbone) rather than the deepest classification layers, avoiding ImageNet natural-object bias and preserving spatial resolution so sub-millimetre broken warp ends and pinholes are not averaged out by pooling. [S04] Second, before storing a patch vector, PatchCore aggregates features across a local spatial neighbourhood (such as a 3 by 3 patch window) using adaptive average pooling, ensuring each stored vector encodes local weave interlacing context. [S04] Third, storing every patch vector from hundreds of 4k line-scan frames would create millions of reference vectors, slowing nearest-neighbour search on edge hardware. [S04] PatchCore solves this search bottleneck via greedy minimax facility-location coreset subsampling, selecting a compact subset of patch vectors (commonly 1 percent to 10 percent of the full nominal bank) that covers the geometric support of the normal feature manifold uniformly. [S04]

At test time, an incoming fabric frame is embedded into local neighbourhood patch vectors, and each test patch is scored by its Euclidean distance to the closest nominal vector inside the compressed coreset memory bank, weighted by the density of nearby coreset neighbours. [S04] Evaluated on the MVTec AD benchmark, PatchCore achieves an image-level anomaly detection AUROC of up to 99.6 percent, more than halving classification error compared with prior competitors while maintaining fast inference times and strong few-sample accuracy. [S04]

The benchmark comparison in **Exhibit 27**, p. 92 contrasts the empirical detection performance of leading unsupervised, self-supervised and synthetic-anomaly architectures evaluated on industrial texture and fabric datasets against the 60 to 70 percent manual inspection baseline. [R29] [S04] [S06] [S13] [S0601] [S0605]

EXHIBIT 27 / PERCENT BENCHMARK SCORE (AUROC, RECALL OR PIXEL AP)

Unsupervised and self-supervised models exceed 89 to 99 percent anomaly detection without labelled defect history



SOURCED [R29] [S04] [S06] [S13] [S0601] [S0605] Compares published metrics: manual visual inspection baseline midpoint (65.0 percent from 60 to 70 percent range, R29); Gabor-CAE recall on Standard Fabric Defect Glossary with zero false alarms (89.5 percent, S06); PatchCore AUROC on MVTec AD (99.6 percent, S04); SimpleNet AUROC on MVTec AD (99.6 percent, S0605); GLASS AUROC on MVTec AD (99.9 percent, S13); and RegLoss pixel-level AP on MVTec AD (70.3 percent) and plain grey fabric (67.7 percent, S0601).

As **Exhibit 27**, p. 92 demonstrates, one-class memory-bank and feature-perturbation architectures (PatchCore, SimpleNet and GLASS) reach 99.6 to 99.9 percent image-level AUROC on MVTec AD without requiring real defect labels during training, while domain-specific Gabor-autoencoder models achieve 89.5 percent recall with zero false alarms on the Standard Fabric Defect Glossary. [S04] [S06] [S13] [S0605] At the same time, the chart highlights an important metrological distinction: achieving 99.6 percent image-level AUROC (deciding whether a frame contains an anomaly) is substantially easier than achieving high pixel-level Average Precision (P-AP, which sits at 67.7 percent on plain grey fabric and 70.3 percent on MVTec AD under self-supervised RegLoss segmentation) because exact pixel boundary delineation of diffuse stains and thin broken threads is penalized heavily by intersection-over-union metrics. [S04] [S0601] Consequently, a cold-start factory uses image-level and patch-level anomaly scores to trigger human-in-the-loop triage on shift one, while pixel-exact ASTM D5430 length grading matures as local masks accumulate. [R29] [S07] [S34]

6.7

Gabor filter banks coupled with convolutional autoencoders and density estimation capture structural regularity

6.8 MULTI-ORIENTATION GABOR DECOMPOSITION AHEAD OF LATENT RECONSTRUCTION

Woven and knitted cotton fabrics are two-dimensional periodic lattices formed by orthogonal warp (0-degree) and weft (90-degree) yarns or diagonal twill ribs. [S06] [S11] [S32] When a plain convolutional autoencoder is trained directly on raw pixel intensities of defect-free cloth, two failure modes arise on the shop floor. [S06] [S12] First, if the autoencoder's bottleneck capacity is too wide, the network reconstructs small point defects, fine broken picks and subtle oil spots alongside the normal weave, causing defect escapes. [S12] [S0604] Second, if the fabric roll experiences slight mechanical skew or rotation on the inspection rollers, a raw pixel autoencoder interprets the shifted phase of the weave grid as a large reconstruction error across the frame, flooding the operator console with false alarms. [S06]

To overcome both sensitivity collapse and rotational fragility, Liu et al. formulated a three-stage one-class fabric defect detection architecture that precedes deep latent learning with a multi-scale, multi-orientation Gabor filter bank. [S06] Gabor filters are orientation-tuned and frequency-tuned linear band-pass filters established in spectral textile inspection for isolating directional yarn periodicity. [S06] [S32] Convolving every incoming fabric image with Gabor kernels spanning multiple spatial frequencies and angular orientations before feeding the response maps into a convolutional autoencoder converts raw pixel grids into directional energy envelopes. [S06] Normal periodic warp and weft interlacing produces stationary filter responses across the frame, whereas a local weave disruption, slub, knot or tear creates a directional energy spike across specific orientation channels. [S06]

6.9 NEAREST-NEIGHBOUR DENSITY ESTIMATION FOR ZERO-FALSE-ALARM THRESHOLDING

Rather than relying solely on pixel-wise mean squared reconstruction error at the output of the convolutional autoencoder, the Gabor-autoencoder architecture feeds the learned latent representations into a non-parametric nearest-neighbour density estimator to locate potential defects on the fabric image. [S06] During one-class training on defect-free cloth, the autoencoder compresses the multi-channel Gabor responses into a low-dimensional latent manifold where normal plain, patterned and rotated fabric patches cluster densely. [S06] At inspection time, patches whose latent representations fall in low-density exterior regions far from their k nearest nominal neighbours are flagged and outlined on the operator display. [S06]

Tested across plain weaves, patterned weaves and rotated fabric samples derived from the Standard Fabric Defect Glossary, this Gabor-autoencoder-density pipeline achieved a true positive rate (recall) of 0.895 (89.5 percent) with zero false alarms (0.0 percent false-positive rate) while training exclusively on defect-free samples. [S06] For a greenfield Cameroonian inspection perch, the zero-false-alarm characteristic reported in that benchmark is operationally significant: when an edge cascade

forwards too many false alarms (as seen in uncalibrated plain autoencoders where up to 85 percent of forwarded frames are false positives), human inspectors suffer alarm fatigue and begin ignoring the monitor. [S06] [S07]

6.10

Reverse knowledge distillation and single-sample motif extraction handle patterned textiles

6.11 ASYMMETRIC TEACHER-STUDENT BOTTLENECKS PREVENT ANOMALY RECONSTRUCTION

When the factory line shifts from unbleached greige cloth or solid-dyed twill to multi-colour stripes, checks and patterned apparel fabrics, neither simple pixel autoencoders nor symmetric student-teacher networks suffice. [S12] [S0604] In standard knowledge distillation for anomaly detection, a frozen pretrained teacher encoder and a trainable student encoder of identical architecture are fed the same raw input image; during nominal training, the student learns to mimic the teacher's multiscale activation maps, and at test time anomalies are detected wherever the student's activations diverge from the teacher's. [S12] [S0604] As Deng and Li demonstrated in their CVPR study on reverse distillation, using symmetric encoder architectures creates a vulnerability: because both networks share identical receptive fields and data flow paths from the raw image upward, the student network often generalizes its convolutional filters to anomalous regions as well, muting the representation discrepancy needed to flag weak defects. [S0604]

Reverse knowledge distillation solves this leakage by inverting the topology into an asymmetric teacher-encoder, one-class bottleneck embedding (OCBE), and student-decoder pipeline. [S12] [S0604] Instead of feeding the raw camera frame to both networks, only the frozen pretrained teacher encoder receives the input image and extracts multiscale spatial feature maps. [S0604] Those teacher features are compressed through the trainable OCBE module into a compact code that preserves normal fabric structure while stripping out anomaly perturbations. [S0604] The student network, structured as an asymmetric decoder, receives only this compact bottleneck embedding and targets restoration of the teacher's multiscale representations from high-level abstract features down to low-level features. [S0604] Because the student decoder never sees raw image pixels directly and must pass through a one-class bottleneck trained on nominal cloth, it fails to reconstruct unseen weave flaws, dye stains or print misregistrations. [S12] [S0604]

In 2024, researchers adapted reverse knowledge distillation specifically to patterned fabric defect detection, highlighting high-level periodic pattern consistency alongside thread-level cues. [S12] Evaluated across MVTec AD, AITEX, TILDA and a dataset acquired directly from an operating textile manufacturing facility spanning diverse colours and weaves, the fabric-specific reverse distillation detector delivered strong anomaly localisation accuracy and fast inference speed without defective training images. [S12]

6.12 DYNAMIC SINGLE-SAMPLE MOTIF PERIOD EXTRACTION IN A SINGLE TRAINING EPOCH

For factories running short-batch patterned or printed fabrics where collecting 50 nominal training frames or running multi-epoch gradient descent is impractical, dynamic motif-based convolutional networks push cold-start efficiency to a single defect-free fabric sample trained in one epoch. [S0602] Introduced for unsupervised patterned textile defect detection, this five-stage methodology executes preprocessing, automatic pattern period extraction, motif patch extraction, dynamic heuristic feature selection and anomaly detection directly from the input domain. [S0602]

Rather than initializing convolutional filter banks with random Gaussian weights and iterating through dozens of backpropagation epochs, the algorithm computes the horizontal and vertical repetition periods of the elementary weave or print motif from one clean sample image. [S0602] It slices the nominal sample into motif-aligned patches and constructs the network's filter sizes, filter counts and weights dynamically from the dominant spatial structures of the input motif, requiring only the number of layers and convolution stride beforehand. [S0602] Validated on the Patterned Fabrics benchmark dataset, this single-sample, single-epoch unsupervised CNN yields competitive precision, recall, accuracy and F1-score compared with multi-epoch unsupervised models at lower computational cost. [S0602] On a Cameroonian line switching between uniform checks and printed shirting, single-sample motif extraction adapts the vision rig to a new pattern repeat almost immediately. [S12] [S0602]

6.13

Controllable anomaly synthesis: GLASS and feature-space noise bridge unsupervised and supervised detection

6.14 FEATURE-SPACE GAUSSIAN PERTURBATIONS IN SIMPLENET AND MANIFOLD-CONSTRAINED SYNTHESIS IN GLASS

Pure one-class models that see only nominal data face a boundary problem: without negative examples during training, the decision boundary around the normal feature manifold can remain loose, allowing weak defects that resemble normal cloth variations (such as faint shade marks, tight reed lines or single-filament tension drops) to hide inside the nominal boundary. [S13] [S0605] Conversely, naive synthetic strategies that paste random natural-image textures onto clean fabric images generate out-of-distribution structural artefacts that look unlike real loom or dyehouse flaws. [S13]

Two architectures resolve this trade-off by synthesizing controllable, near-in-distribution anomalies in feature space and image space: SimpleNet and the Global and Local Anomaly co-Synthesis Strategy (GLASS). [S13] [S0605] SimpleNet uses four components: a pretrained Feature Extractor, a shallow Feature Adapter that transfers local features toward the target industrial domain, an Anomaly Feature Generator that counterfeits anomaly vectors by adding Gaussian noise to normal features, and a binary Anomaly Discriminator. [S0605] Perturbing normal features with additive Gaussian noise trains a sharp discriminator boundary without real defect samples, and discarding the generator at test time allows SimpleNet to reach 99.6 percent AUROC on MVTec AD (cutting error by 55.5 percent over the next best model) at 77 frames a second on a desktop GPU. [S0605]

Building on feature perturbation, Chen et al. developed GLASS to detect weak textile and industrial defects that sit close to normal texture regions. [S13] GLASS combines Global Anomaly Synthesis (GAS) at the feature level with Local Anomaly Synthesis (LAS) at the image level under manifold and hypersphere distribution constraints. [S13] Using Gaussian noise guided by gradient ascent and truncated projection, GLASS pushes synthesized points just outside the normal hypersphere so pseudo-anomalies mimic subtle production flaws. [S13] GLASS achieved 99.9 percent detection AUROC on MVTec AD and was validated in industrial applications for woven fabric defect detection. [S13]

6.15 PHYSICS-INFORMED SYNTHETIC WEAVE DISRUPTIONS FOR CAMEROONIAN COTTON CONSTRUCTIONS

When a greenfield factory in Cameroon wishes to pre-train a YOLOv8 or YOLOv11 bounding-box detector before a new SODECOTON cotton style has accumulated real production defects, GLASS feature co-synthesis can be paired with physics-informed image-level defect rendering on the five-metre nominal enrolment strip. [R03] [S11] [S13] Because the optical geometry of Station A (incoming roll line-scan perch with transmitted backlight and grazing darkfield LED bars) is calibrated and fixed, the software can inject synthetic warp breaks (0-degree single-yarn intensity attenuation streaks), synthetic weft slubs (90-degree localized yarn thickening with simulated darkfield shadow edges) and synthetic dye spots (CIELAB delta-E Gaussian blobs) onto clean nominal frames while GLASS perturbs the latent feature manifold. [S11] [S13] [S14]

To help engineering teams navigate these competing algorithmic families, **Exhibit 28**, p. 97 summarises the input requirements, mathematical mechanisms, published benchmark results and shop-floor trade-offs of the seven primary cold-start anomaly detection families. [S04] [S05] [S06] [S12] [S13] [S0601] [S0602] [S0605]

EXHIBIT 28 / PUBLISHED BENCHMARK METRICS AND ARCHITECTURAL ATTRIBUTES

Technical comparison of cold-start anomaly detection families for textile inspection

Architecture family	Nominal training input	Core mathematical mechanism	Published benchmark performance	Best textile floor fit
PaDiM (Defard et al.)	50-200 clean frames	Pretrained CNN patch embeddings + multivariate Gaussian + Mahalanobis distance	Outperforms prior SOTA on MVTec AD	Fast constant-time inference on aligned plain and twill rolls
PatchCore (Roth et al.)	5-200 clean frames	Mid-level neighbourhood features + greedy coreset memory bank + NN distance	99.6% image AUROC on MVTec AD	Day-one cold-start enrolment across plain, dyed and textured weaves
Gabor-CAE + Density	Defect-free glossary set	Multi-orientation Gabor filter bank + convolutional autoencoder + k-NN density	89.5% recall, 0.0% false alarms	Rolls prone to mechanical skew or rotation requiring zero false alarms
Reverse Distillation	Nominal texture dataset	Asymmetric teacher encoder + one-class bottleneck + student decoder	Validated on MVTec, AITEX, TILDA, mill	Multi-colour patterned weaves and complex dyed constructions
Dynamic Motif CNN	1 defect-free sample	Automatic motif period extraction + input-driven filter weights in 1 epoch	Competitive F1 in single epoch	Rapid style changeovers on yarn-dyed checks, stripes and prints
SimpleNet / GLASS	Nominal frames + noise	Feature adapter + gradient-ascent manifold Gaussian noise + discriminator	99.6% to 99.9% AUROC on MVTec AD	Weak, low-contrast weave flaws and 77 FPS high-speed web screening
RegLoss Self-Supervised	Unlabelled fabric frames	Progressive pixel rearrangement + autocorrelation texture regularity loss	67.7% greige, 63.9% denim, 70.3% MVTec P-AP	Pixel-accurate defect mask segmentation for ASTM 4-point length grading

SOURCED [S04] [S05] [S06] [S12] [S13] [S0601] [S0602] [S0605] Compiled from opened primary sources: PaDiM (S05), PatchCore (S04), Gabor-CAE + density estimator (S06), Reverse Knowledge Distillation for fabrics (S12, S0604), Dynamic Motif CNN (S0602), SimpleNet (S0605), GLASS (S13) and RegLoss self-supervised segmentation (S0601).

The synthesis in **Exhibit 28**, p. 97 clarifies that no single cold-start algorithm dominates every operational constraint simultaneously. [S04] [S06] [S12] [S13] [S0601] [S0602] While PatchCore and GLASS provide near-ceiling image-level recall (99.6 percent and 99.9 percent AUROC on MVTec AD), Dynamic Motif CNNs minimize changeover training to a single defect-free image and one epoch, and Gabor-CAE architectures prioritize zero false alarms on skewed rolls. [S04] [S06] [S13] [S0602] A modular edge software stack keeps both a memory-bank scorer (PatchCore/PaDiM) and a distilla-

tion or motif-period module inside the station container, selecting the execution graph according to whether the incoming batch is a solid-dyed cotton twill or a small-lot printed motif. [S04] [S05] [S12] [S0602]

6.16

Self-supervised texture regularity loss enables pixel-level segmentation without manual masks

6.17 PROGRESSIVE PIXEL REARRANGEMENT AND AUTOCORRELATION REGULARITY CONSTRAINTS

Detecting that a 400 by 400 millimetre camera frame contains an anomaly is sufficient to flash an andon light or queue a thumbnail for operator triage, but automated roll disposition under the ASTM D5430 four-point grading system requires knowing the physical length of the flaw in millimetres. [R29] [S08] [S34] As established in Chapter 2, the four-point system assigns 1 penalty point to defects up to 3 inches (76.2 mm), 2 points from 3 to 6 inches (76.2 to 152.4 mm), 3 points from 6 to 9 inches (152.4 to 228.6 mm), and 4 points above 9 inches or for any through-hole, rejecting rolls that exceed an illustrative threshold of 40 points per 100 square yards. [R29] [S08] Because patch-level embeddings in standard CNN backbones downsample spatial resolution by a factor of 8 or 16, raw Mahalanobis or coreset distance heatmaps produce dilated blobs that overestimate the physical width of a 0.20 millimetre broken warp yarn. [S04] [S15] [S0601]

A 2026 study on self-supervised fabric defect segmentation via texture regularity learning addresses this pixel-level resolution gap without manually painted segmentation masks. [S0601] Existing unsupervised methods largely neglect explicit modelling of texture regularity, an inherent periodic property fundamental to woven and knitted fabrics. [S0601] The architecture introduces two coupled mechanisms inside an end-to-end reconstruction-segmentation dual network: [S0601]

1. **Progressive pixel rearrangement pretext task:** During self-supervised training on nominal fabric images, the pipeline degrades local periodic weave regularity through a progressive pixel rearrangement schedule and compels a reconstruction network to restore the regular texture representation. [S0601]
2. **Texture Regularity Loss (RegLoss):** Rather than supervising reconstruction solely with pixel intensity error, RegLoss quantifies periodicity strength via local autocorrelation analysis and constrains reconstruction quality so the network restores only valid periodic lattice structures. [S0601]

6.18 CONVERTING PIXEL-LEVEL ANOMALY MAPS INTO FOUR-POINT PENALTY LENGTHS

When a test frame containing a real weave flaw passes through this dual reconstruction-segmentation network, the reconstruction branch restores the regular background cotton lattice while erasing the defect, and the segmentation branch localizes anomalies at the pixel level. [S0601] Evaluated on two self-built datasets (plain grey fabric and denim) and on the public MVTec AD benchmark, the RegLoss self-supervised method attains pixel-level Average Precision (P-AP) of 67.7 percent on plain

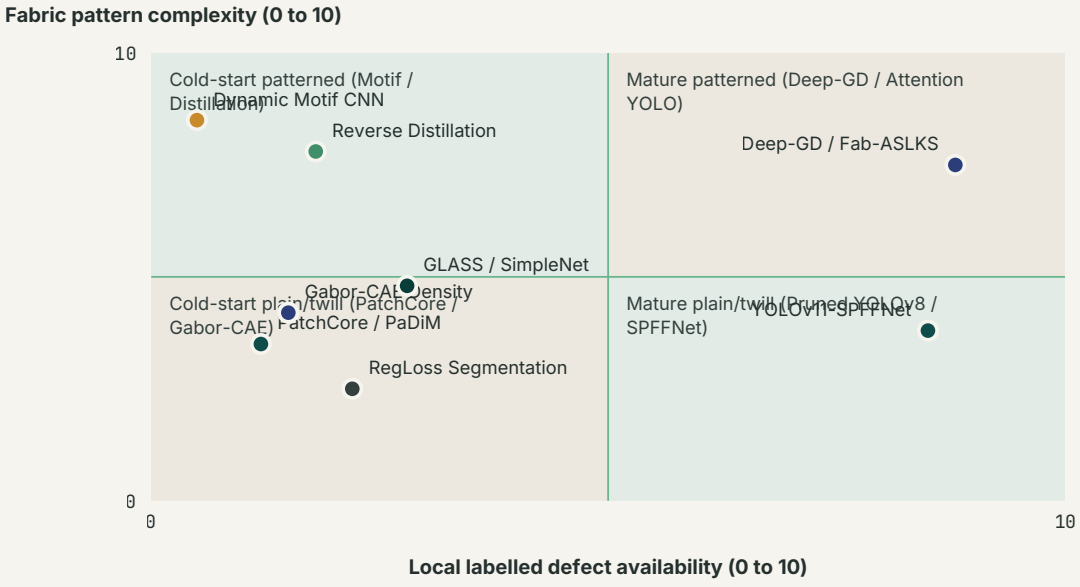
grey fabric, 63.9 percent on denim, and 70.3 percent on MVTec AD, outperforming representative reconstruction-based and feature-distribution-matching baselines. [S0601]

Once the segmentation head outputs a binary pixel mask on the edge GPU, converting the mask into an objective ASTM D5430 penalty score or a cutting-table avoidance polygon requires connected-component geometry. [S08] [S14] [S0601] Given a calibrated rotary shaft encoder pulsing at an illustrative 0.15 millimetres per line-scan row and a cross-web camera pixel pitch of 0.15 millimetres per pixel, a connected component spanning 620 pixels along the warp axis measures 620 multiplied by 0.15 millimetres, which equals 93.0 millimetres (3.66 inches), logging a 2-point major warp defect into the Performance-Management System's Quality Log with its longitudinal metre coordinate and cross-web millimetre offset. [R29] [S14] [S15]

To codify when plant engineers should deploy one-class memory banks, motif-period extractors, synthetic anomaly discriminators or supervised YOLO detectors, **Exhibit 29**, p. 99 maps each architecture across local annotated defect sample availability and background fabric pattern complexity. [S04] [S06] [S11] [S12] [S13] [S0602]

EXHIBIT 29 / DERIVED OPERATIONAL POSITIONING SCORE (0 TO 10 SCALE)

Architecture selection across local defect sample volume and background pattern complexity



DERIVED [S04] [S05] [S06] [S11] [S12] [S13] [S0601] [S0602] Derived decision matrix synthesized from training-data dependencies and texture benchmarks in S04 (PatchCore), S05 (PaDiM), S06 (Gabor-CAE), S11 (YOLOv11-SPFFNet), S12 (Reverse Distillation), S13 (GLASS), S0601 (RegLoss) and S0602 (Dynamic Motif CNN). Horizontal axis scales from zero local defect images (0) to 1,000+ boxes per class (10); vertical axis scales from plain greige weave (0) to complex prints (10).

As **Exhibit 29**, p. 99 illustrates, the four quadrants define an operational progression across the factory lifecycle. [S04] [S11] [S12] [S0602] A greenfield line inspecting unbleached or solid-dyed SODECOTON cotton starts in the lower-left quadrant with PatchCore, PaDiM, Gabor-CAE or RegLoss segmentation; when a short-run print or yarn-dyed check arrives with zero defect history, the station switches to the upper-left quadrant using Dynamic Motif CNN or Reverse Knowledge Distillation; and once

six to twelve months of verified defect crops have accumulated in the Quality Log, the factory graduates its staple weaves into the right-hand quadrants using pruned YOLOv8/YOLOv11-SPFFNet or Deep-GD multi-class classifiers. [R03] [S04] [S06] [S11] [S12] [S0602]

6.19

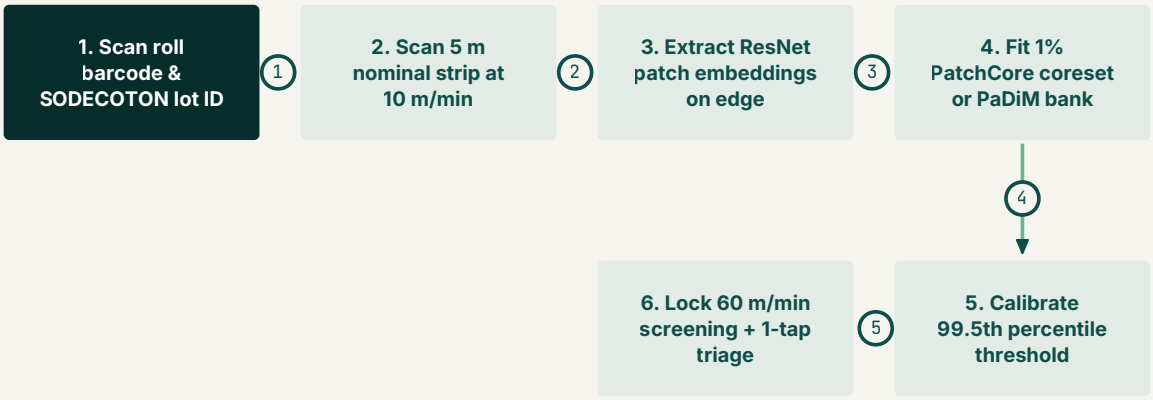
The GFT cold-start protocol: enrol a new fabric roll in fifteen minutes

6.20 STEP-BY-STEP FIFTEEN-MINUTE NOMINAL ENROLMENT ON THE INCOMING PERCH FRAME

To make one-class anomaly detection usable by Cameroonian quality technicians on shift one without a machine-learning specialist at the perch frame, Green Future Tech Limited standardises cold-start roll onboarding into a fifteen-minute shop-floor protocol. [R29] [S04] [S05] [S07] Whenever a new fabric construction, dye shade or supplier lot arrives at Station A, the operator executes the six-step enrolment workflow in **Exhibit 30**, p. 100. [R03] [S04] [S05] [S07]

EXHIBIT 30 / PROCESS SEQUENCE

Fifteen-minute cold-start enrolment workflow from nominal fabric scan to active anomaly triage



1. Min 0-2
2. Min 2-5 (150 clean frames)
3. Min 5-10
4. Min 10-13
5. Min 13-15 (Go live)

DERIVED [R03] [R29] [S04] [S05] [S07] [S34] Derived operational workflow integrating SODECOTON lot intake (R03), 60 metre a minute line-scan optics (R29), PatchCore/PaDiM nominal patch embedding and coreset subsampling (S04, S05), Jetson edge cascade thresholding (S07) and human-in-the-loop touchscreen triage (S34).

Each stage in **Exhibit 30**, p. 100 enforces physical and statistical checks: [R03] [R29] [S04] [S05] [S07]

1. **Lot identification and optical preset selection (Minutes 0 to 2):** The operator scans the roll barcode, linking the roll in the Performance-Management System to its upstream SODECOTON ginning batch, and selects the fabric family preset to set LED strobe current and camera exposure. [R03] [S14]
2. **Supervised five-metre nominal strip scan (Minutes 2 to 5):** The perch roller advances the first 5.0 linear metres of the roll at 10 metres a minute while the quality inspector watches the backlit and brightfield zones to confirm zero weaving flaws, stains or splices across the 150 captured full-width frame strips. [R29] [S04] [S14] [S15]
3. **Frozen-backbone feature extraction (Minutes 5 to 7):** The edge compute node (an NVIDIA Jetson AGX Orin or Xavier module) runs the 150 nominal frames through its frozen WideResNet-50 or ResNet-18 TensorRT FP16 engine to extract levels 2 and 3 patch feature vectors. [S02] [S04] [S07]
4. **Coreset subsampling or Gaussian fitting (Minutes 7 to 11):** For solid weaves and twills, the edge box executes greedy minimax coreset subsampling (1 percent PatchCore bank) or fits PaDiM Gaussians; for repeating prints, it runs the single-epoch Dynamic Motif period extractor or Reverse Distillation bottleneck. [S04] [S05] [S12] [S0602]
5. **Holdout threshold calibration (Minutes 11 to 13):** The software evaluates distance distributions on the final 1.0 metre of the nominal strip (held out from coreset construction) and sets the Stage 1 anomaly trigger at the 99.5th percentile of nominal distances, verified against synthetic GLASS/SimpleNet Gaussian perturbations. [S04] [S07] [S13] [S0605]
6. **Production release at up to 60 metres a minute (Minutes 13 to 15):** The memory bank is hashed, version-stamped in local NVMe storage, and locked into active screening mode up to 60 metres a minute. [R29] [S07]

METHOD

System boundary and cold-start memory-bank sizing equations

To verify that a PatchCore or PaDiM nominal reference bank fits inside embedded GPU memory and executes within real-time line-scan budgets, GFT sizes the cold-start feature bank using four equations anchored on published parameters: [R29] [S04] [S05] [S07]

- **Web tile rate:** For a web width $W = 1,600\text{ mm}$ inspected at $v = 60\text{ m/min}$ ($1,000\text{ mm/s}$) using two 4k line-scan cameras (800 mm each, pixel pitch $p = 0.20\text{ mm/pixel}$), each camera acquires $5,000\text{ lines/s}$, generating an illustrative 78.1 tiles/s of size 512×512 pixels at 60 m/min or 26.0 tiles/s at 20 m/min. [R29] [S14] [S15]
- **Raw nominal patch count (N_{raw}):** With a ResNet-18 backbone at stride 8, each 512×512 tile yields $4,096$ patch embeddings of dimension $D = 512$; scanning an illustrative 5.0 m leader strip (392 tiles) yields $1,605,632$ raw patch vectors. [S04] [S05]
- **Coreset memory footprint (M_{coreset}):** Applying PatchCore greedy coreset reduction at ratio $\gamma = 0.0025$ retains $N_c = 4,014$ vectors, which in `FP16` (2 bytes per value) occupies $4,014 \times 512 \times 2\text{ bytes} = 4.11\text{ megabytes}$ of GPU memory on an NVIDIA Jetson Orin module. [S02] [S04] [S07]
- **Active-learning harvest:** Every frame exceeding the anomaly threshold is cropped for one-tap operator confirmation (`Confirm Defect + Class` or `Dismiss False Alarm`), building a clean YOLOv8/YOLOv11 training set during normal shifts. [S07] [S34]

WORKED EXAMPLE: SIZING COLD-START ANOMALY TRIAGE AND ACTIVE-LEARNING DATASET ACCUMULATION FOR A 3.0 MILLION GARMENT A YEAR PLANT

Consider an illustrative greenfield cut-make-trim facility in Cameroon scaled to produce 3.0 million cotton garments a year across 300 operating days (10,000 garments a day), consuming an estimated 1.50 linear metres of 1.60-metre-wide SODECOTON-origin cotton fabric per garment, or 15,000 linear metres per day (300 rolls per day at 50 metres per roll). [R03] [R14] [R29]

1. **Enrolment overhead per new fabric batch:** Suppose those 300 daily rolls arrive in 6 distinct weave-and-shade lots of 50 rolls (2,500 linear metres) each. Running the fifteen-minute cold-start protocol (**Exhibit 30**, p. 100) on the first roll of each lot takes $6 \times 15\text{ minutes} = 90\text{ minutes}$ per day across two perch frames (45 minutes per frame per shift, or 9.4 percent of shift time), while the remaining 49 rolls per lot reuse the cached coreset bank with zero pause. [S04] [S05]
2. **Cascade forwarding and triage load:** Under S07, a recall-prioritised Stage 1 anomaly detector flags 100 percent of defective frames (20 of 20) while tighter threshold calibration cuts Stage 2 invocations by 29 to 45 percent. [S07] Assuming an illustrative incoming defect density of 12 defects per 100 linear metres (1,800 true defective frames per day out of 146,484 frame strips), calibrating PatchCore/Gabor-CAE to 92.0 percent recall (capturing 1,656 true defects versus 1,080 to 1,260 manually) and a 1.5 percent false-positive rate (2,170 false alarms) sends $3,826$ thumbnails per day across two perch frames over two shifts, or 2.0 one-tap verifications per minute per operator. [R29] [S04] [S06] [S07] [S34]
3. **Time to supervised YOLOv11 fine-tuning:** Logging 1,656 verified defect crops and 2,170 hard-negative patches daily into the Quality Log yields 49,680 positive defect crops and 65,100 hard negatives in 30 operating days, providing roughly 994 verified bounding boxes even for a 2 percent minority defect class and crossing the threshold to fine-tune YOLOv8s or YOLOv11-SPFFNet for Phase 2 without offline annotation labour. [S07] [S09] [S11] [S34]

GFT NOTE

Use one-class anomaly models to bootstrap, not to replace, multi-class defect taxonomy. Unsupervised cold-start models (PatchCore, PaDiM, Gabor-CAE and Reverse Distillation) tell the operator *where* the cloth is abnormal on day one when zero defect history exists. [S04] [S05] [S06] [S12] However, a pure distance score does not tell the monthly Sales and Operations Planning (S&OP) meeting whether the deviation is a spinning nep, a loom broken pick, or a dyehouse oil stain. [R29] [S04] [S18] Every frame flagged by the cold-start detector must therefore prompt a one-tap defect family code from the operator on the touchscreen, ensuring the Performance-Management System's Quality Log receives root-cause categories from shift one while building the labelled dataset for Phase 2 supervised fine-tuning. [R29] [S07] [S18] [S34]

> **WARNING**
>
> **Operational constraint: nominal strip contamination and optical drift.** If the five-metre leader strip scanned during roll enrolment contains an unnoticed continuous warp fault or shade band, PatchCore and PaDiM will encode that flaw into the nominal memory bank and remain blind to identical defects across the lot. [S04] [S05] Similarly, cotton fly accumulation on the LED bar or viewport shifts test embeddings away from the nominal Gaussian mean, raising false-positive rates. [S07] [S15] Three controls resolve this: (1) dual-inspector sign-off on the five-metre leader strip; (2) cross-patch variance outlier rejection during coreset construction; and (3) an automated pixel-histogram check every 30 minutes against the white reference tile to trigger air-knife cleaning or exposure compensation. [S04] [S07] [S15]

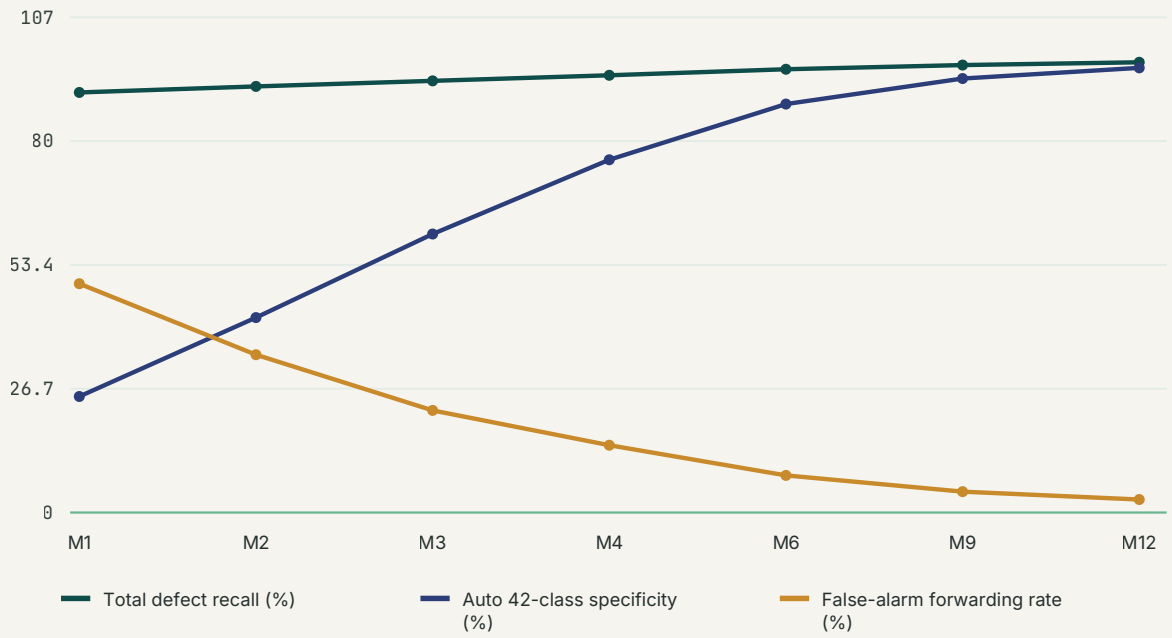
6.21 GRADUATING FROM PHASE 1 UNSUPERVISED ANOMALY TRIAGE TO PHASE 2 MULTI-CLASS YOLO FINE-TUNING

The advantage of pairing unsupervised cold-start vision with Deliverable 2's Performance-Management System lies in how inspection accuracy and class specificity compound over the first twelve months. [R29] [S04] [S07] [S34] During Month 1 (Phase 1 launch), PatchCore or PaDiM carries the primary detection load on incoming SODECOTON cotton rolls, capturing above 90 percent of anomalies while routing flagged frames to human operators for one-tap verification and 42-class tagging. [R03] [R29] [S04] [S05] [S07] During Months 2 and 3, accumulated hard-negative patches dismissed by operators are injected into a SimpleNet or GLASS discriminator layer ahead of the memory bank, cutting false-positive forwarding rates by 29 to 45 percent. [S07] [S13] [S0605] By Months 4 to 6, when the Quality Log holds tens of thousands of verified defect crops, the team fine-tunes Stage 2 YOLOv8s or YOLOv11-SPFFNet weights on local cotton textures, shifting the station to direct 42-class classification with automated ASTM D5430 point scoring while keeping PatchCore/Distillation active in parallel to catch novel faults and enrol seasonal prints. [R29] [S04] [S07] [S11] [S12]

The trajectory in **Exhibit 31**, p. 105 models this twelve-month operational handoff, plotting illustrative defect recall, multi-class automated classification specificity, and operator false-alarm triage load as the factory graduates from day-one cold start to Phase 2 fine-tuned operation. [R29] [S04] [S07] [S11] [S34]

EXHIBIT 31 / PERCENT OF INSPECTED FRAMES OR DEFECTS

Illustrative twelve-month transition from Phase 1 one-class anomaly triage to Phase 2 fine-tuned multi-class detection



ILLUSTRATIVE [R29] [S04] [S07] [S11] [S34] Illustrative operational trajectory anchored on cited benchmark endpoints: manual baseline recall (65 percent midpoint of 60 to 70 percent, R29); day-one PatchCore/cascade recall (>90 percent, R29, S04, S07); initial recall-prioritised cascade false-positive forwarding rate (49.3 percent in S07) dropping via threshold calibration and hard-negative active learning (29 to 45 percent inference reduction in S07); and automated 42-class specificity rising as local verified defect crops fine-tune YOLOv8/YOLOv11 models (S11, S34).

As **Exhibit 31**, p. 105 shows, separating *anomaly recall* from *multi-class specificity* resolves the tension between day-one deployment and long-term factory automation. [R29] [S04] [S07] [S11] Total defect capture jumps above 90 percent in Month 1 because PatchCore and Gabor-CAE require only five metres of nominal cotton to flag departures from normal weave structure, protecting the cutting room from the 20 to 30 percent defect escape rate of manual inspection. [R29] [S04] [S06] While automated 42-class specificity starts at an illustrative 25.0 percent in Month 1 and the recall-prioritised cascade initially forwards up to 49.3 percent false alarms on ambiguous patches, active learning drives false-alarm forwarding below 5.0 percent by Month 9 while lifting automated multi-class classification above 95.0 percent. [S07] [S11] [S34] What Exhibit **Exhibit 31**, p. 105 cannot establish on its own is the exact month-by-month slope for a specific Cameroonian mill, which depends on style changeover frequency, dyehouse shade stability and supervisor discipline during daily huddles. [R29] [S07] [S34]

By combining fifteen-minute nominal roll enrolment at the incoming fabric perch with active-learning capture into the Performance-Management System, a greenfield garment factory eliminates the cold-start excuse completely. [R29] [S04] [S07] The fact-

ory does not wait for defects to happen before it starts measuring quality; it measures what good cotton looks like on the first five metres of cloth, and uses every anomaly encountered thereafter to make both the weave and the model sharper. [R03] [S04] [S06]

A greenfield factory does not need a history of failures to recognise a flaw; five metres of good cloth is enough to teach the camera what belongs in the weave.

CHAPTER SOURCE NOTES

S04 arXiv / IEEE CVPR. 2021-06-15.

S05 arXiv / ICPR. 2020-11-17.

S06 arXiv / Computer Science and Information Technology. 2022-04-20.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

S01 Sensors (MDPI). 2025-04-03.

S11 arXiv. 2025-02-03.

S08 International Journal for Research in Applied Science and Engineering Technology. 2023-05-31.

S12 arXiv. 2024-01-04.

S13 arXiv. 2024-07-12.

S35 arXiv. 2025-05-11.

S09 arXiv. 2025-01-24.

S07 arXiv. 2026-08-12.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

S0603 International Journal of Computer Vision (Springer). 2021-01-06.

S34 IntechOpen. 2025-01-03.

S0602 arXiv. 2023-11-30.

R24 Mordor Intelligence. Publication date unrecorded; inherited source-text check.

R26 Statista. Publication date unrecorded; inherited source-text check.

S0601 Elsevier BV (SSRN). 2026-10-09.

S0605 arXiv / IEEE CVPR. 2023-03-27.

S32 Security and Communication Networks. 2021-05-10.

S0604 arXiv / IEEE CVPR. 2022-01-26.

S14 Applied Sciences (MDPI). 2024-12-23.

S15 Sensors (MDPI). 2013-04-09.

S02 Sensors (MDPI). 2021-10-25.

S18 International Journal of Electrical and Electronics Research. 2024-02-05.

07

Extending computer vision from flat fabric webs to sewn garments catches stitch, seam and dimensional faults in 0.45 seconds

Inspect the defects created by assembly as well as those inherited from cloth. The reported 0.45-second result concerns image processing, while garment handling and dimensional acceptance require separate proof. [S19]

CHAPTER 07
DECISION TEST

Inspect the defects created by assembly as well as those inherited from cloth. The reported 0.45-second result concerns image processing, while garment handling and dimensional acceptance require separate proof. [S19]

Computer vision should follow the transformation from cloth into a garment, because an acceptable roll cannot establish that its cut panels, seams and finished dimensions are acceptable. The engineering answer is a sequence of inspection decisions, each tied to the operation that can create or expose a fault. The title's 0.45-second figure is reported image-processing time from a trouser-classification study, rather than demonstrated elapsed time for detecting every stitch, seam and dimensional fault in a complete garment. This chapter proposes how to extend coverage without confusing an image result with a production guarantee. [S19]

7.1

Incoming cloth inspection cannot release the garment created from it

The inherited Cameroon sector record reports approximately 95 percent of cotton fibre exported raw and approximately 5 percent processed locally. Its original publication date was not recorded. Those proportions explain the interest in local transformation; they do not describe the performance of a GFT factory or identify a garment quality system already in operation. The recommendations below concern a proposed cut-make-trim inspection arrangement, with incoming cloth treated as an input rather than evidence that downstream assembly will conform. [R06]

GFT recommends organising the quality plan around what changes physically. Receipt establishes the condition of the material presented for cutting. Cutting establishes panel identity and geometry. Sewing establishes the intended joins. Finishing establishes the appearance and dimensions presented to the buyer. Each release should answer a bounded question about a recognisable workpiece. A single generic label such as good garment would hide which operations and surfaces were actually observed.

Start the design with an inspection coverage sheet. For each product family, record the critical construction features, the view needed to see them, the operation after which they become visible, and the person authorised to decide an exception. Distinguish front, reverse and concealed surfaces in plain language. A checked exterior should never imply a checked interior. If the construction hides an important join, propose an earlier inspection before closure and retain the operation reference in the finished item's record.

The proposed topology in **Exhibit 32**, p. 109 links these decisions. It is an illustrative reference arrangement, not an implemented GFT line. Arrows indicate the intended passage of material and information. They do not prove that every fault will be visible, that handling will fit production time, or that a station can safely interrupt a sewing machine.

EXHIBIT 32 / ILLUSTRATIVE PROCESS SEQUENCE

Inspection decisions should follow the operations that change the workpiece



1. Roll identity
2. Panel identity
3. Seam history

ILLUSTRATIVE GFT proposed arrangement for a cut-make-trim line. Nodes are inspection responsibilities, not a verified installation or proof of complete surface coverage.

Treat the topology as a proposal to test with production staff. Walk an actual construction sequence with the quality lead before specifying cameras. Ask where a rejected item can be held without losing its identity, where an operator can present an obscured face, and where a confirmation action would interrupt safe handling. The best location is the earliest useful decision point at which the relevant feature can be presented consistently and a practical response remains available.

> **GFT NOTE**
>
> Commission coverage before commissioning speed. Require a signed list of visible features, excluded features, product families and exception owners. Keep manual inspection available for excluded surfaces and uncertain results until local qualification supports a different release rule.

GFT also recommends a vocabulary that separates fault, observation and disposition. A dark region is an observation; a stain is a reviewed interpretation; rework or reject is an authorised disposition. Store these separately. Otherwise a model's provisional guess becomes a permanent quality allegation, and the factory loses the ability to distinguish a camera problem from an actual assembly problem when reviewing recurring events.

7.2

Cutting-table vision needs a verified coordinate chain before it can protect panels

For cutting, GFT proposes carrying the incoming defect map forward as a traceable material record. The proposed record should identify the roll, inspected face, travel direction, distance origin and transverse origin. Cutting staff should confirm the orientation after spreading. A coordinate without that boundary is insufficient: a correct

location in an inspection image may refer to a different physical position after rewinding, trimming or turning the cloth. This is a design requirement, not a claim about an installed projection system.

Do not begin with a visually impressive marker overlay. Begin with a physical registration trial. Mark a harmless reference on a sacrificial length, record it at receipt and follow it through the intended material route. At the cutting table, compare its observed position with the transferred position. Repeat after the handling operations the factory actually intends to use. The production engineer should define an acceptable registration discrepancy in relation to the clearance needed around defects, rather than accept an arbitrary pixel error.

Keep the transfer rule explicit. A longitudinal coordinate may require an offset for trimmed material and a reversal when the roll orientation changes. A transverse coordinate may require mirroring when the observed face changes. Stretch, skew and uncertain layer alignment need separate consideration. GFT recommends recording the transformation and its verification status alongside the original map, so an invalid mapping can be withdrawn without erasing the original observation.

An illustrative coordinate example assumes a mapped fault at 12.40 metres from the original origin, 0.30 metres removed before spreading, and unchanged direction. Its proposed table position is 12.40 minus 0.30, or 12.10 metres. An assumed exclusion margin of 0.05 metres would reserve 12.05 to 12.15 metres longitudinally before considering the fault's own extent. These values demonstrate bookkeeping only. They are neither measured registration accuracy nor an instruction to cut production cloth.

GFT recommends making splice treatment explicit in the same trial. A mapped exclusion can protect a marker only if the spreading team knows which layers contain the affected region and how a splice changes usable material. Record the instruction given to the cutter, the material actually set aside and any panels subsequently replaced. Do not record all mapped exclusions as avoided waste. The useful result is the reconciled outcome after cutting, with the excluded material and replacement work visible.

Where a proposed projection device is used, test the displayed outline against physical references across the usable table, including its edges. An accurate central alignment should not stand in for an unchecked outer region. Have the operator confirm that the projected instruction is visible under working illumination and remains associated with the current lay. If the camera or projector moves, withdraw the previous registration approval until the intended verification is repeated.

The panel check should have its own reference, preferably the approved pattern revision and an agreed presentation method. Compare the intended perimeter, not a loosely similar silhouette. Specify whether notches, drill marks, grain orientation and paired-panel identity are in scope. An outline that looks correct can still belong to the wrong size or revision. GFT recommends making identity confirmation a prerequisite for geometric comparison and displaying the revision clearly at the station.

For a proposed multi-layer lay, do not infer that an overhead view verifies concealed panels. Agree whether the method checks the upper surface, selected separated panels or every panel after separation. Record that choice in the inspection scope. If

the workflow cannot preserve individual panel identity, limit the release statement to the bundle and sampling boundary actually observed. The purpose is to avoid a precise-looking contour drawing becoming an unsupported claim about hidden material.

The outline's proposed pattern-defect performance figure is not reproduced here because its publisher page could not be opened in this session. No verified result in this chapter establishes automatic splice projection or cut-panel metrology at production speed. Those functions remain proposed engineering tasks, with acceptance dependent on a local coordinate trial, agreed presentation and a comparison against the approved physical pattern.

7.3

Sewing-station vision should follow the visible stitch path and retain human disposition

Published airbag work demonstrates a framework that follows seam paths and checks stitch colour patterns to identify skipped stitches, missed stitches and superimposed seams in lockstitch and double-thread chainstitch constructions. This is relevant evidence for inspecting structured joins. It is not evidence that an apparel installation inherits the same performance on a different thread, fabric, seam construction or camera view. [S22]

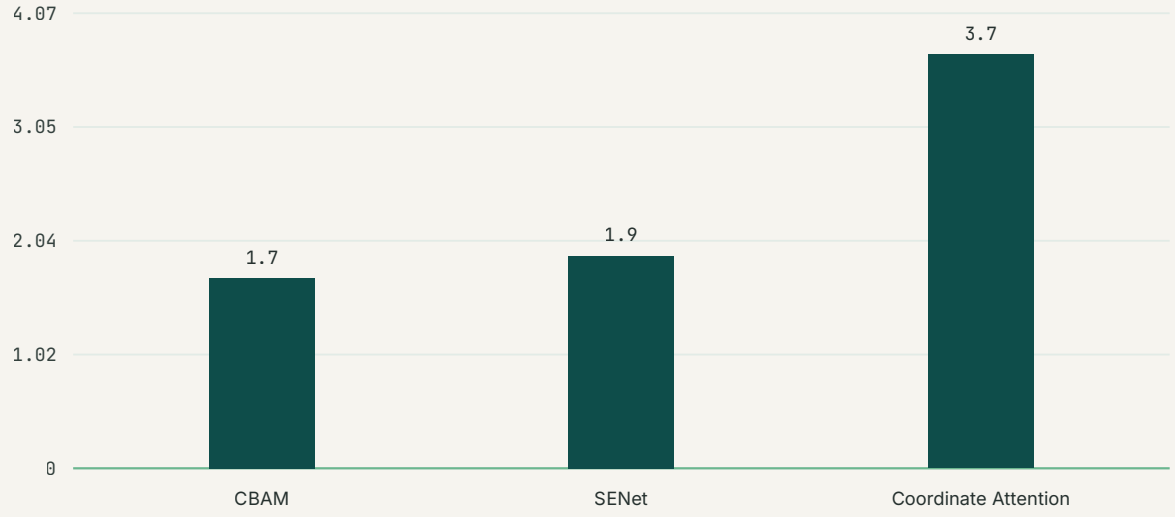
For an apparel pilot, GFT recommends selecting a critical seam whose exit from the sewing operation can be presented without requiring the operator to defeat guards or reposition hands unsafely. Define the observed seam segment relative to the operation and garment identity. Keep the camera trigger, sample association and review image connected. An alert that cannot identify which item and which seam segment it concerns offers little useful control once the workpiece leaves the station.

The proposed image-validity gate should precede defect classification. Check whether the expected region is visible, the image is sufficiently sharp for the selected stitch feature, and the presentation belongs to the approved product recipe. If visibility is inadequate, return an inspection-unavailable state. Do not encode a missing image as no defect. The quality lead should determine whether unavailable coverage requires re-presentation, manual review or a production hold for that particular construction.

The garment-sewing attention study reports improvements in mAP@0.5 of 1.7 percentage points with CBAM, 1.9 with SENet and 3.7 with Coordinate Attention in its comparison. Its proposed YOLOv8-FPCA method also changes the detection head and loss treatment. The ablation in **Exhibit 33**, p. 112 preserves a common metric; it does not compare those gains with classification accuracy from unrelated datasets. [S17]

EXHIBIT 33 / PERCENTAGE-POINT IMPROVEMENT IN MAP@0.5

Coordinate Attention gives the largest reported gain in this sewing ablation



SOURCED [S17] Reported attention comparisons in the opened primary article. Values are changes in a detection metric on that study's data, not absolute accuracy or a forecast for a factory.

For procurement, request a reproducible comparison on the factory's held-out seam examples before treating an architectural choice as valuable. Keep the capture conditions, test garments and annotation rules fixed between model candidates. Ask the supplier to provide the complete detection results, including misses and false alarms, rather than a hand-picked image gallery. GFT recommends evaluating the operating threshold against the planned review workload and the consequences of escapes, not choosing it solely to maximise a summary metric.

Use a conservative initial response. Show the suspected location, preserve the item identifier and route the item for confirmation. A repeated alert on the same seam segment should not generate repeated defective-garment counts. Conversely, a garment with several distinct faults should retain each reviewed fault. Agree the event-merging rule with the quality lead and keep the original images available for diagnosing whether merged observations concerned the same physical defect.

A machine-stop function needs a separate machinery assessment and qualified integration. The quality model should not be granted direct control merely because its demonstration runs quickly. For the initial proposed deployment, GFT recommends an operator-visible advisory and controlled hold process. Define who can clear the hold, what evidence is needed, and how a station resumes after a camera obstruction or product change. These are proposed operating controls rather than assertions of achieved reliability.

7.4

Seam-puckering assessment needs controlled presentation and an agreed grading reference

The opened seam-puckering abstract describes seamline location using edge detection and a Hough transform, followed by rotation, cropping, projection-profile extraction and LOESS smoothing. Extracted features feed a neural classifier for grading. The study supports a route towards repeatable appearance assessment. Its public abstract does not establish the spectral-analysis method suggested elsewhere in the outline, so that separate method is not attributed to it here. [S21]

GFT recommends treating a puckering station as a measurement procedure, not simply another defect camera. Define the seam length presented, garment support, lighting arrangement, pressing state and permitted operator handling. Retain a reference image of acceptable presentation. If an operator smooths the fabric differently on each pass, an apparent improvement in the score may describe presentation rather than a change in sewing quality. The qualification exercise should deliberately repeat placement as well as image capture.

The proposed grading reference should come from the buyer's approved appearance requirements and a recorded set of boundary examples. Ask reviewers to label the same retained samples independently and reconcile disagreements before training or selecting thresholds. Preserve the disagreement record. A model trained on unresolved labels may reproduce that disagreement while displaying a single confident grade. The practical objective is consistent decisions against an agreed requirement, with borderline cases still available for human assessment.

Avoid declaring a cause from appearance alone. For a confirmed puckering event, the mechanic should inspect the relevant sewing settings and materials under the plant's approved troubleshooting procedure. The camera record can show the location and recurrence of the appearance; it does not itself prove that thread tension, feed setting or needle condition caused it. GFT recommends separating the observed symptom from the mechanic's diagnosis and from the corrective action tried.

For the pilot, retain before-action and after-action samples under the same presentation procedure. Compare reviewed outcomes and repeat the placement check. If the appearance changes only when the operator handles the item differently, do not record a process improvement. If the correction holds across independently presented items, document the scope of that evidence. This creates a useful engineering record without turning a local adjustment into a general claim about all fabrics or seams.

7.5

Finished-garment evidence supports local trials rather than universal accuracy promises

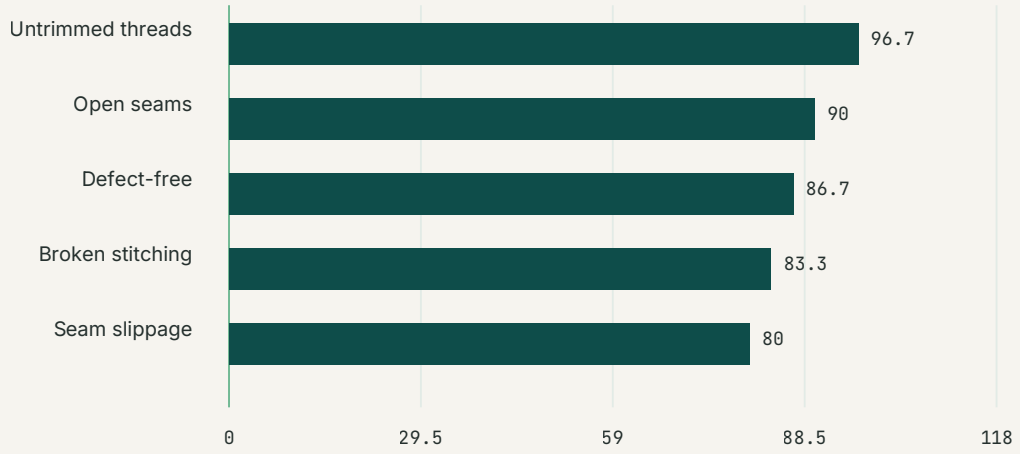
Deep-GD reports 97.16 percent overall classification accuracy across 12 categories using HKBU imagery, bilateral filtering, a Squeeze-and-Excitation network and a Bayesian-optimised Random Decision Forest. Its categories include fabric-texture

and stain conditions. That result does not establish complete inspection of assembled garment interiors or dimensional conformity. Use it as an example of a classification approach, with its dataset boundary retained. [S18]

The Ethiopian trouser study reports 1,000 collected images, a VGGNet-19 feature extractor with an artificial neural network, and 86.7 percent held-out accuracy. The reported processing time is 0.45 seconds per sample. The per-class percentages in **Exhibit 34**, p. 114 are reproduced as the abstract's reported performance; their labels are not silently converted into defect recall, precision or customer escape probability. [S19]

EXHIBIT 34 / PERCENT, REPORTED PER-CLASS PERFORMANCE

Reported trouser performance varies materially by class



SOURCED [S19] Publisher abstract for the Addis Ababa trouser study. Class denominators and confusion matrix were not available in the opened abstract. These values are not a factory acceptance guarantee.

The public abstract does not expose the complete confusion matrix, device specification or elapsed handling cycle. GFT therefore recommends obtaining those details before relying on this evidence for equipment sizing. Keep photographs of the same physical garment and closely related production bursts together when partitioning a local trial. This proposed partition rule reduces the chance that the evaluation rewards recognition of repeated material or presentation rather than successful inspection of unseen production.

Before collecting a local evaluation set, GFT recommends agreeing how disputed examples will be decided. Ask reviewers to inspect the physical item where it remains available, because a photograph may omit the very feature being disputed. Retain an uncertain label when a reliable decision cannot be made. An unresolved example can test the review process without being treated as a confirmed positive or negative in headline model metrics.

Keep the trial sequence independent of supplier demonstrations. Prepare a sealed evaluation list, record the approved software and threshold, and retain all outputs before reviewing aggregate performance. If a revised model is proposed after inspecting failures, use fresh held-out material for the next acceptance decision.

Retain the earlier result as development evidence. This proposed procedure makes improvement visible while preventing repeated tuning on the final evaluation set from masquerading as independent proof.

For the local trial, propose a qualification sheet covering product family, size, colour, surface, defect definition, presentation and reviewer adjudication. Include acceptable items as well as retained faulty examples. Record which classes are represented inadequately and exclude them from automatic release. Select acceptance criteria before examining the final trial results. Otherwise the team can keep adjusting the definition of success until an attractive result appears, without improving the actual release decision.

An illustrative trial ledger might contain 200 confirmed defective items and 800 confirmed acceptable items, with 180 defects flagged and 40 acceptable items flagged. On these assumed counts, recall is 180 divided by 200, or 90 percent. Precision is 180 divided by 220, or approximately 81.8 percent. The acceptable-item false-alarm rate is 40 divided by 800, or 5 percent. These invented counts explain denominators, not expected GFT performance.

GFT recommends reviewing that illustrative ledger as a workload as well as an algorithm result. The review desk would receive 220 flagged items and the nominally clear group would still contain 20 defective items. An automatic-release proposal must confront that residual risk explicitly. Add no-read and missing-image cases to the operating ledger rather than excluding them from the denominator simply because no model score exists.

> **WORKED EXAMPLE: IMAGE LATENCY LEAVES HANDLING TIME UNRESOLVED**
>
> An illustrative serial cell assumes 5.00 seconds to identify and present a garment, 0.45 seconds for image processing taken from the published trouser study, and 3.00 seconds to review the result and move the item. Assuming no overlap, total cycle time is 8.45 seconds. Dividing 3,600 seconds by 8.45 gives approximately 426 items an hour before breaks, re-presentation or downtime. The handling assumptions are illustrative; the published processing result alone does not establish this cell capacity. [S19]
>
> For an illustrative sensitivity case, assume presentation rises to 8.00 seconds with all other components unchanged. The serial total becomes 11.45 seconds and the same division gives approximately 314 items an hour. These are arithmetic capacities under assumed handling, not measured throughput. Use timed local work to replace both presentation assumptions before deciding how many stations a line needs. [S19]

The worked example defines a serial boundary deliberately. If a supplier proposes overlapping presentation and computation, request a timed demonstration of the actual queue and item association. Count re-presentation, result acknowledgement and physical diversion in the trial. A quick processor offers little practical benefit when unresolved items accumulate faster than the reviewer can handle them. GFT recommends retaining time-stamped events for each transition so the limiting activity can be identified directly.

7.6

Dimensional release requires uncertainty small enough for the buyer's tolerance

Garment measurement is a separate claim from defect classification. The opened university record describes segmentation with UNet and key-point extraction; it reports average measurement errors of 1.27 centimetres for blazers, 0.747 for dresses and 1.012 for skirts. Those results concern photographed garments and do not validate a proposed calibrated flat-lay station. They also give no basis for promising sub-millimetre repeatability on a production table. [S0701]

GFT proposes a fixed measurement plane, a verified scale reference and an approved garment-presentation procedure for a flat-lay trial. Define each measurement by its physical endpoints and construction state. Body length, chest width and sleeve length need written conventions that the buyer and quality lead understand in the same way. Resolve whether the garment is relaxed, fastened, pressed or otherwise conditioned before comparing camera measurements with the tech pack.

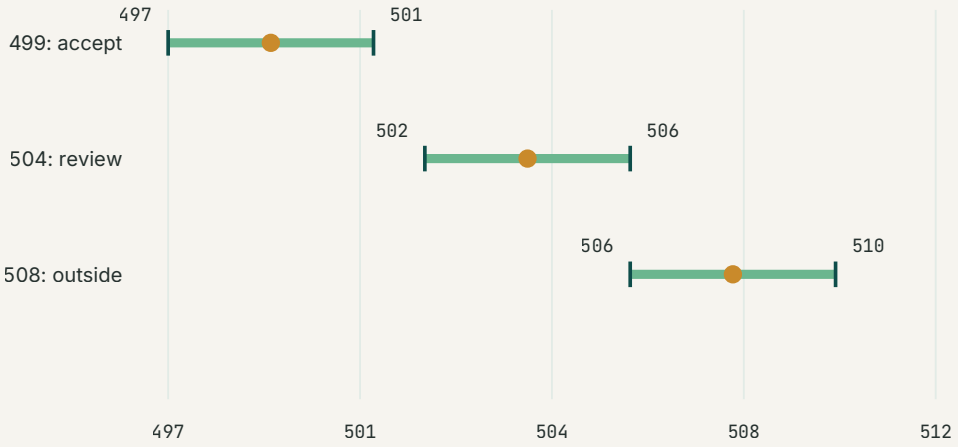
The proposed measurement budget should include scale calibration, lens effects, landmark choice, fabric placement and repeated operator handling. Pixel density alone is insufficient as an acceptance argument. Use a stable reference artefact to examine the optical mapping, then repeated garments to examine the full procedure. Keep these exercises separate in the record: a successful reference check cannot establish repeatability when a soft garment is removed and laid down again.

> **METHOD**
>
> For an illustrative dimensional rule, assume a nominal chest width of 500 millimetres and buyer limits of 495 to 505 millimetres. Assume a locally justified conservative uncertainty bound of 2 millimetres. Accept automatically only when the full measurement interval lies inside the buyer limits, giving a guarded reading band of 497 to 503 millimetres. The assumed bound must be established by the intended measurement procedure before use; it is not a demonstrated camera specification.

The illustrative example in **Exhibit 35**, p. 117 shows why an apparently in-tolerance reading can still require review. Each interval represents an assumed reading plus or minus the same conservative bound. It is a decision illustration, not a confidence interval derived from observations. No actual buyer specification, station accuracy or GFT production measurement is asserted.

EXHIBIT 35 / ILLUSTRATIVE CHEST WIDTH, MILLIMETRES

An uncertain measurement near a tolerance boundary should be reviewed



ILLUSTRATIVE Assumed buyer limits 495 to 505 mm. Each assumed reading has a conservative plus or minus 2 mm bound. A full interval inside the limits qualifies for automatic acceptance under the illustrated rule; overlap requires review.

Under that illustrative rule, the reading of 504 millimetres spans 502 to 506 and overlaps the upper buyer limit. Route it for agreed remeasurement rather than automatic acceptance. The reading of 508 spans 506 to 510 and lies wholly outside the assumed limits. That distinction prevents uncertain overlap from being confused with clear evidence of dimensional nonconformity. Neither state should be decided by a sewing-defect classifier's confidence score.

GFT recommends a measurement-system trial in which reviewers repeat placement and measurement without seeing their earlier answers. Compare the camera procedure with the agreed reference method and examine persistent bias separately from scatter. Investigate results by measurement point and product family. An average across easy and difficult landmarks can conceal an unusable point, just as a successful overall defect score can conceal a poorly detected class.

Calibration recovery also needs an explicit proposed procedure. Record the last successful reference check and identify production assessed since that point if the next check fails. The quality lead should decide the affected release boundary from the evidence available, rather than simply recalibrating and erasing the failed state. Preserve the original readings, the failed check and the recovery decision. This makes a calibration failure a traceable quality event with an accountable disposition.

For the operator interface, GFT recommends displaying the measured dimension, applicable limits and review reason together. Keep unavailable, borderline and outside-limit states visibly different. Allow a reviewer to record a repeat measurement without overwriting the initial result. If the revised outcome changes disposition, retain who authorised the change and which method resolved it. The resulting history should explain the decision to a production supervisor without requiring access to model internals.

Release only the dimensions that the qualification actually supports. A contour may support an external width while leaving a folded collar or internal construction unmeasured. Preserve a per-point result and a reason for any unavailable measurement. If the buyer changes the measurement convention or tolerance, review the release rule before reusing the previous recipe. The record should identify the approved specification revision alongside the calibration and presentation versions.

7.7

Needle engineering should close the loop from repeated symptoms to verified correction

Vision should support preventive engineering by making repeated symptoms inspectable. GFT recommends connecting confirmed stitch faults to the machine, operation, material recipe and maintenance event. Use the connection to ask a focused question: did the symptom persist after the approved intervention under comparable conditions? Avoid describing correlation with a needle change as proof that wear caused every fault in the cluster.

The proposed maintenance procedure should retain the needle specification, change reason, responsible role and verified restart sample. Follow the machine and needle supplier's applicable instructions when setting replacement intervals and adjustment methods. No universal needle life is specified here. A replacement programme should be justified for the actual materials and operations, with scheduled checks and symptom-triggered review documented as different reasons for intervention.

After a change, require the quality lead to inspect the restart output using the agreed method. Record both a successful restart and a failed restart. A cleared dashboard alarm is not proof that the garment is acceptable, and absence of further camera alerts is insufficient if the view has become obstructed. The proposed restart gate should therefore confirm image availability as well as the reviewed construction feature.

GFT recommends separating maintenance learning from individual performance scoring. Begin with the physical fault, construction and process conditions. Ask operators what changed and retain their explanation as an attributed observation, rather than automatically converting a camera event into blame. A pilot depends on candid reporting of presentation difficulties, borderline examples and workarounds. Those reports are useful engineering inputs even when the model appears confident.

For commissioning, propose distinct sign-offs for visibility, detection, handling and correction. The quality lead owns defect definitions and release scope. Production owns presentation and material movement. Maintenance owns safe mechanical intervention. The digital lead owns event identity, model version and reproducible results. Assign an accountable owner for resolving conflicts between these boundaries. This proposed responsibility split prevents a successful software demonstration from substituting for an accepted factory procedure.

A proposed daily review should reconcile material movement with inspection events. Examine items awaiting confirmation, items held after maintenance and items whose identity was lost during re-presentation. An unresolved hold must not disappear

because a new shift signs in or a garment is photographed again. Assign a physical holding location and an accountable role, then close the record only when the reviewed item has a documented disposition.

The handover package should include the approved recipes, retained reference examples, unresolved class gaps, response instructions and restart checks. State where manual review remains necessary and what evidence would permit an expanded scope. Keep power resilience and wider information-system design with their dedicated engineering workstreams; this chapter's acceptance boundary is whether the right garment feature is observed, judged and acted upon in time to support production.

The practical recommendation is to pilot a visible critical seam and a bounded finished-item check, measure the handling burden, and expand only when the evidence supports the next feature. A fast image result is useful within that sequence. It becomes a production control only when identity, coverage, disposition and corrective action are verified together. That is how measurement can remain attached to the material as cloth becomes a garment.

The seam is controlled when the signal changes the next decision.

CHAPTER SOURCE NOTES

S19 SAGE, Textile Research Journal. 2026-09-06.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

S22 arXiv, author-deposited manuscript. 2015-10-25.

S17 JVE International, Journal of Measurements in Engineering. 2024-09-09.

S21 SAGE, Textile Research Journal. 2016-10-14.

S18 FOREX Publication, International Journal of Electrical and Electronics Research. 2024-02-05.

S0701 MDPI, Applied Sciences; author record at Kaunas University of Technology. 2022-04-28.

08

Running vision inference at the edge keeps quality control alive when factory internet links drop or grid power fluctuates

In Cameroon's Douala–Dibamba corridor or northern Garoua belt, a cloud-dependent camera is a single point of failure; edge inference makes vision self-contained.

CHAPTER 08
DECISION TEST

In Cameroon's Douala-Dibamba corridor or northern Garoua belt, a cloud-dependent camera is a single point of failure; edge inference makes vision self-contained.

Running computer-vision defect detection on embedded edge processors mounted directly at the inspection frame is the only architecture that guarantees continuous quality control in a greenfield Cameroonian garment factory. [R29] In Green Future Tech Limited's Cameroon textile programme (prepared 31 January 2024, with sources refreshed in 2026), the digital strategy is designed explicitly around two physical site constraints: intermittent electrical power and variable wide-area connectivity. [S0801] A vision pipeline that streams raw camera frames to an external cloud server stalls whenever the wide-area link congests or drops, blinding quality assurance while fabric rollers and sewing lines continue to run. [S07] By executing pre-trained neural networks locally on compact industrial accelerators and transmitting only lightweight pass-or-fail dispositions, defect class codes and cropped anomaly thumbnails to the on-premise data store, the factory decouples real-time inspection from external network availability. [S02] Coupled with inline direct-current battery buffering, read-only operating filesystems and store-and-forward database queues, edge inference ensures that neither a telecommunication outage nor a grid brownout can allow defective cloth to escape into the cutting and sewing rooms. [S14]

8.1

Cloud-streamed video fails under Cameroonian bandwidth and power constraints

High-speed optical textile inspection generates a continuous raw data stream that cannot be transported reliably over commercial broadband connections in Sub-Saharan Africa. [S07] Automated fabric roll inspection operates at web speeds of up to 60 metres a minute (1,000 millimetres a second) across more than 40 distinct defect types, lifting capture accuracy from a manual baseline of 60 to 70 percent to above 90 percent. [R29] Resolving single-thread warp and weft breaks across a 1,600-millimetre fabric width requires either dual 4,096-pixel line-scan cameras scanning 800 millimetres each or an equivalent contact image sensor array operating at a spatial pitch of 0.20 millimetres per pixel. [S15] At 1,000 millimetres a second, a 0.20-millimetre pixel pitch demands a line-trigger frequency of 5,000 lines a second (5 kilohertz), which is an illustrative optical calculation from cited line speeds. [R29]

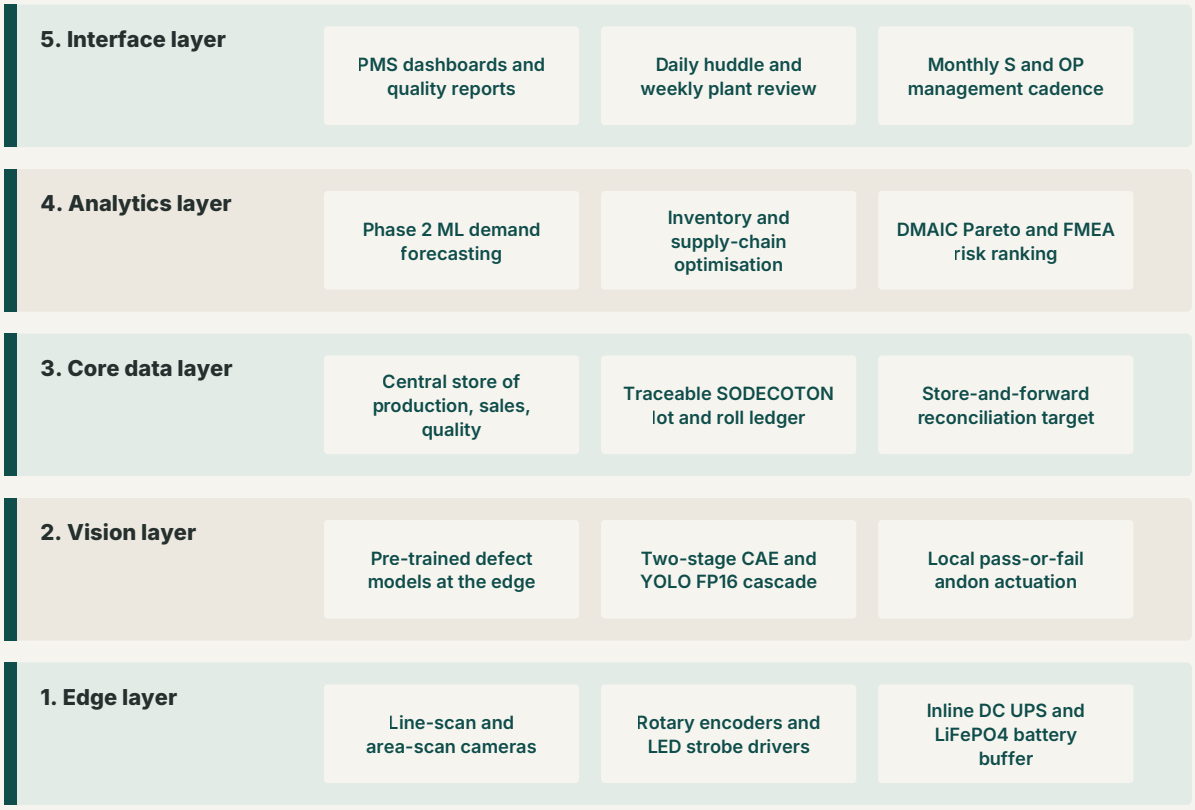
When captured as uncompressed 8-bit red-green-blue (RGB) colour imagery to preserve the 6.7 percentage-point classification advantage that colour channels hold over greyscale on dyed and printed cloth (86.7 percent against 80.0 percent in published garment benchmarks), a single dual-camera fabric perch produces 8,192 pixels multiplied by 3 bytes multiplied by 5,000 lines per second. [S19] That single perch generates 122.88 megabytes a second, or 983.04 megabits a second of sustained raw video throughput, as an illustrative bandwidth derivation from cited sensor dimensions. [S15] Adding two cutting-table gantry cameras and eight sewing-line area-scan stations processing garments in 0.45 seconds per sample pushes aggregate raw camera throughput past 1.2 gigabits a second across an illustrative 3.0 million garment a year cut-make-trim plant. [S19]

Attempting to backhaul a 1.2 gigabit a second video stream to a remote cloud region is incompatible with Cameroon's telecommunications and electrical baseline. [S0802] According to World Bank World Development Indicators data refreshed in October

2026, fixed broadband subscriptions in Cameroon stood at 2.13 per 100 people in 2022, down from 2.23 in 2021 and 2.76 in 2020. [S0802] Industrial parks outside central business districts face asymmetric uplink bandwidth, variable packet latency and periodic fibre cuts. [S0802] On the electrical side, national access to electricity reached 72 percent of the population in 2023 and 2024 (up from 64.3 percent in 2020), yet electric power transmission and distribution losses absorbed 24.82 percent of total generated output across 2022, 2023 and 2024. [S0801] High grid losses reflect feeder overloading, transformer trips and voltage sags that routinely interrupt unbuffered routers, optical network terminals and factory switchgear. [S0803]

EXHIBIT 36 / ARCHITECTURAL LAYERS

GFT five-layer reference architecture keeps vision inference at the edge and surfaces insights in the PMS



SOURCED [R29] [S07] [S14] Implements the five-layer Reference Architecture from Green Future Tech Limited's Cameroon textile programme (Digital Strategy, Section 5). Vision inference executes locally at Layer 2 so quality control survives wide-area network drops.

To eliminate wide-area bandwidth dependency, Green Future Tech Limited's Reference Architecture structures factory information flow into five distinct layers, shown in **Exhibit 36**, p. 122. [S14] At the bottom sits the Edge layer, comprising the line-scan cameras, area-scan sewing cameras, rotary shaft encoders and controlled light-emitting diode (LED) illumination bars. [S15] Immediately above it sits the Vision layer, which executes pre-trained defect detection models on embedded processors mounted within two metres of the camera head. [S07] Because inference happens at

the Vision layer before any network traversal, the 983.04 megabit a second raw optical stream never leaves the inspection frame's shielded camera cable, based on our illustrative station calculation. [S14]

Instead of streaming video, the Vision layer emits a compact JavaScript Object Notation (JSON) event record containing the timestamp, fabric roll or garment bundle identifier, station code, pass-or-fail flag, defect class code, bounding-box pixel coordinates and, when a flaw is detected, a compressed 64-kilobyte thumbnail crop for supervisor verification. [S07] That architectural boundary compresses network traffic by more than four orders of magnitude, from nearly 1 gigabit a second of raw pixels to an illustrative average of less than 75 kilobits a second of structured metadata per perch. [S07] Those lightweight records land in the Core data layer (an on-premise relational store of production, quality and sales data), which feeds the Analytics layer (Phase 2 demand forecasting and inventory optimisation, where AI at scale cuts forecast errors by 20 to 50 percent and inventory costs by 35 to 40 percent) and the Interface layer (the Performance-Management System dashboards). [R27]

8.2

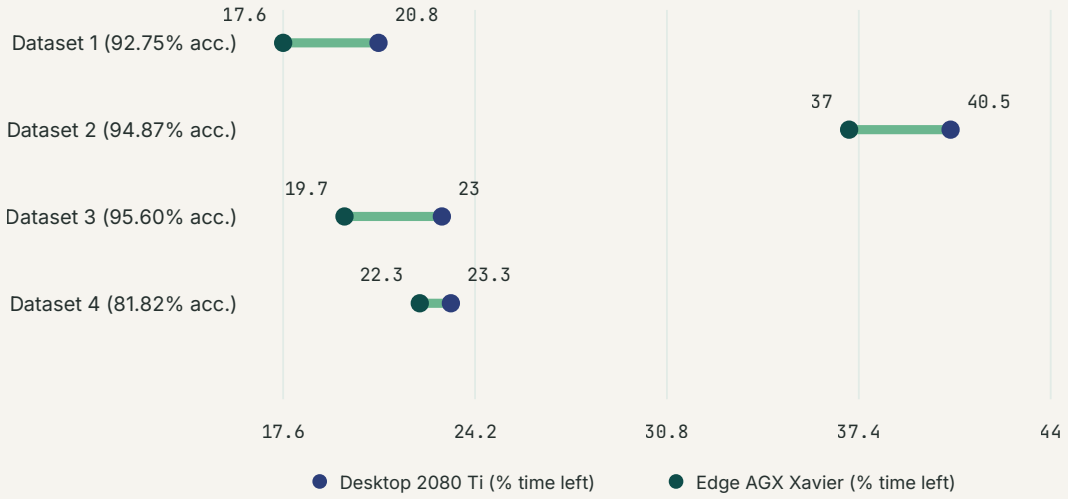
Bayesian network pruning and quantisation fit deep vision models onto embedded accelerators

Executing deep convolutional and transformer vision models inside a 10-to-30-watt thermal envelope at the inspection frame requires systematic model compression so that edge latency stays below physical line-speed deadlines. [S02] Uncompressed two-stage and single-stage object detectors trained on workstation graphics processing units (GPUs) contain tens of millions of redundant 32-bit floating-point (FP32) parameters that saturate the memory bandwidth and thermal ceiling of embedded system-on-module (SoM) boards. [S01] If an edge processor takes 250 milliseconds to evaluate a fabric frame while the roller advances at 60 metres a minute, the web moves 250 millimetres between evaluations, creating blind inspection gaps. [R29]

Published textile-inspection research demonstrates that structured Bayesian network pruning and reduced-precision quantisation remove parameter redundancy while preserving above-90-percent defect classification accuracy. [S02] In an optimization study by Xie et al. (2021), researchers trained a modified convolutional neural network on four textile defect datasets and applied Bayesian pruning to eliminate low-salience filters before compiling the pruned graph for an embedded NVIDIA Jetson AGX Xavier module. [S02] Across the four test datasets, Bayesian pruning reduced inference time on a desktop NVIDIA GeForce RTX 2080 Ti GPU to 20.84 percent, 40.52 percent, 23.02 percent and 23.33 percent of the unpruned baseline duration. [S02] When deployed onto the embedded NVIDIA Jetson AGX Xavier accelerator, the same pruned networks reduced inference time even further, to 17.56 percent, 37.03 percent, 19.67 percent and 22.26 percent of the unpruned embedded execution time. [S02]

EXHIBIT 37 / PERCENT OF UNPRUNED INFERENCE TIME

Bayesian network pruning cuts inference time to 17.56 to 37.03 percent on embedded AGX Xavier hardware



SOURCED [S02] Compares remaining inference latency as a percentage of the unpruned baseline model across four fabric inspection datasets on desktop GPU (GeForce RTX 2080 Ti) and embedded edge SoM (NVIDIA Jetson AGX Xavier). Post-processing classification accuracies reached 92.75%, 94.87%, 95.60% and 81.82% across Datasets 1 to 4 respectively (S02).

Crucially, as illustrated in **Exhibit 37**, p. 124, compressing the network by nearly four-fifths did not degrade inspection reliability below industrial thresholds on standard weaves. [S02] After coupling the pruned convolutional backbone with morphological post-processing, detection accuracy across the four datasets reached 92.75 percent, 94.87 percent, 95.60 percent and 81.82 percent respectively, with the three standard woven datasets comfortably clearing the 90 percent automated benchmark. [S02] The greater relative speedup observed on the embedded AGX Xavier (down to 17.56 percent of baseline latency on Dataset 1 compared with 20.84 percent on the desktop card) occurs because pruning shrinks intermediate activation tensors so they fit inside the embedded module's on-chip cache and shared low-power memory bus, avoiding external memory stalls. [S02]

8.3 REDUCED-PRECISION TENSORRT COMPILATION DOUBLES THROUGHPUT ON ENTRY-LEVEL EDGE MODULES

Beyond structural filter pruning, compiling trained PyTorch or Open Neural Network Exchange (ONNX) weights through NVIDIA TensorRT into 16-bit floating-point (FP16) or calibrated 8-bit integer (INT8) kernels activates hardware tensor cores on embedded Jetson modules. [S07] Converting weights and activations from 32-bit floating point to 16-bit floating point halves memory footprint and doubles arithmetic throughput with negligible loss in mean Average Precision (mAP) across YOLOv5, YOLOv8 and YOLOv11 architectures. [S01] On entry-level edge boards designed for developing-economy mills, such as the NVIDIA Jetson Nano (operating within a 5-to-10-watt power envelope), TensorRT FP16 compilation allows a complete two-stage inspection pipeline to sustain 13.45 frames a second in continuous operation. [S07] On higher-tier industrial modules such as the NVIDIA Jetson AGX Orin (configurable from 15 to 60 watts), the same TensorRT-optimised detector executes in approximately 32 milliseconds end-to-end per frame (over 31 frames a second), providing ample headroom for 60 metre a minute fabric roll line-scan inspection. [S07]

For stationary garment inspection stations at the end of sewing lines, compute requirements are even more modest. [S19] Empirical factory trials on a 1,000-image trouser dataset in Addis Ababa demonstrated that a VGGNet-19 feature extractor paired with an artificial neural network classifier processes each garment image in 0.45 seconds on standard local hardware, achieving 86.7 percent overall accuracy across broken stitching, open seams, seam slippage and untrimmed threads while running 60 times faster than manual visual checks. [S19] When migrated to a fanless Jetson Orin Nano at the pressing and packing table, a 0.45-second inference cycle easily outpaces a human operator placing garments on the flat-lay frame every 6 to 10 seconds, as an illustrative station ergonomics benchmark. [S19]

8.4

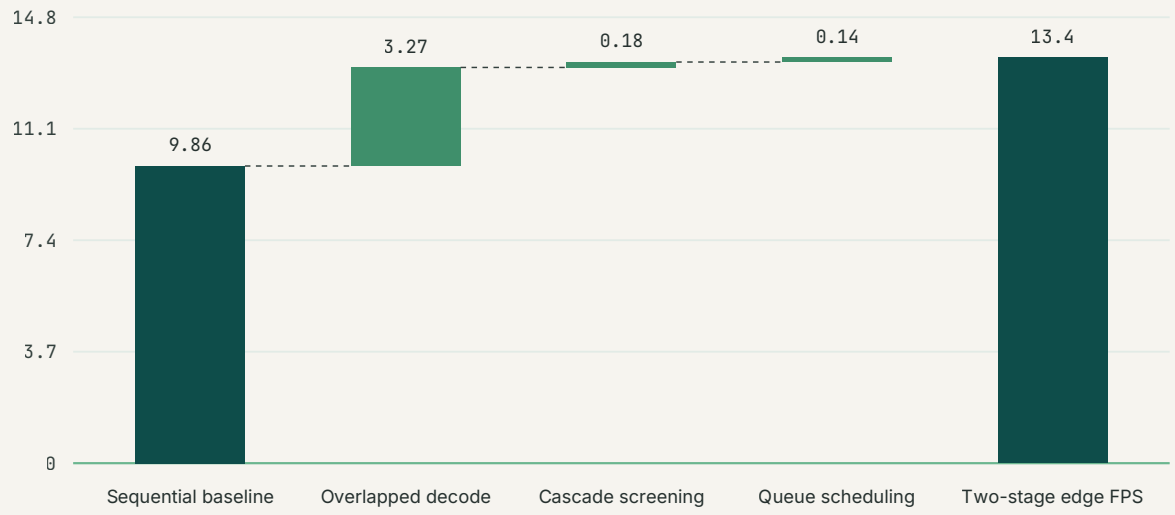
Two-stage edge cascades filter nominal fabric before invoking heavy detectors

Even after pruning and FP16 quantisation, running a multi-scale object detector such as YOLOv8s or YOLOv11 on every single camera frame wastes electrical power and compute cycles because the vast majority of fabric passing under the lens is defect-free. [S07] In a well-controlled garment and textile plant, overall product defect rates are managed toward an industry tolerance of 2 to 5 percent, and even in unoptimised printing lines where failures spike, defective pieces represent around 7.75 percent of output (107 defective units out of 1,380 inspected garments in published Indonesian factory data). [S24] On a continuous fabric roll perch, more than 90 to 95 percent of captured frames contain only regular nominal warp-and-weft texture. [S26]

To exploit defect sparsity at the edge, recent low-cost inspection engineering pairs an ultra-lightweight first-stage anomaly screener with a second-stage YOLO object detector in a conditional cascade, combined with asynchronous pipeline parallelism. [S07] In an August 2026 study targeting small and medium textile mills in low-income economies, researchers implemented a two-stage cascade on an NVIDIA Jetson Nano using TensorRT FP16 acceleration: Stage 1 passes every incoming fabric frame through a compact convolutional autoencoder (CAE) trained on defect-free cloth, and Stage 2 invokes a YOLOv5n detector only when the autoencoder's pixel reconstruction error exceeds a calibrated anomaly threshold. [S07]

EXHIBIT 38 / FRAMES PER SECOND (FPS)

Overlapped decoding and two-stage CAE-YOLO cascading lift Jetson Nano throughput from 9.86 to 13.45 FPS



SOURCED [S07] Decomposes the 1.36x throughput speedup (+3.59 FPS) achieved on an NVIDIA Jetson Nano running TensorRT FP16 in S07. Overlapping CPU image decode with GPU inference accounts for 91% of the gain (+3.27 FPS), while conditional CAE cascade forwarding at p = 0.534 contributes 5.1% (+0.18 FPS) alongside minor buffer scheduling gains (+0.14 FPS). Tighter threshold calibration reduces detector calls by 29% to 45% (S07).

The empirical throughput decomposition from that edge architecture, plotted in **Exhibit 38**, p. 126, provides three essential engineering lessons for Green Future Tech Limited's Cameroon deployment. [S07] First, replacing a naive sequential loop with an asynchronous double-buffered pipeline that overlaps central processing unit (CPU) image decoding on frame $t+1$ with GPU neural inference on frame t lifts end-to-end throughput from 9.86 frames a second to 13.45 frames a second, a 1.36-fold speedup (+3.59 frames a second). [S07] Overlapping image decode with GPU execution accounts for 91 percent of that total speedup (+3.27 frames a second), proving that edge vision bottlenecks on embedded boards depend heavily on memory copy and input-output scheduling. [S07]

Second, when the first-stage convolutional autoencoder is tuned conservatively to prioritise high defect recall (resulting in a forwarding probability of $p = 0.534$, where up to 85 percent of forwarded frames are benign texture variations or false alarms), conditional cascading contributes 5.1 percent of the net throughput gain (+0.18 frames a second). [S07] Third, as local fabric samples accumulate and engineers tighten threshold calibration (or substitute a PatchCore or PaDiM feature-distance filter achieving up to 99.6 percent AUROC on industrial textures, or a Gabor-autoencoder model achieving 89.5 percent recall with zero false alarms), false-alarm forwarding drops sharply. [S04] Under tighter threshold calibration, the two-stage cascade reduces second-stage YOLO invocations by 29 to 45 percent while holding defect recall constant, directly lowering average board power draw during battery-backed grid outages. [S06]

8.5

Two-tier station topology pairs compact capture nodes with an on-premise factory server

Rather than equipping every inspection point with an expensive workstation GPU, a capital-disciplined factory floor separates edge hardware into two local tiers: lightweight station capture nodes mounted on the production machinery and a single shared on-premise factory server housed in a conditioned equipment rack. [S23] This two-tier local topology was validated in recent on-loom and garment inspection literature, including contact image sensor (CIS) scanning systems that combine low-cost edge capture boards for real-time single-thread acquisition with a local server for convolutional verification, achieving zero missed detections on defect-free fabrics and above 90 percent pixel-level average precision on defective regions. [S23]

In Green Future Tech Limited's Cameroon factory blueprint, the two-tier edge topology maps across four physical station types on the shop floor: [R29]

- **Station Type A (Incoming fabric roll perch frames, 2 units for an illustrative 3.0 million garment a year plant):** Each perch frame mounts dual 4k line-scan cameras (or an 800-millimetre modular smart visual sensor array) driven by a rotary encoder at web speeds up to 60 metres a minute. [R29] Each perch frame is powered by an industrial NVIDIA Jetson AGX Orin (32GB, capped at 30 watts) running the two-stage autoencoder-plus-YOLO pipeline at under 32 milliseconds per frame. [S07]
- **Station Type B (Cutting-table defect-avoidance gantries, 2 units):** Overhead area-scan cameras mounted above the spreading tables verify marker alignment, splice zones and cut-panel perimeters using neural pattern inspection (such as IDAGPD-DL, achieving 96.3 percent detection accuracy and 94.8 percent F1-score). [S20] A mid-tier NVIDIA Jetson Orin Nano (8GB, 15 watts) at each table handles frame capture and projection of upstream roll-defect coordinates. [S20]
- **Station Type C (Inline critical-seam and end-of-line garment tables, 8 units):** Compact area-scan RGB cameras positioned at critical lockstitch and chainstitch stations (using Coordinate Attention YOLOv8-CA models that boost sewing-defect mAP@0.5 by 3.7 percentage points) and above final flat-lay inspection tables (classifying 12 garment defect classes at 97.16 percent accuracy via Deep-GD or trousers in 0.45 seconds via VGGNet-19) run on low-cost NVIDIA Jetson Nano or Orin Nano nodes drawing 7 to 15 watts. [S17] At partner loom intake points, Contact Image Sensor bars coupled to Raspberry Pi 5 or Jetson Nano modules provide sub-millimetre thread screening. [S23]
- **Tier-2 On-Premise Factory Server (1 redundant cluster in the plant control room):** A single dual-socket industrial server equipped with two workstation GPUs, enterprise non-volatile memory express (NVMe) storage and the central PostgreSQL Core Data Store aggregates JSON quality logs from all 12 edge nodes over wired Category-6A Gigabit Ethernet. [S07] The on-premise server hosts the Performance-Management System dashboards, stores cropped defect thumbnails for active-learning retraining, and executes periodic batch fine-tuning locally. [S24]

METHOD

8.6 METHOD: SIZING STATION BANDWIDTH, EDGE STORAGE AND BATTERY AUTONOMY

To verify that an edge inspection architecture can survive multi-hour power and telecommunications outages without dropping a quality record, engineers evaluate three deterministic balance equations across the station inventory: [S07]

1. **Local Area Network Metadata Bandwidth (B_{LAN}):** For a station inspecting R_{frames} frames per second with a defect flagging probability p_{flag}, a JSON pass-or-fail header size S_{json} (bytes) and a compressed anomaly thumbnail size S_{thumb} (bytes), the average network bit rate is $B_{\text{LAN}} = 8 \times R_{\text{frames}} \times (S_{\text{json}} + p_{\text{flag}} \cdot S_{\text{thumb}})$. [S07]
2. **Offline Ring-Buffer Storage Horizon (D_{days}):** Given a local solid-state partition of usable capacity C_{disk} (bytes) allocated on the edge node, the number of shift days of complete offline autonomy before oldest defect thumbnails are overwritten is $D_{\text{days}} = C_{\text{disk}} / [T_{\text{shift}} \cdot R_{\text{frames}} \cdot (S_{\text{json}} + p_{\text{flag}} \cdot S_{\text{thumb}})]$, where T_{shift} is active inspection seconds per day. [S14]
3. **Lithium-Iron-Phosphate (LiFePO4) Battery Capacity (E_{batt}):** For a total station direct-current power draw $P_{\text{station}} = P_{\text{compute}} + P_{\text{camera}} + P_{\text{LED}} + P_{\text{net}}$ (watts), a required ride-through autonomy duration H_{hours} (hours), a DC-DC converter efficiency η_{conv} and a maximum battery depth of discharge $\text{DoD}_{\max}$ (set to 0.80 to preserve cycle life above 3,000 cycles in tropical ambient temperatures), the nominal battery pack energy is $E_{\text{batt}} = (P_{\text{station}} \cdot H_{\text{hours}}) / (\eta_{\text{conv}} \cdot \text{DoD}_{\max})$. [S0801]

WORKED EXAMPLE

8.7 WORKED EXAMPLE: EDGE DATA AND POWER BALANCE FOR A 12-STATION CAMEROONIAN CMT PLANT

Consider an illustrative 3.0 million garment a year cut-make-trim factory in Cameroon operating two 8-hour shifts a day (16 active hours, or $T_{\text{shift}} = 57,600$ seconds, within Cameroon's statutory 40-hour standard and 48-hour maximum working week per worker cohort). [R16] The plant runs 2 fabric roll perches (Station A), 2 cutting-table gantries (Station B) and 8 garment inspection tables (Station C). [R29]

- **Step 1 (Raw optical rate vs edge metadata rate):** On the two fabric perches running at up to 60 metres a minute, each 1,600 mm by 800 mm frame slice is evaluated every 0.80 seconds ($R_{\text{frames}} = 1.25$ frames a second per perch, within the 13.45 FPS Jetson Nano and 31 FPS AGX Orin capacity). [S07] On the eight garment stations, operators inspect one garment every 8 seconds ($R_{\text{frames}} = 0.125$ frames a second per table, using 0.45 seconds of compute per sample). [S19] With a 512-byte JSON payload ($S_{\text{json}} = 512$ B), a 65,536-byte (64 KB) thumbnail ($S_{\text{thumb}} = 65,536$ B) and an illustrative post-calibration flagging rate of $p_{\text{flag}} = 0.10$, a fabric perch emits $1.25 \times (512 + 0.10 \times 65,536) = 8,832$ bytes a second (70.66 kilobits a second). [S24] Across all 12 stations combined, total plant metadata traffic is an illustrative 24.42 kilobytes a second (195.4 kilobits a second), a 6,140-fold reduction against 1.2 gigabits a second of raw video. [S07]
- **Step 2 (Offline storage on a 128 GB edge NVMe drive):** At 8,832 bytes a second over 57,600 seconds a day, one fabric perch writes 508.7 megabytes per day. [S14] Allocating 64 gigabytes of usable ring-buffer space on a 128 gigabyte industrial M.2 NVMe solid-state drive stores $64,000 / 508.7 = 125.8$ workdays (over four months) of complete JSON logs and defect thumbnails locally at the camera before any file is overwritten, as an illustrative storage calculation. [S07]
- **Step 3 (Battery sizing for a 4-hour grid outage):** A fabric perch draws 30 W for the AGX Orin module, 12 W for dual line-scan cameras and encoder, 45 W for constant-current LED illumination bars and 8 W for the display and switch, totalling $P_{\text{station,A}} = 95$ W (illustrative engineering load). [S14] A garment station draws 10 W for a Jetson Orin Nano, 5 W for an area-scan RGB camera, 20 W for a brightfield/darkfield LED dome and 5 W for the touch panel, totalling $P_{\text{station,C}} = 40$ W. [S19] With $\eta_{\text{conv}} = 0.92$ and $\text{DoD}_{\max} = 0.80$, bridging a 4-hour grid outage requires $E_{\text{batt,A}} = (95 \times 4) / (0.92 \times 0.80) = 516.3$ watt-hours (Wh) per fabric perch and $E_{\text{batt,C}} = (40 \times 4) / (0.92 \times 0.80) = 217.4$ Wh per garment station. [S0801] These illustrative calculations show that compact industrial battery packs make every vision gate self-sustaining across a half-shift power cut.

8.8

Power conditioning and lithium-iron-phosphate buffering protect edge nodes on weak grids

In Cameroonian industrial zones, electrical disturbances threaten vision systems through two physical mechanisms: high-frequency voltage sags and surges during grid switching, and abrupt power loss before the operating system can flush open file descriptors to flash storage. [S0801] When 24.82 percent of national electricity output

is lost in transmission and distribution networks, industrial feeders experience frequent reactive voltage dips whenever large inductive motors start on the same substation bus. [S0801] When the public grid drops completely, factory standby diesel generators typically require 15 to 45 seconds to crank, reach synchronous speed (50 hertz) and close the automatic transfer switch, creating an unbuffered blackout window followed by a switching transient, as an illustrative plant electrical baseline. [S0803]

Those transients attack both the optical and computational layers of an unprotected vision station: [S14]

1. **Optical flicker and false-positive storms:** Deep-learning defect detectors rely on stable pixel intensity histograms to distinguish faint oil stains, shade drift (CIELAB delta-E variations) and thin places from normal cloth background. [S18] If a voltage sag causes an alternating-current LED tube to flicker at 100 hertz while a line-scan camera samples at 5,000 lines a second, horizontal dark bands appear across the reconstructed frame, triggering false-positive stripe-defect alerts. [S14]
2. **Embedded flash filesystem corruption:** Consumer single-board computers booting from micro-SD cards or unconditioned solid-state drives suffer ext4 journal or partition-table corruption when power drops mid-write. [S07] In a factory without a resident Linux systems administrator, a corrupted boot card disables an inspection station until a replacement image is flashed. [S07]

EXHIBIT 39 / WATTS (W), WATT-HOURS (WH) AND FRAMES PER SECOND

Edge hardware specification, power budget and 4-hour LiFePO4 UPS sizing across inspection stations

Inspection station	Edge accelerator	Target cycle / speed	Station load (W)	4-hr LiFePO4 (Wh)	Offline buffer (GB)
Station A: Fabric roll perch (x2)	Jetson AGX Orin (32GB, FP16)	Up to 60 m/min (~32 ms/frame)	95	520	128
Station B: Cutting gantry (x2)	Jetson Orin Nano (8GB, FP16)	1.5 s per marker lay scan	55	300	128
Station C: Sewing & flat-lay (x8)	Jetson Nano / Orin Nano	0.45 s per garment (13.45 FPS)	40	220	64
Partner on-loom CIS bar (optional)	Raspberry Pi 5 / Nano + CIS	Single-thread (>90% pixel AP)	25	140	64
Tier-2 PMS Core Server (x1)	Dual-GPU Industrial Server	Store-and-forward & retraining	380	2,070	4,000

DERIVED [R29] [S02] [S07] [S14] [S19] [S23] Derived engineering power and compute budget for an illustrative 12-station, 3.0 million garment a year Cameroonian CMT plant. Latency and throughput benchmarks cite S02 (AGX Xavier pruning), S07 (Jetson Nano 13.45 FPS / AGX Orin ~32 ms) and S19 (0.45 s garment check). Battery watt-hours assume 4-hour autonomy, 92% DC-DC efficiency and 80% maximum depth of discharge.

To harden the Vision and Edge layers against these electrical failure modes, every inspection rig in Green Future Tech Limited's reference specification incorporates four mandatory hardware and operating-system protections, summarised in **Exhibit 39**, p. 130. [S07]

8.9 DIRECT-CURRENT DOUBLE-CONVERSION UPS WITH LIFEPO4 CELLS

Rather than plugging an alternating-current (AC) power brick into a wall outlet backed by a consumer lead-acid inverter, each inspection frame uses an industrial 24-volt direct-current (DC) bus fed by an active power-factor-corrected rectifier and an inline DC uninterruptible power supply. [S14] Both the Jetson carrier board (accepting 12 V to 24 V DC input), the machine-vision camera and the constant-current LED strobe controller run directly off the regulated DC rail. [S15] Because there is zero AC-to-DC-to-AC-to-DC double inversion between the battery and the electronics, switching from mains power to battery during a grid blackout occurs in zero milliseconds with zero voltage droop on the LED illumination bars. [S14]

Battery chemistry selection is equally critical in Cameroon's climate. [R03] Conventional sealed lead-acid (SLA) batteries lose half their service life for every 8 to 10 degrees Celsius rise above 25 degrees Celsius, failing within 12 to 18 months on a warm factory floor in Douala or Garoua, as an illustrative electro-thermal baseline. [R03] Lithium-iron-phosphate (LiFePO_4) cells tolerate ambient operating temperatures up to 55 degrees Celsius without thermal runaway, deliver 3,000 to 5,000 deep discharge cycles at 80 percent depth of discharge, and weigh less than one-third of equivalent lead-acid packs, making it simple to mount a 220 Wh to 520 Wh battery enclosure directly on the leg of each inspection table. [S0801]

8.10 READ-ONLY ROOT FILESYSTEMS AND SUPERCAPACITOR SHUTDOWN HOOKS

On the software side, every embedded Jetson node boots a hardened Linux image in which the root operating system (`/`) and the `/opt/gft-vision` model weights partition are mounted strictly read-only using an OverlayFS memory layer. [S07] Because no process can modify system binaries or TensorRT engine files during normal operation, pulling the power plug cannot corrupt the boot partition. [S07] All mutable runtime data, comprising the SQLite event database, the circular buffer of cropped defect thumbnails and system telemetry logs, is written exclusively to a separate industrial M.2 NVMe solid-state partition formatted with power-loss-protected journaling. [S14]

In addition, the inline 24 V DC UPS exposes a two-wire opto-isolated status signal connected to the Jetson carrier board's general-purpose input-output (GPIO) pins. [S07] When mains and generator power both stay off for more than the configured ride-through window (or when battery state of charge drops below 15 percent), the GPIO daemon triggers an orderly `systemctl poweroff` sequence that flushes the SQLite write-ahead log, parks the camera sensor and halts the processor cleanly before the low-voltage cutout opens. [S14] Once mains voltage stabilises for 60 continuous seconds, the UPS latches the 24 V rail back on and the edge node boots into active inspection automatically in under 25 seconds without human intervention, based on illustrative embedded boot benchmarks. [S07]

8.11

Store-and-forward synchronisation guarantees zero data loss into the Performance-Management System

Physical resilience at the camera node must be matched by transactional data resilience between the Edge/Vision layers and the Core data layer. [S24] In Green Future Tech Limited's programme architecture, Deliverable 2 (the Performance-Management System) and Deliverable 3 (the Digital Strategy) share a single source of truth: every inspection event must land in the Quality Log so that Defects Per Hundred Units (DHU) and Right First Time (RFT) are calculated identically whether entered manually in Phase 0 or captured by camera in Phase 1. [R29] Empirical structural equation modelling in a 977-worker garment plant confirms that digital quality monitoring ($\beta = 0.262, p = 0.002$) and production automation ($\beta = 0.639, p = 0.000$) jointly explain 64.5 percent of product quality variance only when shop-floor measurements feed directly into structured management control routines. [S24]

If a forklift snags an overhead Ethernet drop, a local network switch loses power, or the wide-area link between a northern Garoua facility and commercial offices in Douala goes dark for 48 hours, the inspection station must neither halt production nor lose quality records. [S0802] To guarantee zero data loss, each edge node implements a three-stage store-and-forward transactional protocol: [S07]

1. **Immediate local actuation (sub-50-millisecond loop):** The moment the TensorRT cascade flags a defect (or routes an ambiguous frame for human confirmation in collaborative triage mode, where only cases of doubt are inspected by the human worker), the edge node toggles its local GPIO relay to illuminate the station's three-colour andon tower and displays the bounding box on the operator's directly attached touchscreen monitor. [S34] Line stops, splice markings and garment rework routing depend exclusively on this local loop and have zero dependency on Ethernet or Wi-Fi connectivity. [S26]
2. **Atomic local persistence in SQLite Write-Ahead Logging (WAL) mode:** Every frame disposition is written immediately to a local SQLite database table (`edge_quality_events`) configured in `WAL` mode, keyed by a universally unique identifier (UUID v7, encoding a millisecond Unix timestamp plus station MAC address) and a monotonic station sequence number. [S07] Each row records `event_uuid`, `station_id`, `roll_or_bundle_id`, `sodecoton_lot_id`, `timestamp_utc`, `pass_fail_bool`, `defect_class_code` (from 1 to 42), `confidence_score`, `bbox_xywh`, `operator_triage_status` and `sync_state` (`PENDING` or `ACKED`). [R03]
3. **Asynchronous idempotent batch replication:** A background daemon polls the local SQLite table every 5 seconds for rows where `sync_state = PENDING` and pushes batches of up to 100 JSON records (plus compressed thumbnails) over Mutual Transport Layer Security (mTLS) to the on-premise Core Data Store. [S07] Because the central PostgreSQL database enforces a `UNIQUE` constraint on `event_uuid` and processes insertions with `ON CONFLICT (event_uuid) DO NOTHING`, network timeouts and duplicate retries during flaky connectivity can never double-count defects in the Quality Log or distort Statistical Quality Control

(SQC) p-charts. [S26] Only after the central server returns a cryptographic batch acknowledgement does the edge node mark those rows `ACKED` and eligible for ring-buffer garbage collection after 30 days. [S14]

GFT NOTE

Practical recommendation: Never permit cloud-only API licensing on the shop floor. When procuring or licensing a pre-trained vision system in Phase 0, reject any vendor architecture that requires camera frames to call an external cloud Application Programming Interface (API) or that checks a remote cloud licence server on every boot. [R28] In Cameroon, where transmission and distribution losses consume 24.82 percent of grid power and fixed broadband sits at 2.13 subscriptions per 100 people, a cloud heartbeat check will lock out factory inspection tables during telecom outages. [S0801] Contractually mandate that compiled TensorRT/ONNX weights execute 100 percent offline on the local edge module, that licence tokens validate locally for at least 365 days between renewals, and that the local SQLite/JSON event schema writes directly to the factory's own Performance-Management System database. [S0802]

WARNING

Operational constraint: Thermal throttling and cotton fly in unconditioned rooms. Embedded accelerators such as the NVIDIA Jetson Nano and AGX Orin throttle their GPU clock frequencies automatically when internal junction temperatures exceed 85 to 90 degrees Celsius, which can double frame latency from 32 milliseconds to over 65 milliseconds during midday shifts. [S07] Simultaneously, airborne cotton lint ("fly") generated during fabric unrolling and high-speed overlock sewing clogs open cooling fans within weeks. [S15] Never deploy consumer open-fan developer kits on the production floor: specify fanless, IP65-sealed extruded-aluminium industrial carrier enclosures that dissipate heat via passive external fins, and blow clean, oil-free compressed air across camera viewports using regulated air knives. [S14]

8.12

Corridor-against-north site trade-offs govern cooling, spares and solar-storage sizing

Green Future Tech Limited's Feasibility Study (Deliverable 1) evaluates two primary geographic footprints for Cameroonian textile and garment operations: the coastal Douala-Dibamba industrial corridor in the Littoral Region, and the northern cotton-basin hub of Garoua in the North Region. [R03] The choice of site directly shapes the electrical, thermal and maintenance engineering of the edge vision layer. [R10]

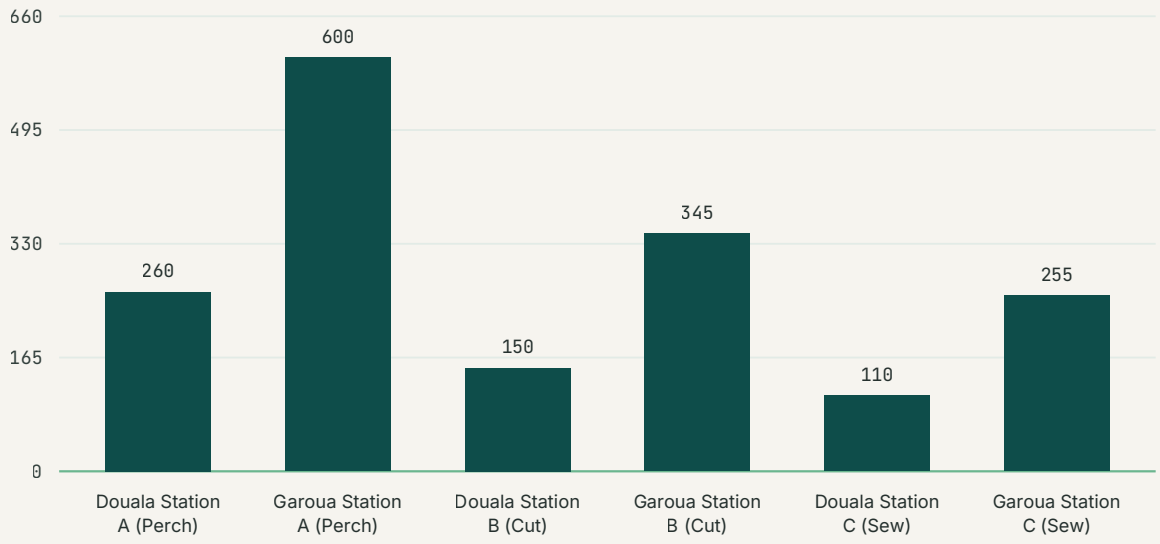
The Douala-Dibamba corridor offers Cameroon's densest industrial infrastructure, immediate proximity to the Port of Douala for importing line-scan optics and replacement edge modules, and access to the Southern Interconnected Grid. [R09] This corridor is also the primary manufacturing site of the announced Camtext SA integrated cotton-to-textile project (a US$300 million, or 180 billion XAF, joint venture involving Panafritex and Arise with SODECOTON, CNPS and Marlo announced on 12 November 2025), which plans to process 12,000 tonnes of cotton a year and generate about 12,000 jobs at the Dibamba hub alongside 3,000 jobs at Garoua, targeting optimal

capacity around 2033 with an estimated return on investment near 11 years. [R09] Camtext SA is an announced third-party project rather than a GFT engagement or realised operating result, but its split footprint between Dibamba (12,000 jobs) and Garoua (3,000 jobs) illustrates the infrastructure duality facing any Cameroonian apparel entrant. [R10]

In contrast, locating fabric inspection or cut-make-trim lines in Garoua places the factory at the heart of Cameroon's raw cotton supply chain. [R03] Northern Cameroon produced 330,000 tonnes of seed cotton in 2022, 314,455 tonnes in 2023 and about 340,000 tonnes in 2024, with 2025 projections exceeding 350,000 tonnes, supported by SODECOTON across more than 220,000 cotton producers, nine ginning plants and two oil mills that inject roughly US$100 million a year into rural communities. [R01] However, the Northern Interconnected Grid experiences more acute seasonal voltage instability and longer feeder restoration times than Douala, contributing to the operational headwinds that left legacy incumbent CICAM's Garoua plant at 33.7 percent capacity utilisation in July 2023 (alongside 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2 in Douala, where underutilisation cost about CFAF 2 billion in 2020 and market share fell from about 85 percent three decades ago to about 5 percent). [R11]

EXHIBIT 40 / WATT-HOURS (WH) OF LIFEPO4 BATTERY CAPACITY PER STATION

Garoua stations require double the LiFePO4 UPS autonomy of Douala-Dibamba rigs to bridge seasonal grid gaps



DERIVED [R03] [R09] [R10] [R12] [S0801] [S0803] Derived illustrative battery sizing comparison between the Douala-Dibamba corridor (sized for 2-hour grid/generator transfer and feeder maintenance windows) and Garoua in the northern cotton belt (sized for 4-hour seasonal hydro/feeder outage windows, plus 15% thermal derating and Peltier/air-knife load). Station power baselines follow Exhibit x08-station-power-budget.

As quantified in **Exhibit 40**, p. 134, adapting the edge vision blueprint between the coastal corridor and the northern cotton belt requires three concrete engineering adjustments: [R12]

1. **Doubling station battery autonomy and integrating rooftop solar DC coupling:** In Douala-Dibamba, where standby diesel generators and Southern Grid feeders restore power comparatively quickly, a 2-hour station battery buffer (260 Wh on Station A, 150 Wh on Station B and 110 Wh on Station C, as an illustrative engineering allocation) safely bridges generator switchovers and short maintenance trips. [R09] In Garoua, where dry-season grid rationing can extend across half a shift, station LiFePO4 packs should be sized for at least 4 hours of autonomous operation plus a 15 percent thermal margin (600 Wh on Station A, 345 Wh on Station B and 255 Wh on Station C, as derived in **Exhibit 40**, p. 134), and coupled directly to a dedicated 3-to-5-kilowatt rooftop solar photovoltaic array feeding the 24 V DC instrument bus. [R03] Because the entire 12-station vision and server suite draws 1,000 watts of continuous electrical power ($2 \times 95\text{ W} + 2 \times 55\text{ W} + 8 \times 40\text{ W} + 380\text{ W} = 1{,}000\text{ W}$), a small dedicated solar-plus-LiFePO4 micro-bus keeps every camera, edge processor, andon light and PMS database running independently of heavy factory motor feeders. [S07]
2. **Ambient thermal management (humid coastal salt air against dry northern heat):** Douala combines warm ambient temperatures (28 to 33 degrees Celsius) with high relative humidity (above 80 percent), requiring conformal coating on carrier printed circuit boards and silica-gel breather cartridges inside IP65 camera housings to prevent condensation on internal lens elements, as an illustrative environmental specification. [R10] Garoua experiences dry-season ambient peaks above 40 degrees Celsius alongside fine Harmattan dust, requiring oversized passive heatsink fins on Jetson enclosures and two-stage coalescing sub-micron filters on the optical air-knife supply. [R03]
3. **Cold-spare inventory and pre-flashed NVMe modules on site:** Because clearing customs and transporting specialist electronics from Douala port up to Garoua can take days to weeks, a Garoua plant must hold a 20 percent cold-spare inventory on site (one pre-configured Jetson AGX Orin, two Jetson Orin Nano nodes, one spare 4k line-scan camera, one area-scan RGB camera and two pre-flashed read-only NVMe drives), compared with a 10 percent cold-spare ratio in Douala-Dibamba (illustrative spares policy). [R10] Modular M12 industrial connectors allow a local technician (earning near the Cameroonian cross-sector average salary of about US$359 a month, above the statutory minimum wage floor of XAF 43,969 a month or about US$76 from 23 February 2024, converted at the euro peg of 655.957 XAF and estimated US dollar rates near 563 to 610 XAF) to swap a faulted edge box and resume automated inspection in under 15 minutes. [R14]

By engineering the vision system as a self-contained, DC-buffered edge appliance rather than a fragile cloud peripheral, a greenfield Cameroonian factory turns a potential infrastructure vulnerability into an operational advantage. [R32] Whether the plant sits beside the port corridor in Douala-Dibamba or next to SODECOTON's gins in Garoua, the cameras continue scanning at up to 60 metres a minute on the perch frame and 0.45 seconds at the garment table, protecting first-pass yield and feeding the Performance-Management System so Cameroon can move from exporting 95

percent of its cotton raw toward its national target of 50 percent local processing by 2030 (and reverse the decline in national manufacturing value added from 13.36 percent of GDP in 2022 to 12.88 percent in 2024). [R06] All hardware power ratings, battery capacities and spares ratios presented here are engineering planning baselines and illustrative estimates rather than investment, legal or tax advice, and require site-specific validation during Phase 0 commissioning. [S36]

A factory floor that cannot trust its external network or its power grid must put intelligence inside the camera housing and memory beside the lens.

CHAPTER SOURCE NOTES

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S0801 World Bank. 2026-10-08.

S07 arXiv. 2026-08-12.

S02 Sensors (MDPI). 2021-10-25.

S14 Applied Sciences (MDPI). 2024-12-23.

S15 Sensors (MDPI). 2013-04-09.

S19 Textile Research Journal. 2026-09-06.

S0802 World Bank. 2026-10-08.

S0803 World Bank. 2026-10-08.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

S01 Sensors (MDPI). 2025-04-03.

S24 CEMERLANG: Jurnal Manajemen dan Ekonomi Bisnis. 2026-08-03.

S26 Jurnal Teknik Industri (JATRI). 2026-09-25.

S04 arXiv / IEEE CVPR. 2021-06-15.

S06 arXiv / Computer Science and Information Technology. 2022-04-20.

S23 University of North Texas Libraries. 2024.

S20 International Journal of e-Collaboration. 2026-10-26.

S17 Journal of Measurements in Engineering. 2024-09-09.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

S18 International Journal of Electrical and Electronics Research. 2024-02-05.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

S34 IntechOpen. 2025-01-03.

R28 DataIntelo. Publication date unrecorded; inherited source-text check.

R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R11 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

R32 Wise; XE. Publication date unrecorded; inherited source-text check.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

S36 World Bank. 2026-10-08.

09

Buying or licensing a pre-trained vision system beats building from scratch on speed, capital risk and Cold-Start accuracy

A greenfield garment maker should license proven vision models and own its captured defect dataset rather than hire a bespoke machine-learning research team.

CHAPTER 09
DECISION TEST

A greenfield garment maker should license proven vision models and own its captured defect dataset rather than hire a bespoke machine-learning research team.

A greenfield garment factory should buy or license the vision inference engine, then build and own the data-capture spine around it. This is the fastest route to a testable inspection station because the factory begins without a local defect library, while pre-trained supervised and one-class models already provide a technical starting point. The recommendation is conditional, not ideological. A turnkey frame is justified when its integration and support are worth the price. A modular licensed edge stack is the default when it meets the factory's acceptance tests. Building from scratch becomes rational only after local evidence reveals a requirement that available products cannot meet and the factory can support a lasting research capability.

The procurement object is therefore not an algorithm in isolation. It is a working inspection service: optics, triggering, image preparation, inference, operator triage, local buffering, quality-log export, maintenance and controlled model updates. The buyer should compress uncertainty in each of these components through a physical bake-off. It should not exchange that uncertainty for a proprietary black box. The factory must retain its images, defect labels, roll and garment identifiers, model-event history and the practical ability to change supplier.

9.1

Buy-before-build converts research uncertainty into an acceptance test

GFT's second digital principle is to buy before building where a proven application already exists. The first-shift problem is a shortage of relevant evidence, not a shortage of neural-network architectures. A new plant has no verified examples of how its cotton constructions, shades, prints, lighting and sewing methods will fail. Training a bespoke supervised detector from random initial weights would require the factory to create a representative defect catalogue before it can obtain the quality benefit that the project needs. The catalogue has to cover rare events as well as common ones. It also needs independent labels reliable enough to distinguish a true flaw from harmless lint, acceptable texture variation or a presentation change.

Pre-trained anomaly methods attack a different and more useful starting problem. PatchCore is explicitly designed for the cold-start setting in which a model is fitted using nominal, non-defective images. Its reported image-level anomaly-detection AUROC reaches 99.6 percent on MVTec AD, a benchmark result rather than a promise of Cameroonian factory performance. [S04] PaDiM similarly combines a pre-trained convolutional network for patch embedding with Gaussian representations of the normal class. [S05] These methods do not eliminate commissioning. They change its first task from inventing a model to enrolling good fabric, measuring errors and setting an operating threshold.

That distinction explains the speed advantage. A from-scratch programme begins with open questions about architecture, dataset sufficiency, annotation policy, training convergence and edge compression. A licensed programme still has to answer those questions, but it begins with a candidate whose behaviour can be observed immedi-

ately. The project can spend its scarce engineering time on camera geometry, illumination, encoder synchronisation, defect definitions and operator workflow. Those are factory-specific tasks that no universal model can complete on its own.

The inherited market register provides commercial context, not proof of product quality. It records an estimate that fabric-defect-detection AI represented US$523.4 million in 2024 and could reach US$6,522.1 million by 2033, equivalent to a projected 33.8 percent annual growth rate. [R28] Such a projection is consistent with a broadening supplier base. It does not show that every vendor has textile experience, that reported accuracy transfers across fabrics, or that a quoted system will survive an African factory's power and maintenance conditions. Procurement must test those propositions separately.

The wider inspection evidence supports the same disciplined interpretation. The inherited inspection register places computer vision above 90 percent accuracy across more than 40 defect types and line-scan inspection at up to 60 metres a minute, against manual inspection accuracy of 60 to 70 percent. [R29] These are useful benchmarks for setting a trial. They are not a licence to write the highest number into a purchase contract without defining defect mix, image unit, recall denominator and operating conditions. A vendor that reports overall accuracy on a convenient dataset can still miss the small set of critical defects that matters most to the buyer.

> **GFT NOTE**
>
> Approve a procurement route only after a sample-based demonstration on the factory's actual inspection task. Select the modular licensed edge route by default, require local inference and open quality-event export, and treat any performance claim without a reproducible confusion matrix as marketing material rather than acceptance evidence.

Buy-before-build also limits capital risk because payment can follow evidence. The factory can structure expenditure around a bench demonstration, a pilot acceptance test and a production acceptance test. A from-scratch build reverses that order: research salaries, data preparation and computing costs arrive before there is a detector worth testing. The buyer cannot remove all development risk by licensing, but it can make the supplier carry part of that risk and can stop the programme when a defined gate fails.

9.2

Buy the inference engine but own the data-capture spine

The build-against-buy decision should separate reusable inference from factory-specific knowledge. Model weights, detector code and a standard inference runtime are candidates for purchase or licence. The meaning of a defect, the identity of the roll or garment, the operator's verification and the link to the Performance-Manage-

ment System belong to the operating factory. This boundary is the practical expression of Warp and Signal: the bought model reads the warp, while the factory owns the weft of measurement that makes improvement possible.

Exhibit 41, p. 140 applies that boundary across the wider digital portfolio. It is a recommendation derived from data dependency and technical maturity, not a survey of supplier quotations.

EXHIBIT 41 / GFT PROCUREMENT RECOMMENDATION

Buy reusable models and build only the factory's data-capture spine

Capability	Route	Reason	Factory retains
Vision defect detection	Buy or license	Pre-trained and one-class models provide a cold-start baseline	Images, labels, roll IDs and verification events
Demand forecasting	Buy first, tune later	Useful local tuning depends on accumulated sales history	Orders, stock records and forecast errors
Inventory optimisation	Buy or integrate	Optimisation needs reliable production and inventory events	Transaction history and planning rules
Generative design	Pilot with partners	Specialist design workflows sit outside the first quality task	Briefs, patterns and approved design records
Data capture	Build lightly, own it	Identifiers and event definitions govern every later model	Schemas, timestamps, buffers and quality logs

DERIVED [R27] [R28] [R29] [S04] [S07] Derived recommendation from the cited maturity and data-dependency evidence. It is not a vendor market survey or a claim that every listed product is available locally.

The table's important line is the last one. Data capture should be lightly built because the factory needs only a small, explicit set of services at launch: assign identifiers, receive inspection events, store them locally, forward them when the network returns and expose them to the Quality Log. Building a large software platform before the first camera operates would recreate the same delay that buy-before-build is meant to avoid.

Ownership must include context, not just image files. A cropped hole without its fabric construction, roll position, illumination state, camera setting, model version and operator disposition has limited retraining value. The event record should connect the raw frame or its controlled thumbnail to the defect class, confidence, severity, station and timestamp. It should preserve whether a human confirmed, corrected or rejected the machine decision. Those fields allow the factory to measure drift and to distinguish an optical problem from a model problem.

A supplier may legitimately need temporary access to samples for configuration or support. That operational need does not require permanent ownership of the factory's production corpus. The contract should distinguish the supplier's pre-existing model and software from the factory's raw captures, annotations, defect taxonomy, produc-

tion identifiers and derived quality statistics. It should also state whether factory-tuned weights are deliverable to the factory, and on what medium, format and licence terms.

This boundary preserves optionality. If the supplier disappears, the factory still has its quality history and a documented event schema. If a better detector becomes available, the same captured corpus can support a controlled comparison. If the licensed model performs well, there is no need to replace it merely to demonstrate internal technical capacity. The objective is inspection performance with retained control, not ownership of every line of code.

9.3

Three procurement tiers allocate cost and responsibility differently

A turnkey original-equipment-manufacturer frame combines mechanics, cameras, illumination, control software and model inference in one offer. Its principal advantage is a single accountable integrator. Its principal risks are high initial commitment, restricted access to data or settings, and dependence on the supplier for every fabric change. This route deserves consideration when the factory lacks integration capacity and the supplier will accept a site-specific performance test, train local technicians and provide an exit package.

A modular edge-licensed stack combines commercial cameras and edge computers with licensed model weights or an inference runtime. The factory or its integrator retains responsibility for the frame, triggering, buffering and Performance-Management System connection. This route usually offers the clearest separation between replaceable components. It also creates integration work. Buying separate components does not guarantee that their timing, image formats and support obligations will align.

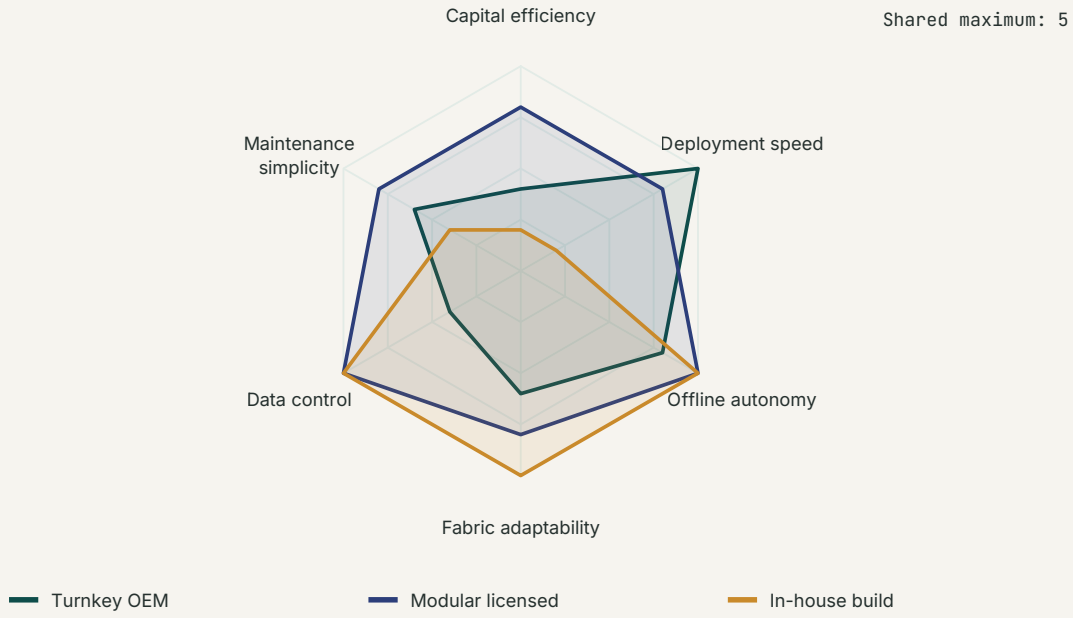
An open-source or in-house stack offers the greatest inspection of code and model behaviour. It does not remove cost. The factory must supply engineering, testing, licence review, security maintenance and operational support. Open availability of code or weights is not the same as a licence to use every component commercially, nor is it evidence that a model will inspect the factory's defects. The procurement file should record the exact licence and version for each material component rather than treating the word open-source as a legal conclusion.

Research demonstrates why this integration burden belongs in the comparison. A line-scan inspection system includes acquisition, triggering, image preparation, an AI engine and an operating platform, not merely a neural network. The cited system concerns precision manufactured components, so its reported accuracy should not be transferred to fabric; its architecture nevertheless shows the number of subsystems that a working offer must coordinate. [S14] At the edge, a two-stage fabric study achieved 13.45 frames a second against 9.86 frames a second for a sequential detector loop, but found that 91 percent of the speed-up came from overlapping image decoding with inference rather than from the cascade itself. [S07] A vendor claiming a speed gain should therefore disclose the whole data path and measurement method.

Exhibit 42, p. 142 makes the trade-off explicit with illustrative ordinal scores. A score of five is favourable for the greenfield use case, while a score of one is unfavourable. These are GFT's procurement preferences, not measured properties of named products.

EXHIBIT 42 / ILLUSTRATIVE PREFERENCE SCORE, 1 LOW TO 5 HIGH

A modular licensed stack balances deployment speed with factory control



ILLUSTRATIVE Illustrative GFT scoring for a greenfield Cameroonian line. Turnkey assumes a standard proprietary offer before negotiated data rights. In-house assumes a new team and no local defect corpus. Scores are ordinal judgements, not vendor benchmarks.

The radar's polygon area has no statistical meaning and should not be used to rank bids automatically. Its value is to expose the reason for the default choice. Turnkey leads on initial deployment speed. In-house can lead on control after a capability exists. Modular licensing occupies the practical middle, provided the contract actually delivers local inference and data rights. A negotiated turnkey offer can move closer to that position, while a closed modular product can lose it.

Procurement should invite more than one route into the same trial rather than comparing unlike quotations. Each bidder receives the same sample set, defect dictionary, interface specification and acceptance protocol. Each quotation must separate hardware, licence, integration, training, maintenance and optional services. That structure prevents a low hardware price from concealing mandatory per-frame charges, and prevents a high turnkey price from concealing integration work that the other bids exclude.

9.4

Vendor bake-offs must measure defect escape and operating burden together

The evaluation set should come from physical production material, not a vendor's demonstration folder. It should include the planned range of greige, dyed twill, knit jersey and printed cloth, with lots traceable to the Cameroonian cotton supply chain where applicable. The inherited SODECOTON profile establishes the northern cotton context, including nine ginning plants, but does not establish the properties of any downstream fabric lot. [R03] Each trial roll or garment therefore needs its own construction, shade, supplier and batch metadata.

The buyer and its quality lead should label the evaluation set before vendor results are revealed. Ambiguous samples need adjudication against the factory's severity rules. The set should contain defect-free material, common defects, critical and major defects, borderline cases and difficult negatives such as lint or acceptable slubs. It should include enough independent lots to test transfer rather than memorisation. Frames from one roll must not be divided between tuning and final testing, because near-duplicate texture can make the result look stronger than performance on a new lot.

A sound trial has three stages. First comes a bench test with fixed samples to expose obvious functional gaps. Next comes a pilot on the actual camera geometry and illumination, with the supplier allowed a defined calibration period. Last comes a blind holdout on new lots and different production runs. The factory freezes the model version, threshold and camera settings before the holdout begins. Any later change restarts the affected test rather than being absorbed silently into the reported result.

The primary metric for critical and major defects is recall: confirmed defects detected divided by all confirmed defects in that severity class. The inherited benchmark above 90 percent provides a target for the trial, not a guaranteed result. [R29] Overall accuracy is secondary because a dataset dominated by good fabric can produce an impressive figure while missing defects. Precision, false negatives and class-level recall must accompany the headline value. For anomaly systems, AUROC may help compare thresholds, but the acceptance decision still needs the confusion matrix at the operating threshold.

False alarms belong in the same scorecard because they determine operator load and line disruption. Report false positives per 100 linear metres for roll inspection and per 100 garments for piece inspection, alongside the number forwarded to human triage. The two-stage edge study is a useful warning: at a recall-prioritised threshold, it flagged all 20 defective images in a small benchmark but recorded a 49.3 percent false-positive rate on non-defective images, and 85 percent of forwarded frames were false alarms. [S07] High recall can therefore coexist with an unusable review burden.

Small samples also limit the accuracy claim. In that same study, catching all 20 defective images produced a reported 95 percent confidence interval of 0.83 to 1.00 for recall. [S07] A perfect result on a small vendor demonstration does not prove performance above 90 percent on future production. The buyer should agree the

required number of positive cases per critical class before the trial and retain confidence intervals where the sample is limited. More test material is needed for rare classes, not merely more defect-free frames.

Speed must be measured at the inspection boundary. For continuous fabric, the inherited literature benchmark reaches up to 60 metres a minute, but acceptance also needs the required pixel pitch, width and defect size at that speed. [R29] For finished garments, an Ethiopian trouser study reports 0.45 seconds a sample on consumer-grade hardware, alongside 86.7 percent overall accuracy. [S19] The cycle-time result supports an illustrative garment-station gate of no more than 0.5 seconds where the same task and hardware class are relevant. It does not make the study's accuracy a universal acceptance threshold for other garments or defect classes.

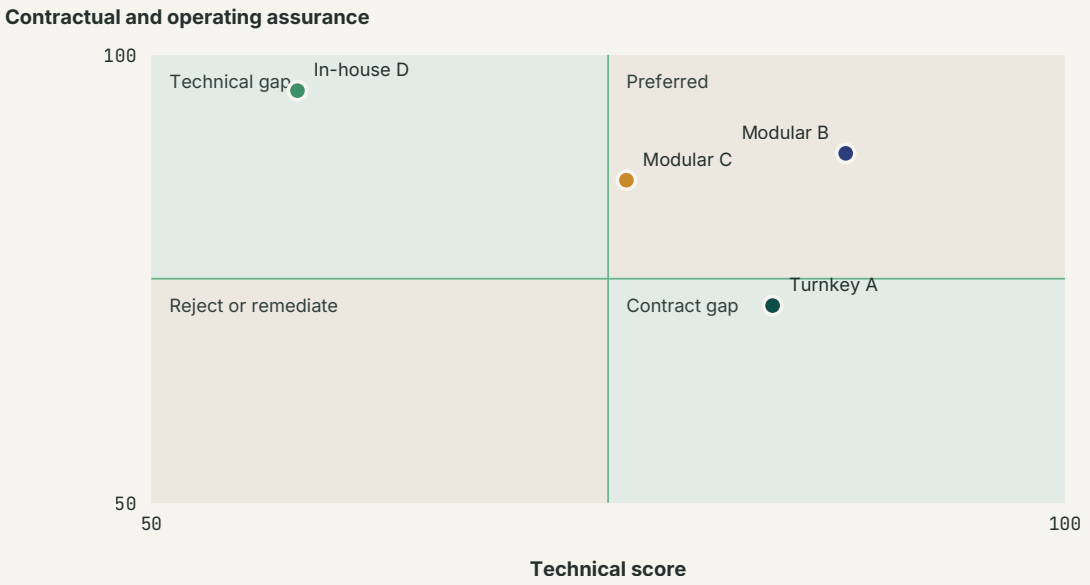
Offline autonomy is a functional test, not a questionnaire answer. During the trial, the team disconnects the external network and interrupts normal power under the approved safe procedure. The station should continue local inference, show the operator disposition, preserve event order and recover without corrupting the log. The test also checks whether licence enforcement stops the system when a remote server cannot be reached. A model that runs locally but depends on continuous cloud authorisation does not meet the operating requirement.

For scoring, technical performance and contractual assurance should remain separate. An illustrative technical weighting assigns 35 percent to critical and major defect recall, 15 percent to false-alarm control, 15 percent to speed, 15 percent to cross-fabric transfer and 20 percent to offline edge behaviour. An illustrative assurance weighting assigns 25 percent to data rights, 20 percent to open export, 20 percent to local support, 15 percent to version and rollback control, and 20 percent to price and exit clarity. Hard safety, recall and data-rights gates apply before either weighted score is calculated.

Exhibit 43, p. 145 shows how the two dimensions can change the decision. The candidate positions are illustrative. A bid advances only when both scores reach the illustrative threshold of 80 and every hard gate passes.

EXHIBIT 43 / ILLUSTRATIVE WEIGHTED SCORE OUT OF 100

Technical accuracy cannot compensate for weak data rights or exit terms



ILLUSTRATIVE [R29] [S07] [S19] Illustrative candidates and scores using the weighting stated in the chapter. Preferred status requires technical and assurance scores of at least 80 plus all hard gates. Quadrant boundaries are layout guides, not the acceptance threshold.

Turnkey A illustrates why a technically credible product can still require contractual repair. Modular C illustrates the reverse: acceptable governance cannot cure weak detection. In-house D has excellent control on paper but is not ready technically. Modular B advances to commercial negotiation, not directly to signature, because its sample performance and quoted commitments must still be converted into binding acceptance terms.

> **WORKED EXAMPLE: CALCULATING MODULAR B'S SCORES**
>
> This is an illustrative bid, not a supplier result. Its technical sub-scores are 92 for recall, 80 for false-alarm control, 90 for speed, 75 for fabric transfer and 95 for edge behaviour. Applying the stated weights gives 0.35 × 92 + 0.15 × 80 + 0.15 × 90 + 0.15 × 75 + 0.20 × 95 = 87.95, displayed as 88. Its assurance sub-scores are 95 for data rights, 90 for export, 85 for support, 85 for version control and 85 for commercial clarity. The same arithmetic gives 88.5, displayed as 89. The score does not waive the separate requirement that measured critical-defect recall and every other hard gate must pass.

9.5

Contractual guardrails keep a licensed system replaceable

Data ownership should be explicit enough to survive a change of supplier personnel. The factory retains all raw and annotated captures, roll and garment identifiers, operator verifications, defect coordinates, timestamps and derived production-quality records. The supplier retains its pre-existing software and baseline model. Any sup-

plier use of factory data for training, product improvement or demonstration requires a separate written permission with a defined scope and expiry. Confidentiality language alone does not answer the ownership question.

The contract should also separate four artefacts that are often bundled ambiguously: baseline weights, factory-tuned weights, inference code and training code. A buyer may reasonably accept that training code remains proprietary while requiring delivery of the tuned weights and a local inference package. Where weights cannot be exported, the fallback should be a perpetual right to the last accepted version, an offline installer, complete configuration records and an escrow or continuity arrangement proportionate to the operational risk. The objective is continued inspection and migration, not unnecessary transfer of the supplier's entire research platform.

Local inference needs a precise contractual definition. Normal production must not require per-frame cloud calls, remote authorisation for every session or transmission of full production images. Remote support should be optional, logged and controlled by the factory. If connectivity is unavailable, the accepted version must continue to operate and store events for later synchronisation. Licence servers, update services and analytics dashboards should be listed as separate dependencies so that a claimed edge system cannot conceal a remote point of failure.

An open interface schedule should accompany the commercial schedule. It defines the quality-event fields, identifier formats, timestamp convention, defect dictionary, confidence representation, thumbnail reference, model version and operator disposition. It also defines export in a documented machine-readable form and the destination Performance-Management System tables. A bulk export offered only on request is weaker than a routine local interface available on every shift. The buyer should test a sample export before final acceptance.

Version control requires the same discipline as detection performance. The supplier identifies the accepted model, runtime and defect-mapping versions. Updates arrive with release notes and a rollback path. No update changes the operating threshold or defect definitions silently. The factory can test a candidate on a held-out sample and run it in shadow mode before promotion. The contract should also require notice when a component or model reaches end of support, with enough time to procure a replacement or migrate.

Service terms should describe spares, response routes, local technician training and the division of responsibility for optics, mechanics, computing and software. Acceptance payments follow measured gates, with a retention or warranty mechanism for faults that appear only during sustained production. Exit terms include a final export, removal of supplier access, delivery of configuration and weights where agreed, and a practical handover test. These clauses are procurement recommendations, not legal advice; local counsel should adapt ownership, liability, escrow and continuity provisions to the governing law and the actual transaction.

9.6

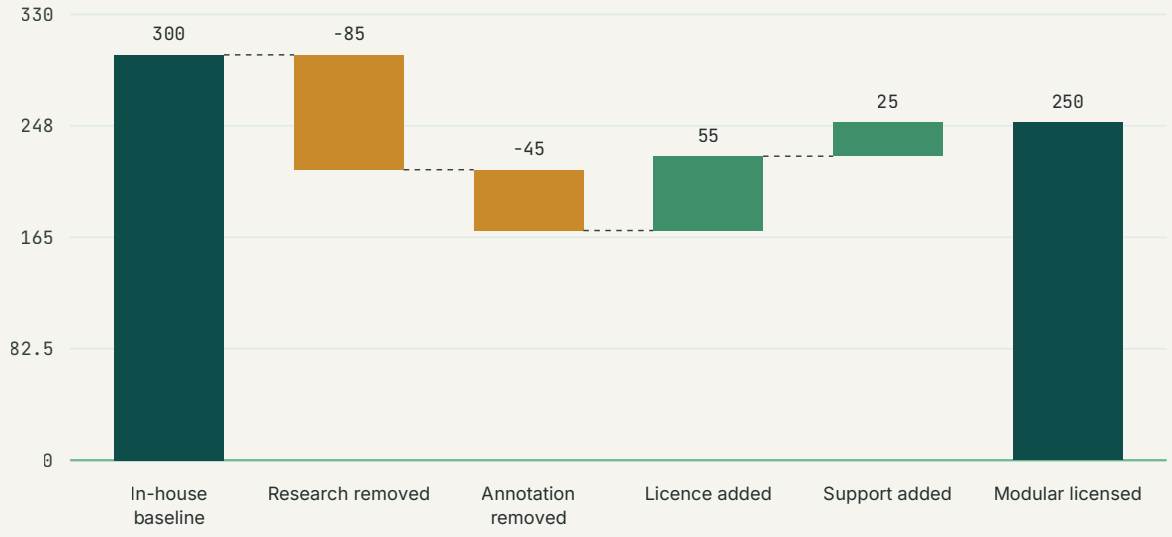
Five-year cost comparison must include delay and exit, not just licence price

Total cost of ownership should compare the same inspection boundary over the same period. For the paper's illustrative 3.0 million garment a year plant, that boundary covers the specified roll and garment inspection stations, while station count remains a quotation input. Hardware, cameras, illumination and mechanical frames belong in every route where they are required. Model research, annotation, integration, licences, support, technician time and updates then appear as separate lines. The model should also expose delay before accepted inspection, production disruption during changeovers, and exit costs. Omitting the common hardware makes a bespoke build look artificially cheap; omitting integration makes a modular licence look artificially complete.

Exhibit 44, p. 147 uses illustrative cost units to show the method without pretending to have supplier quotations. It begins with an in-house five-year baseline of 300 units, removes bespoke research and initial annotation that are not required in the licensed route, then adds licence and vendor-support costs. Common hardware, integration, maintenance and delay exposure remain inside the balance.

EXHIBIT 44 / ILLUSTRATIVE FIVE-YEAR COST UNITS

Licensing removes research cost but retains a five-year support floor



ILLUSTRATIVE Illustrative model, not a quotation or forecast. In-house baseline 300 = research 85 + annotation 45 + common hardware and integration 80 + maintenance 50 + delay exposure 40. Modular licensed closing 250 = 300 - 85 - 45 + licence 55 + vendor support 25. All other components remain unchanged.

The arithmetic yields 250 cost units for the modular route, 50 units below the illustrative in-house baseline. The absolute saving is not the main conclusion. More important is where uncertainty sits. Research and initial annotation are removed before acceptance, while licence and support payments can be tied to a working

system. The model deliberately leaves delay exposure unchanged, so it does not manufacture an advantage from an assumed faster launch. A real procurement model should test that variable explicitly.

Delay exposure needs a contribution basis rather than an arbitrary penalty. An illustrative sensitivity could assign two cost units of monthly contribution at stake and compare a six-month delay with a three-month delay. The difference is 2 x (6 - 3) = 6 cost units. This is not a garment-margin forecast. It demonstrates how to convert commissioning time into the same unit framework as licences and support. The plant's actual model should use contribution from the relevant production volume, net of variable costs and constrained by demonstrated saleable capacity.

The comparison also needs downside cases. A low licence quote may exclude the second camera, additional defect classes, offline operation, local support or data export. A turnkey quote may include these items and therefore be less expensive once the modular integration is fully loaded. An in-house route may win when its research produces a measurable recall advantage on a defect class with high escape cost. Procurement should request sensitivity tables for fabric variety, station count, support term, update frequency and exit year rather than defend one central total.

Residual value and reversibility matter at the end of the horizon. Standard cameras and edge computers may serve another station or supplier. A proprietary frame may have little alternative use, while a portable dataset retains value across model changes. Conversely, an internally developed model has limited value if the team leaves and no documentation exists. These are not reasons to assign invented resale percentages. They are reasons to record what remains usable, who can operate it and what migration would cost.

This cost analysis is not investment advice. It is a procurement model to be populated with dated quotations, local tax treatment, financing assumptions and the factory's measured quality economics. The following chapter examines the operating savings themselves; the present decision is narrower. It asks which acquisition route reaches verified inspection with the least avoidable research exposure while preserving future control.

9.7

Factory-tuned weights should replace the baseline only after evidence matures

The licensed baseline is the beginning of local learning, not its end. Every verified defect and false alarm can improve the factory's understanding of its fabrics. Reverse-distillation research has evaluated unsupervised fabric anomaly detection across MVTec AD, AITEX, TILDA and a textile-manufacturing dataset, showing the relevance of methods designed for patterned textures and domain transfer. [S12] Network-pruning research further shows that trained fabric models can be optimised for embedded hardware, with dataset-specific rather than uniform accuracy outcomes. [S02] These findings support a later tuning stage, not an immediate from-scratch programme.

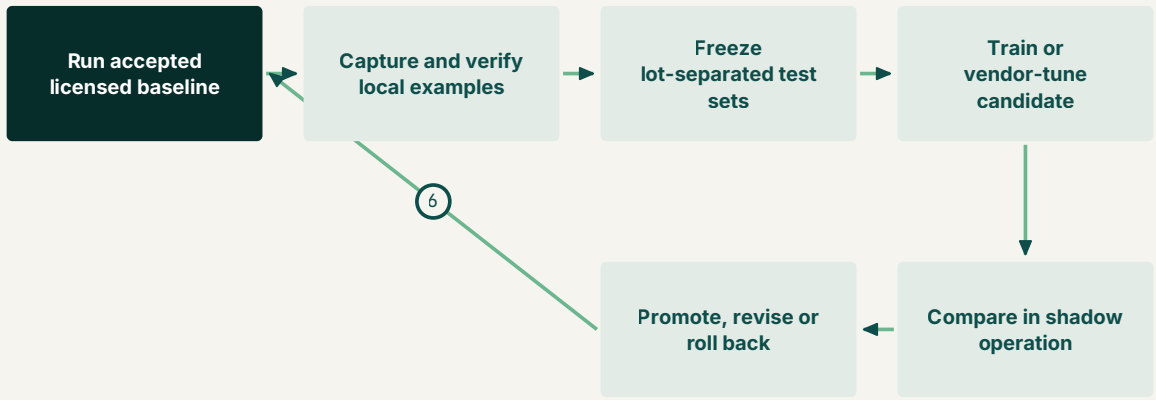
An illustrative planning window of six to twelve months can provide a review cadence, but the calendar must not be the promotion gate. The useful trigger is dataset coverage: enough independently verified examples across priority defect classes, fabric families, lots, shifts and operating conditions. Rare critical classes may remain unsuitable for supervised tuning even after a year. They can continue under one-class screening and human review while better-represented classes move to a local detector.

Data preparation determines whether tuning evidence is credible. Deduplicate near-identical frames, resolve operator disagreements and preserve the original disposition history. Split training, validation and final test sets by roll or production lot so that the same fabric run cannot appear on both sides of the evaluation. Keep a time-based holdout to test new production. Record optical and model settings with every sample. Without these controls, a larger dataset can merely produce a larger leakage problem.

Exhibit 45, p. 149 sets out the proposed transition. The candidate model does not replace the accepted baseline until it proves itself on frozen data and in shadow operation.

EXHIBIT 45 / GFT PROPOSED MODEL-TRANSITION WORKFLOW

Factory-tuned weights earn deployment through a gated comparison



6. retain on failure

ILLUSTRATIVE [S02] [S04] [S12] Illustrative workflow derived from cited cold-start, fabric anomaly and model-optimisation methods. It is a proposed control process, not a reported implementation at a GFT client.

The comparison uses the same class-level recall, false-alarm and speed measures as procurement. It also checks calibration by fabric family and verifies that a gain on common defects does not conceal worse performance on a critical rare class. Shadow mode records what the candidate would have done while the accepted system continues to control disposition. Promotion requires a named operational owner, a signed model package, updated documentation and a tested rollback. A failed candidate returns to analysis without interrupting production.

This stage changes the make-against-buy question. The factory now possesses something it lacked at launch: a verified local corpus and operating measures. It may ask the incumbent to tune the model, commission another supplier to run a challenger, or fine-tune an exportable model internally. The choice should be made on evidence and rights rather than habit. In all three cases, the camera and data spine can remain stable if the original interfaces were designed correctly.

Buy-before-build therefore does not freeze the factory into permanent dependence. It purchases a measurable starting capability, creates time to learn and reserves the option to build where local evidence justifies it. The discipline is to buy the reusable inference engine, own the measurement system and make every later model earn its place on the line.

Buy the first detector, own the evidence, and build only when the factory's own data proves why.

CHAPTER SOURCE NOTES

S04 arXiv / IEEE CVPR. 2021-06-15.

S05 arXiv / ICPR. 2020-11-17.

R28 DataIntelo. Publication date unrecorded; inherited source-text check.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

S07 arXiv. 2026-08-12.

S14 Applied Sciences (MDPI). 2024-12-23.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

S19 Textile Research Journal / SAGE Publications. 2026-09-06.

S12 arXiv. 2024-01-04.

S02 Sensors (MDPI). 2021-10-25.

10

In a low-wage garment economy, computer vision pays back by saving fabric and preventing rework rather than replacing inspectors

Low nominal wages weaken the case for labour substitution. The investable return comes from material yield, avoided rework and protected capacity, each measured without counting the same garment twice.

CHAPTER 10
DECISION TEST

Low nominal wages weaken the case for labour substitution. The investable return comes from material yield, avoided rework and protected capacity, each measured without counting the same garment twice.

Computer vision earns its place on a low-wage garment floor by stopping valuable material and productive time from being lost, not by making inexpensive inspection jobs disappear. The correct investment boundary is therefore the whole disposition of a defect: what the camera finds, what a person confirms, what the cutting or sewing team does next, and what loss is actually avoided. GFT recommends a base case with no inspector layoffs, no unverified buyer-penalty savings and no margin from recovered capacity unless demand can absorb it. Fabric recovery, prevented rework and fewer seconds must each reconcile to a physical record.

The central illustrative model in this chapter reaches a simple payback of 15.6 months on two inspection stations. Fabric recovery supplies 70.3 percent of its gross annual benefit, while conditional inspector reallocation supplies 2.9 percent. Removing the fabric bucket extends the same model's payback to about 203 months. These are constructed results from stated assumptions, not a GFT installation result, a supplier quotation or a forecast for Cameroon. Their purpose is to show which measurements decide whether the investment works.

That distinction also corrects a common modelling error. A higher defect-detection score is not itself cash. It becomes cash only when an earlier decision avoids a cut panel, an unpicked seam, a downgrade or a verified external failure. Conversely, a fast model can destroy value through false alarms, unnecessary recutting and review queues. The economic case must therefore sit beside the model's confusion matrix and the factory's material ledger, not inside a sales presentation about accuracy.

10.1

Low nominal wages make labour substitution a weak investment thesis

The inherited evidence register records Cameroon's XAF 41,875 monthly benchmark for 2023 rising to XAF 43,969 from 23 February 2024, approximately US$76 at the register's rounded conversion. The opened decree text clarifies the scope that a simple wage comparison can conceal. Its 2023 tiers were XAF 41,875 for state agents governed by the Labour Code, XAF 45,000 for agriculture and related activities, and XAF 60,000 for other sectors. The later XAF 43,969 figure concerns the state-agent tier, not a universal floor that a private garment investor should silently apply. [R14] [S1002]

The base case here consequently uses XAF 60,000 a month as the residual private-sector floor, pending confirmation of the establishment's legal classification, actual pay grades and any applicable collective agreement. This is a modelling boundary, not legal advice. [S1002] Real employer cost will also differ from the statutory gross wage through payroll charges, allowances, supervision, attendance and overtime. Those items require a payroll schedule rather than an assumed multiplier.

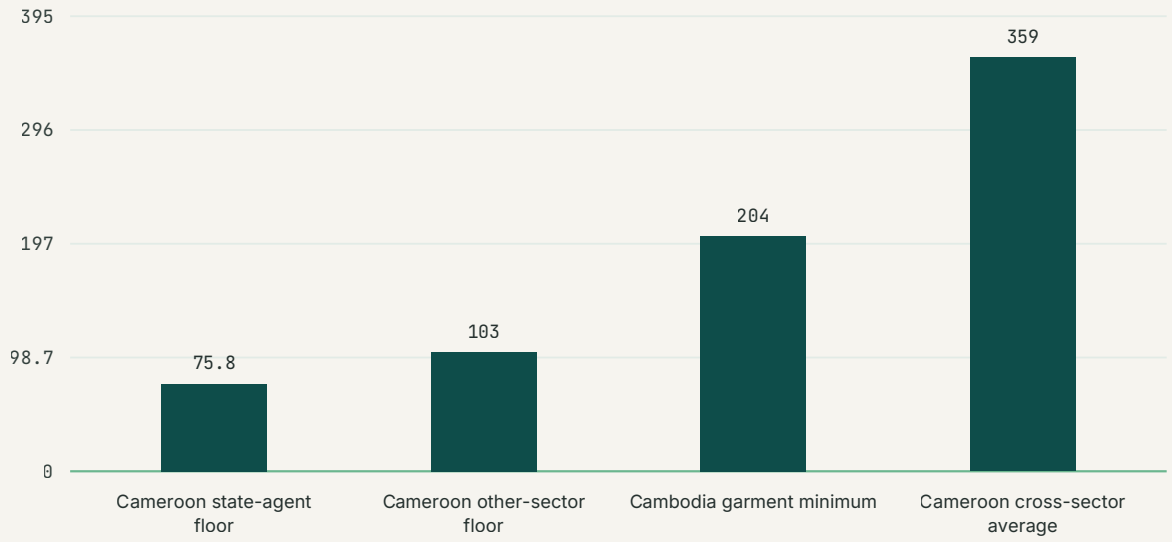
The comparison with Cambodia remains useful even after that correction. The inherited register records a 2024 Cambodian garment-sector minimum of US$204 a month, while a separate cross-sector estimate for Cameroon sits near US$359 a month. The first is a sectoral statutory minimum and the second is an average across

sectors. Neither should be read as an inspector's actual take-home pay or total employer cost. Their value is to test whether a proposed camera investment depends on a wage arbitrage that the evidence does not support. [R15] [R17]

Exhibit 46, p. 153 places the inherited and correctly scoped figures on one approximate dollar scale solely for readability. It uses an illustrative planning conversion of XAF 580 per US dollar, close to the ratio implied by the register's rounded XAF 43,969 and US$76 pair. This is not a current foreign-exchange quotation. The original CFA franc amounts govern the model, and the different dates and wage definitions remain visible. [R14]

EXHIBIT 46 / US DOLLARS A MONTH, ILLUSTRATIVE FX

Low monthly wages cannot carry a camera case based on headcount removal



DERIVED [R14] [R15] [R17] [S1002] Derived comparison, not a current FX quote or like-for-like pay survey. Cameroon state-agent figure XAF 43,969 / 580 = 75.8, consistent with R14's rounded US$76. Cameroon other-sector floor XAF 60,000 / 580 = 103.4 from the opened 2023 decree text. Cambodia garment-sector minimum US$204 for 2024 and Cameroon cross-sector average US$359 retain their different dates and definitions.

A simple annual calculation shows the constraint. Four illustrative full-time inspector-equivalent roles at XAF 60,000 a month cost XAF 2.88 million a year before other employer costs. Even if all four roles disappeared, which this chapter does not recommend, their gross wages alone would take about 17.4 years to repay an illustrative XAF 50 million installation before any operating expense. With the model's XAF 12 million annual operating allowance, labour removal on that scale never pays back. These calculations demonstrate why headcount substitution cannot be the main thesis.

The operating evidence also argues for retaining human review. The low-cost edge study positions its system as assisted triage, not autonomous acceptance. At its recall-prioritised threshold, 85 percent of forwarded frames were false alarms in the reported small benchmark. A factory that budgets no verification labour merely because inference is automatic would omit a real workload. The economic design should reduce repetitive scanning while preserving accountable disposition, borderline review and feedback into the process. [S07]

Labour still has economic value, but it should enter as a conditional reallocation benefit. An inspector who spends less time watching acceptable fabric can verify exceptions, investigate recurring faults, coach operators or strengthen final review. That time is valuable if the reassigned work removes another constraint. It is not automatically a cash saving, and it is not a redundancy plan. The model therefore counts only half of four illustrative roles as realisable capacity value and tests the result without that bucket.

10.2

Fabric recovery carries the largest addressable value

Material is the stronger economic candidate because it is consumed whether or not the factory has spare demand for its labour. A Tanzanian cutting study places material consumption at approximately 50 to 70 percent of total garment manufacturing costs. That is an external benchmark from another production system, not a measured Cameroonian cost share. It nevertheless explains why a small improvement in net material recovery can outweigh several low-wage inspection roles. [S1001]

The same study reports ten cutting tests with average efficiency of 78.67 percent, wasted pieces of 19.2 percent and unnoticed waste of 2.03 percent. Those categories belong to cutting consumption and should not be relabelled as defect losses. Vision cannot recover ordinary marker fallout, end loss or every spreading error. The investment case must isolate the subset of loss that better detection and disposition can physically change. [S1001]

The inherited inspection register summarises manual detection at 60 to 70 percent and computer vision above 90 percent. Using a 65 percent manual midpoint and a conservative 90 percent vision threshold creates a 25 percentage-point detection increment for scenario arithmetic. It does not create a 25 percentage-point saving in fabric cost. Only additionally confirmed defects whose revised disposition avoids net loss belong in the benefit calculation. [R29]

For incoming fabric, GFT uses the relationship B_f = Q x d x (c_v - c_m) x r x F. Q is garments produced in the period, d is the fabric-origin defect share on a garment-equivalent basis, c_v and c_m are verified vision and manual capture rates, r is the net recoverable fraction of value on newly intercepted units, and F is fabric cost per garment equivalent. The term r is the economic heart of the equation. It must absorb splice losses, marker constraints, recutting, handling and any supplier recovery rather than assume that every detection saves a whole garment. R29 informs the capture inputs and S1001 informs the material-cost context; the remaining central-case inputs are illustrative. [R29] [S1001]

Three physical routes can create that recovery. A severe fault can be removed before spreading, with the resulting splice and end losses recorded. A marker can sometimes move a panel away from a mapped fault when product rules permit. A verified supplier defect can support a claim or replacement when the contract and evidence allow it. These routes have different values and must not be added together for the same material. If a recut panel consumes new cloth, the net benefit is the avoided downstream loss minus the replacement material and work.

GFT NOTE

Approve the investment against a weekly fabric-disposition reconciliation, not against detections. For every newly intercepted fabric fault, record inspected length or garment equivalents, the manual comparison result, the cutting instruction, replacement cloth, recovered supplier value and final panel outcome. Use one disposition per physical unit. Count only the net avoided loss after splice, recut, handling and claim costs. Keep ordinary cutting waste in a separate ledger so marker inefficiency cannot masquerade as vision value.

False alarms need their own debit. An acceptable roll flagged as faulty can consume review minutes, trigger an unnecessary splice or remove usable area from a marker. A no-read or lost roll coordinate can also force manual reinspection. These outcomes reduce r and may make it negative on some styles. The model should therefore report recovery by fabric family and disposition route, not one blended percentage that hides a failing print or a poorly transferred roll map.

The recovered amount should be valued at avoidable cost, not retail price. If a panel is recut before sewing, the relevant saving may be the sewing and finishing cost avoided plus any net cloth preserved, not the garment's full selling price. If the same cloth would have been scrapped at incoming inspection under either method, there is no incremental benefit. This counterfactual discipline prevents an accurate camera from receiving credit for a loss that the old process would also have prevented.

10.3

Rework and seconds destroy capacity as well as minutes

Rework is more than an extra wage entry. In the model's operating logic, a repaired garment returns to an already balanced sequence, occupies handling space and competes with new work for attention. The direct calculation should nevertheless begin with what can be observed: verified repair minutes, the people paid to perform them, and any additional material or machine time. A broad claim that every repaired garment doubles total factory cost would be neither necessary nor defensible.

A jacket-line Six Sigma study found broken threads and puckering together represented 23 percent of 16 defect types. The finding supports concentrating corrective work on a Pareto subset rather than spreading effort evenly. It does not mean that a Cameroonian line will have the same defect mix or that removing those classes will remove the same share of cost. The local model needs class-specific minutes and dispositions. [S25]

For direct rework labour, GFT uses B_r = Q x s x p x t x w. Here s is the exclusive share of production entering the rework-eligible cohort after fabric-recovery units have been removed, p is the proportion of that rework avoided through earlier control and verified corrective action, t is direct repair minutes per avoided rework unit, and w is labour cost per minute. This p is not the camera's detection increment. Detection may initiate the correction, but a mechanic's adjustment, changed work method or supplier action must actually prevent recurrence. All central-case rates are illustrative.

The wage conversion is explicit. At XAF 60,000 a month and a 40-hour week averaged over 52 weeks and 12 months, the illustrative unburdened wage is XAF 5.77 a minute, calculated as 60,000 / (40 × 52 / 12 × 60). R16 supplies the standard-week context and S1002 the private-sector wage floor. Actual payroll cost per productive minute should replace this figure, including paid non-production time where the accounting boundary requires it. [R16] [S1002]

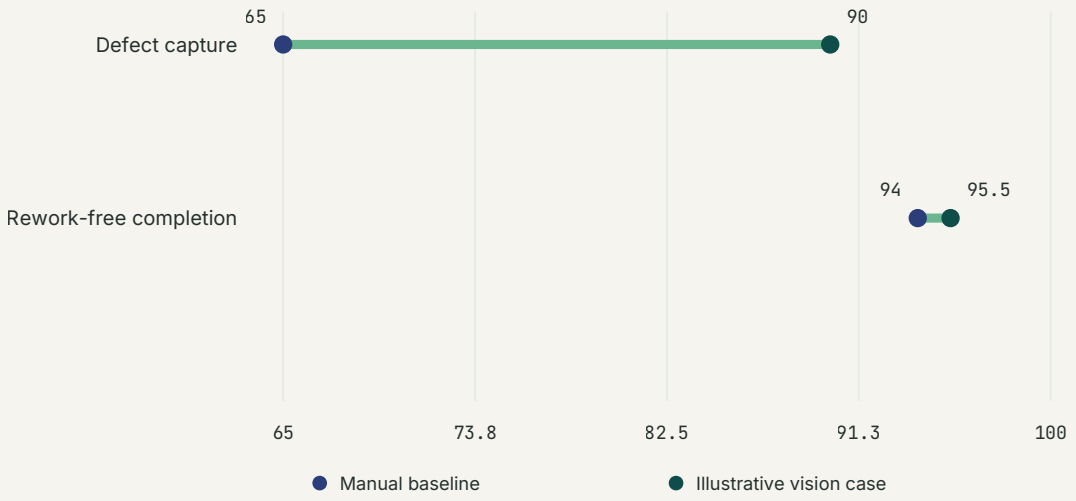
The saved minutes should not automatically receive a second benefit as increased output. If demand, upstream supply or another station limits production, released sewing time may not become saleable capacity. The base case therefore includes direct avoided repair labour but zero additional contribution margin from higher output. A separate capacity case can be added only when the production plan identifies unsatisfied demand and the true bottleneck. This prevents one repaired garment from generating both a full labour saving and an invented margin on a garment the factory cannot sell.

Seconds require a different ledger because downgrade loss is not the same as repair time. GFT defines B_s = Q x a x g x L, where a is the exclusive second-quality cohort after the previous buckets, g is the reduction associated with verified control, and L is the avoidable loss per downgraded unit. L should be the realised first-quality value minus second-quality recovery, less any selling and handling costs that differ between the routes. It is not the full selling price and not the fabric cost by default.

The illustrative dumbbell in **Exhibit 47**, p. 156 makes the asymmetry visible. Detection capture changes by 25 percentage points in the scenario, while rework-free completion changes by only 1.5 points because the assumed exclusive rework cohort and avoided share are much smaller. The chart does not claim that better detection caused a measured yield improvement. It shows why a large model metric and a modest economic effect can both be true.

EXHIBIT 47 / PERCENT, WITH METRIC DEFINITIONS RETAINED SEPARATELY

A large detection gain does not translate one-for-one into first-pass yield



ILLUSTRATIVE [R29] Illustrative scenario. Defect capture uses 65 percent as the midpoint of R29's 60 to 70 percent manual range and a conservative 90 percent vision threshold against R29's above-90 summary. Rework-free completion assumes an exclusive 6 percent rework cohort and a 25 percent reduction: baseline 94.0 percent, case 95.5 percent. These are not observed factory results.

An Indonesian printed T-shirt case recorded 107 defective units among 1,380 produced, equivalent to 7.75 percent by direct division. It also linked those defects to material waste and found points outside statistical control limits. The case demonstrates why defect prevalence, process stability and waste should be measured together. Its 7.75 percent is not transferred into this chapter's base case; it informs the upper end of the sensitivity range only as an external observed example. [S26]

Capacity context deserves the same restraint. The inherited register reports CICAM plant utilisation between 10.83 and 48.9 percent in a July 2023 account and an underutilisation cost of about CFAF 2 billion in 2020. Those figures describe another organisation and several possible causes of underuse. They do not establish a recoverable amount for vision or a utilisation baseline for a new garment plant. The lesson is to protect effective capacity with measured first-pass quality, not to attribute every idle minute to defects. [R12]

A Central Java garment study provides a further boundary. Digital quality control and production automation jointly explained 64.5 percent of product-quality variance in its structural model, with reported associations of 0.262 and 0.639 respectively. These questionnaire-based results support integrated quality and production management. They do not price a defect, prove causation for another site or justify a fixed percentage increase in Overall Equipment Effectiveness. [S24]

The practical COPQ ledger should therefore have separate rows for fabric recovery, direct rework labour, extra rework material, downgrade loss, inspection review, false-alarm disposition and capacity effects. Each row needs a unit, an owner and a counterfactual. A garment appears in one primary savings row, while any genuinely separate additional cost is shown as a component rather than another full-garment benefit. That structure is more useful than a single Cost of Poor Quality percentage whose denominator changes whenever production ramps.

10.4

Buyer-failure savings remain contingent evidence, not a payback plug

Institutional uniforms and repeat domestic orders can make quality failures commercially serious, but the base case should not invent a chargeback schedule. A rejected lot, an emergency replacement, a return or a lost tender has a different contractual cause and cost. None belongs in annual benefits until the factory has buyer terms, an observed baseline and a defensible estimate of how the inspection change alters the event probability.

The Ethiopian trouser study reports image processing of 0.45 seconds a sample and a substantial throughput improvement over manual inspection. That result supports testing faster classification. It does not demonstrate elapsed garment handling, complete surface coverage or prevention of buyer-lot rejection. The economic station record must add presentation, review, diversion, re-presentation and queue time before claiming a production-rate benefit. [S19]

Final acceptance and process control also answer different questions. A sampled buyer audit can accept or reject a lot after production, while inline vision can flag a changing process before every item is complete. A camera may reduce the chance of defective work reaching the final audit, but the saving arises only if the factory avoids a real downstream disposition. Better inspection coverage alone is not proof of fewer claims.

For that reason, the central model assigns zero value to avoided AQL failure, air freight, returns, reputational damage and future tender access. Zero does not mean those effects are unimportant. It means their frequency and unit costs are unknown here. When evidence becomes available, the external-failure calculation can use lots exposed, the measured change in verified rejection probability and the actual incremental cost per event, with legal and commercial review of the contract terms.

A similar rule applies to quality premiums. A buyer might pay more for documented inspection, or a factory might win work it would otherwise lose, but those are commercial scenarios rather than automatic vision benefits. Counting both a higher price and the full margin on newly won volume would overstate the inspection effect if product design, delivery or relationship management also contributed. The attribution should be agreed before the sale is included.

GFT recommends an external-failure register linked to lot, buyer requirement, defect class, verified cause, disposition and incremental cost. The baseline should include accepted lots as well as failures so a rate can be calculated. Until that register supports a probability, management should treat avoided external loss as upside outside the payback decision. This makes the investment harder to exaggerate and easier to reconcile after commissioning.

10.5

An illustrative plant pays back only when material recovery is counted

The payback calculation begins with annual volume rather than nameplate capacity. The illustrative factory has a nameplate of 3.0 million garments a year and operates at 70 percent in the central case, giving Q = 2.1 million garment equivalents. Fabric cost F is XAF 3,000 per garment equivalent. The incoming fabric-origin defect share d is 5 percent, manual capture c_m is 65 percent and vision capture c_v is a conservative 90 percent. Only 45 percent of the value on newly intercepted units is recovered after all disposition losses. R29 informs the capture range and S1001 the material-cost context; the other inputs are explicit scenario assumptions. [R29] [S1001]

The illustrative sewing assumptions are an exclusive rework-eligible share s of 6 percent, an avoided share p of 25 percent and direct repair time t of 12 minutes a unit. The separate exclusive seconds cohort a is 3 percent, with an illustrative reduction g of 20 percent and downgrade loss L of XAF 900 a unit. Four inspector-equivalent roles are reallocated, with half their annual floor wage counted as useful capacity. These cohorts are defined as mutually exclusive for the primary savings calculation.

The capital allowance is an illustrative XAF 50 million for one incoming-fabric station and one finished-garment station, including cameras, illumination, edge compute, mechanical mounting, integration and commissioning. It is not a supplier quote and does not assert a market price. Annual operating cost is an illustrative XAF 12 million for licensing and support, calibration, replacement lighting and components, local technical time, review overhead and data retention. Procurement must replace both allowances with dated, scoped quotations and local workload measurements.

WORKED EXAMPLE: TWO STATIONS ON A 3.0 MILLION GARMENT NAMEPLATE

At the illustrative central volume of 2.1 million garments, annual fabric recovery is B_f = 2,100,000 × 0.05 x (0.90 - 0.65) x 0.45 x XAF 3,000 = XAF 35.438 million. This is the value remaining after the assumed recovery fraction, not 25 percent of the factory's fabric bill. R29 supports only the capture-range context. [R29]

Illustrative direct rework labour uses the XAF 5.769 unburdened minute cost derived above: B_r = 2,100,000 × 0.06 × 0.25 × 12 × 5.769 = XAF 2.181 million. Illustrative downgrade recovery is B_s = 2,100,000 × 0.03 × 0.20 x XAF 900 = XAF 11.340 million. These are exclusive disposition cohorts, so the model does not count the same garment as both full rework and full downgrade value.

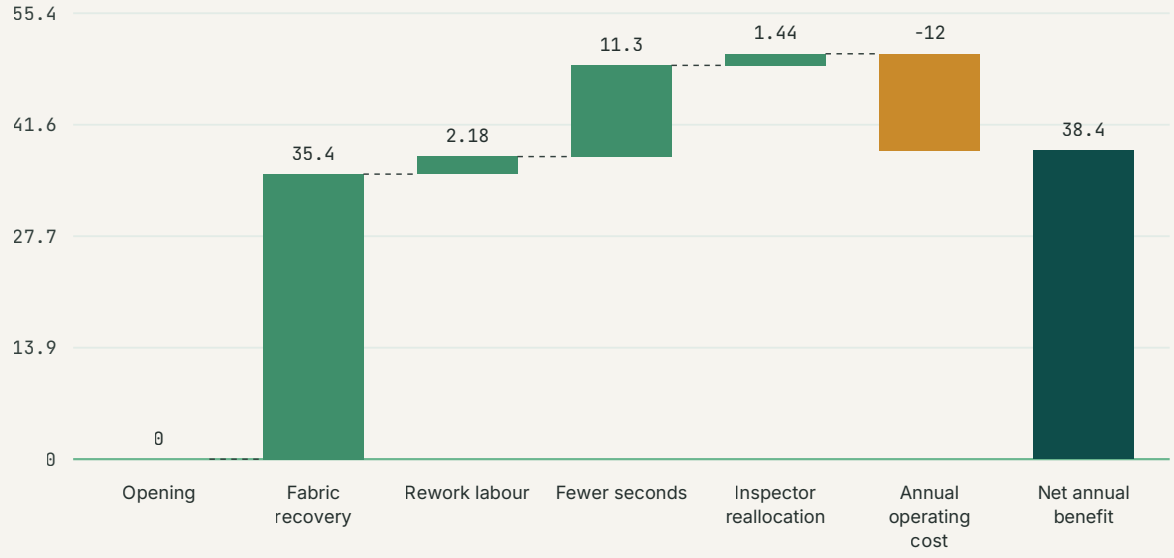
Conditional inspector reallocation is B_l = 4 x XAF 60,000 × 12 × 0.50 = XAF 1.440 million a year, using the private-sector floor in S1002. It assumes the released time performs useful work and does not assume redundancies. The illustrative gross benefit is therefore XAF 35.438 + 2.181 + 11.340 + 1.440 = XAF 50.398 million a year.

Subtracting the illustrative XAF 12 million operating allowance leaves XAF 38.398 million net annual benefit. Simple payback is XAF 50 million / XAF 38.398 million x 12 = 15.6 months. At the illustrative XAF 580 per US dollar planning conversion, gross benefit is about US$86,894 and net annual benefit about US$66,204. These dollar figures are translations of the same scenario, not additional benefits or a current FX estimate.

Exhibit 48, p. 160 reconciles those buckets. The waterfall starts at zero, adds each gross source, subtracts annual operating cost and ends at net annual benefit. It does not deduct the capital allowance again, because capital is used in the payback ratio. Treating both the full capital allowance and an annual depreciation charge as cash costs in the same simple-payback calculation would mix accounting boundaries.

EXHIBIT 48 / XAF MILLION A YEAR, ROUNDED TO THREE DECIMALS

Fabric recovery dominates the illustrative annual value bridge



DERIVED [R16] [R29] [S1001] [S1002] Derived illustrative scenario, not a forecast. Q = 3,000,000 × 70% = 2,100,000. Fabric = Q x 5% x (90% - 65%) x 45% recovery x XAF 3,000 = 35.4375. Rework = Q x 6% x 25% x 12 minutes x XAF 5.769231 a minute = 2.180769, with the minute cost derived from XAF 60,000 / (40 × 52 / 12 × 60). Seconds = Q x 3% x 20% x XAF 900 = 11.34. Reallocation = 4 x XAF 60,000 × 12 × 50% = 1.44. Gross = 50.398269. Net = gross - 12 operating cost = 38.398269. R29 informs capture only; S1001 informs material-cost context; R16 and S1002 inform wage arithmetic. Other inputs are assumptions and cohorts are exclusive.

The counterfactual tests reveal the model's dependency. Removing the illustrative XAF 1.440 million labour-reallocation bucket changes payback from 15.6 to 16.2 months. Removing the XAF 35.438 million fabric-recovery bucket leaves only XAF 2.961 million net annual benefit and extends simple payback to about 203 months. A labour-only case has XAF 1.440 million gross benefit against XAF 12 million annual operating cost, so it has no positive payback. These are scenario arithmetic, not measured returns.

The composition tells the same story. In the illustrative gross benefit, fabric recovery contributes 70.3 percent, fewer seconds 22.5 percent, direct rework labour 4.3 percent and inspector reallocation 2.9 percent. Management should therefore spend its measurement effort on recovery fraction, downgrade value and physical disposition before refining a small labour assumption to false precision.

Simple payback is only a screening measure. This illustrative model excludes discounting, tax, financing, working-capital effects, replacement cycles and terminal value. It holds volume and defect prevalence constant within each scenario, includes no external-failure saving and gives no value to data assets or future product premiums. A full investment review would add those financial dimensions only where evidence exists. This planning method is not investment, legal or tax advice.

10.6

Published benchmarks guide tests but cannot certify a factory's return

The Ethiopian trouser study estimates a 12 to 18 month return for medium-sized factories. The central illustrative result of 15.6 months sits inside that interval, but this is arithmetic coincidence rather than independent validation. The opened abstract does not disclose a capital and operating-cost bridge, and its image-classification station is not the same as this chapter's two-station allowance. Its estimate should be used to frame questions about scope and payback, not to mark the Cameroonian model as observed. [S19]

That distinction suggests a disciplined evidence ladder. First, establish the manual counterfactual with verified defect and disposition records. Next, demonstrate local capture and false-alarm performance on retained production examples. Then measure the physical outcomes that create value: recovered cloth, avoided repair minutes and changed downgrade destinations. Only after those measurements should management annualise benefits and compare them with quoted ownership costs.

The low-cost edge study reinforces the same gate from the operating side. Its triage design can reduce repetitive work, but its reported false-alarm forwarding shows why review and calibration remain cost centres. A vendor demonstration with high recall is not yet a labour budget. The factory should time the complete review workflow and include that workload in annual operating cost before accepting a payback claim. [S07]

GFT's procurement gate is consequently conditional. Scale only when the verified net benefit meets the sponsor's threshold under actual utilisation, the recovery fraction remains positive after all cutting consequences, and manual review can absorb the alarm load. If those conditions fail, change the optical setup, defect scope or station count before adding capital. Do not repair a weak operating case by adding speculative buyer penalties.

10.7

Sensitivity makes utilisation and recovery quality the investment gates

The illustrative sensitivity model retains every central assumption except utilisation and fabric-origin defect share. Annual volume changes with utilisation, fabric recovery changes with both volume and defect share, and the direct rework and seconds buckets scale with volume. Inspector reallocation and annual operating cost remain fixed. Payback is XAF 50 million divided by positive net annual benefit, multiplied by 12. This is a constructed comparative model, not a statistical confidence interval.

Exhibit 49, p. 162 shows the interaction. At 30 percent utilisation and a 2 percent incoming defect share, the fixed cost base leaves only XAF 1.310 million net annual benefit and an illustrative payback of 458 months. At 85 percent utilisation and an 8 percent defect share, the same assumptions produce 8.0 months. The wide range is precisely why a payback statement without volume, prevalence and disposition assumptions is not decision-ready.

EXHIBIT 49 / ILLUSTRATIVE SIMPLE PAYBACK, MONTHS

Low utilisation defeats payback even when defect prevalence rises

	2% defects	4% defects	6% defects	8% defects
30% utilisation	458	81.2	44.6	30.7
50% utilisation	65.1	31	20.4	15.2
70% utilisation	35	19.2	13.2	10.1
85% utilisation	26	14.9	10.4	8

Cell scale: 8 to 458

DERIVED [R12] [R16] [R29] [S26] [S1001] [S1002] Derived scenario matrix. For each cell Q = 3,000,000 x utilisation. Net annual benefit = Q x defect share x (90% - 65%) x 45% x XAF 3,000 + Q x 6% x 25% x 12 x XAF 5.769231 + Q x 3% x 20% x XAF 900 + XAF 1.44 million - XAF 12 million. Payback = XAF 50 million x 12 / net. All rates except the R29 capture context are illustrative. S26's 107 / 1,380 = 7.75% observed case motivates testing an 8% upper scenario, not adopting its prevalence. R12 supplies only a warning against transferring another plant's utilisation. Rounded to one decimal; negative net benefit has no payback.

At the central 5 percent defect share, changing only utilisation gives illustrative paybacks of 57.6 months at 30 percent, 34.5 months at 40 percent, 15.6 months at 70 percent and 12.3 months at 85 percent. A ramping factory should therefore avoid installing every station against nameplate volume on day one. A smaller initial scope can limit fixed operating cost while the factory learns which dispositions genuinely recover value.

Fabric price changes alter the same investment gate. Holding the other central inputs fixed, illustrative fabric cost of XAF 2,000 a garment equivalent produces XAF 26.586 million net annual benefit and 22.6-month payback. XAF 3,000 gives 15.6 months, and XAF 4,500 gives XAF 56.117 million net benefit and 10.7 months. The correct input is the avoidable cost of the actual material mix, including only those logistics and processing elements that change under the counterfactual.

The defect-rate axis needs equal caution. A higher observed defect share can increase addressable value, but only if the additional detections are real and recoverable. A difficult print that creates both more true faults and more false alarms may increase review, recutting and downtime faster than gross benefit. The recovery fraction r must therefore be remeasured by fabric family rather than held at the illustrative 45 percent because the central spreadsheet requires it.

The sensitivity results support three practical responses. At low utilisation, reduce fixed scope and prove one high-value route. At high material cost, prioritise incoming inspection and cutting coordination. At high sewing-rework prevalence, assign cor-

rective action to the verified station and class before expecting camera detection to become saved minutes. These are different operating interventions, and the station configuration should follow the dominant loss rather than a generic vision package.

The final investment record should preserve the assumptions beside the realised values each month. Replace nameplate volume with reconciled production, estimated capture with audited capture, assumed recovery with net disposition value, and allowance costs with invoices and measured review time. Report simple payback only after net annual benefit is positive and externally supportable. If the realised case diverges, change the operating decision rather than hiding the divergence inside a blended quality score.

Vision pays back in the cloth preserved and the work not repeated, not in the inspector removed.

CHAPTER SOURCE NOTES

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

S1002 Services du Premier Ministre, decree text reproduced by Maathis. 2023-03-21.

R15 wage.is. Publication date unrecorded; inherited source-text check.

R17 just-style. Publication date unrecorded; inherited source-text check.

S07 arXiv. 2026-08-12.

S1001 Tanzania Journal of Science, University of Dar es Salaam. 2021-12.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S25 OPSI, metadata and abstract through GARUDA. 2024-06-30.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

S26 Jurnal Teknik Industri (JATRI). 2026-09-25.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S24 CEMERLANG: Jurnal Manajemen dan Ekonomi Bisnis. 2026-08-03.

S19 SAGE Publications, abstract metadata supplied through Crossref. 2026-09-06.

11

Vision only improves quality when camera bounding boxes drive hourly line huddles and upstream corrective action

Counting defects without fixing their root cause merely automates scrap accounting; closed-loop integration wires edge cameras directly into the Performance-Management System.

CHAPTER 11
DECISION TEST

Counting defects without fixing their root cause merely automates scrap accounting; closed-loop integration wires edge cameras directly into the Performance-Management System.

Computer vision improves a garment factory only when a detection changes a decision before the same fault repeats. The destination is therefore not a model dashboard. It is the Performance-Management System, where a verified camera event becomes a defect count, an hourly exception, an assigned action and, eventually, evidence that the action worked. GFT's recommended closed loop has four controls: one quality-event schema from manual entry and camera inference, statistical signals that distinguish a process shift from ordinary variation, management cadences with named decision rights, and upstream feedback to cutting, maintenance and suppliers. A bounding box without those controls is only a more precise record of loss.

This chapter treats integration as an operating design rather than a software connection. It defines what the edge node should write, when a person must confirm it, how Defects Per Hundred Units and Right First Time remain comparable, and which meeting may change which setting. It also separates process control from customer acceptance. A stable process can still make unacceptable garments, while a noisy process can occasionally pass a final audit. The Performance-Management System must show both conditions rather than hiding them behind one score.

11.1

Deliverable 2 gives every detection a management destination

The programme's six KPI families, Production, Quality, Delivery, Cost, People and Sustainability, form one management record. Vision does not need a seventh family or a parallel scorecard. A verified sewing defect belongs in Quality. Its rework minutes affect Cost and People. A late corrective action affects Delivery. Fabric avoided through marker placement affects Sustainability. The same event can inform several families, but it should be captured once and referenced by a common identifier. That discipline prevents the camera supplier, quality lead and production manager from reporting different versions of the same shift.

The destination also determines the latency required. A model confidence score may be useful to a vision technician, but it cannot lead a line huddle. The huddle needs a small set of decisions: which defect rose, at which station, on which style or fabric lot, who is investigating, and whether production may continue. The event contract must therefore preserve operational context alongside the image result. Camera coordinates without a bundle identifier cannot guide cutting. A defect class without a station identifier cannot guide maintenance. A thumbnail without the final human disposition cannot be treated as confirmed truth.

Published garment evidence supports this management integration, while also warning against a simple technology claim. In a Central Java garment study, digital quality control had a reported path coefficient of 0.262 with $p = 0.002$, production automation had a coefficient of 0.639 with reported $p = 0.000$, and the two predictors jointly explained 64.5 percent of product-quality variance. These are study-specific associations from a questionnaire and structural model, not guaranteed effects for a Cameroonian factory. They support combined measurement and production control, not unattended camera acceptance. [S24]

Exhibit 50, p. 166 keeps those empirical anchors beside two other garment studies. Their measures are not interchangeable. Path coefficients, defect shares and defective-unit counts answer different questions. The exhibit deliberately retains those definitions instead of converting them into a single performance score.

EXHIBIT 50 / REPORTED STUDY RESULTS

Garment studies support integrated control rather than isolated detection

0.262

Digital quality control

Path coefficient; reported p = 0.002 in S24

0.639

Production automation

Path coefficient; reported p = 0.000 in S24

64.5%

Quality variance explained

Joint result for both predictors in S24

23%

Two dominant jacket defects

Broken threads and puckering among 16 types in S25

107 / 1,380

Defective T-shirts

Observed units in the S26 case

SOURCED [S24] [S25] [S26] S24 reports SEM-PLS associations in one Indonesian garment company. S25 reports the share of two dominant categories among 16 jacket defect types. S26 reports defective printed T-shirts in another case. These results are not Cameroonian benchmarks or GFT operating outcomes.

A separate Six Sigma garment study found that broken threads and puckering together represented 23 percent of 16 defect types. The value of that finding is not the percentage itself. It is the management sequence: Pareto analysis concentrated attention, fishbone analysis organised possible causes, and the improve stage assigned countermeasures. Vision can make the Pareto current and visual, but the root-cause work still determines whether recurrence falls. [S25]

GFT NOTE

Connect the camera to the existing operating rhythm before adding analytics. GFT recommends one shared quality-event identifier, one defect dictionary and one action register across the edge node, Quality Log and huddle board. The first acceptance test is operational: a supervisor can open an hourly exception, see the verified image and context, assign a corrective action, and later prove closure without re-keying the event.

11.2

One Quality Log preserves meaning from spreadsheet to edge inference

The manual crawl phase and automated phase must write the same logical record. During manual operation, a quality lead enters units inspected, defective units, defect counts, defect class, station, style and disposition. When vision starts, the edge node supplies the detection fields and a person verifies exceptions according to risk. The storage technology may change from workbook to local buffer and central database, but the definitions of numerator, denominator, time period and responsibility must not change. Otherwise the apparent improvement could be nothing more than a change in accounting.

Three measures need explicit separation. Defects Per Hundred Units is total verified defects divided by units inspected, then expressed on a per-hundred-unit basis. It allows several defects on one garment. The defective-unit proportion is garments with at least one verified defect divided by garments inspected. Right First Time is units completed without repair or rework divided by units completed, expressed as a percentage. A garment with two defects counts twice in DHU but once in the defective-unit proportion. Mixing these measures would corrupt both the p-chart and the incentive calculation.

The recommended quality event carries a timestamp, roll or bundle identifier, garment or style identifier where available, station, defect class, severity, model confidence, image reference, coordinates, verifier disposition and final physical disposition. Model confidence is metadata, not a quality KPI. The verifier disposition records whether the detection was accepted, rejected or left pending. The physical disposition records what happened to the item: pass, repair, recut, downgrade or scrap. Keeping those two decisions separate stops a correctly detected fault from being mistaken for a completed corrective action.

An edge study of two-stage fabric inspection explicitly positions its system as assisted triage rather than autonomous acceptance. That boundary matters for the Quality Log. Clearly acceptable frames may pass under an approved rule, while flagged or ambiguous cases receive human review. The log must then record the review state, because unreviewed model alarms are not equivalent to verified defects. [S07]

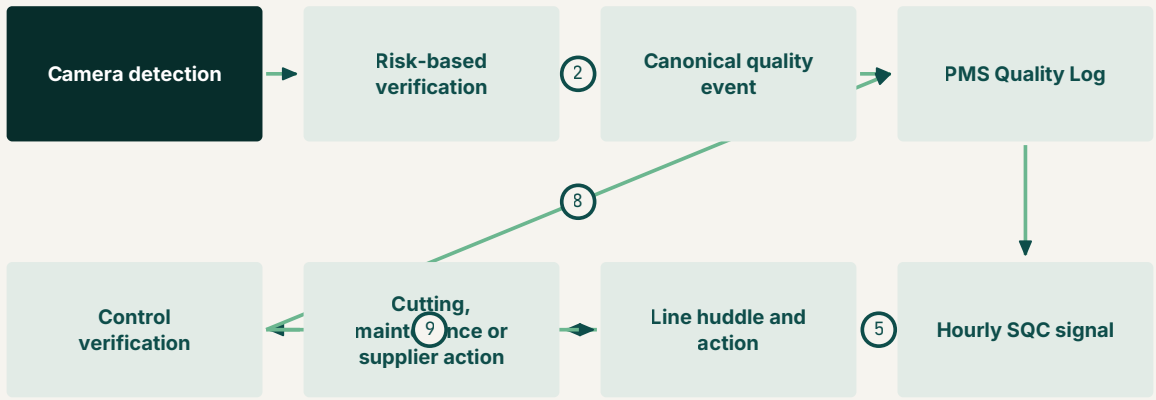
The inherited inspection register summarises manual defect detection at 60 to 70 percent accuracy and computer vision above 90 percent across more than 40 defect types. Even where a factory reproduces that detection gain, it describes classification

performance rather than process improvement. The PMS must show whether verified detections lead to containment, corrected settings and fewer repeat events. Otherwise a better detector can simply produce a longer list of unresolved faults. [R29]

Exhibit 51, p. 168 shows the proposed path from image to controlled process. The return path is essential. If an action does not restore the signal, the event remains open and escalates. If it does restore the signal, the closure evidence becomes part of the control plan.

EXHIBIT 51 / PROPOSED OPERATING FLOW

A detection becomes valuable only after action and verification close the loop



2. confirmed disposition
5. exception only
8. closure evidence
9. if still unstable

DERIVED [R29] [S07] [S24] [S25] [S26] GFT system design derived from cited vision triage, digital quality, DMAIC and SQC methods. It is a proposed architecture, not evidence of an implemented GFT result.

Exhibit 52, p. 169 translates that flow into a compact field contract. It is an illustrative GFT design, not a claim that a factory has already implemented these fields. Its purpose is to make ownership and review cadence inspectable before software procurement begins.

EXHIBIT 52 / ILLUSTRATIVE PMS FIELD MAPPING

One event can inform six KPI families without creating six records

Verified output	Shared field	KPI family	Decision and cadence
Units inspected	units_inspected	Quality	Supervisor checks denominator hourly
Defect count and class	defect_count, class_code	Quality	Quality lead reviews Pareto daily
Units without rework	rft_numerator	Production, Quality	Line lead reviews RFT hourly
Repair minutes	rework_minutes	Cost, People	Production manager reviews daily
Late action flag	action_due_state	Delivery	Plant review escalates weekly
Fabric avoided or scrapped	fabric_disposition_mass	Sustainability, Cost	Cutting and purchasing review weekly

ILLUSTRATIVE [S24] Illustrative GFT schema and decision-rights design. S24 supports integrated digital quality management generally, but does not specify these fields or cadences.

The schema also needs rejection rules. An event with an unknown station, an impossible timestamp, a missing denominator or an unresolved bundle should enter an exception queue rather than silently distorting the KPI. Late-arriving buffered events should retain their original production timestamp and carry a separate ingestion timestamp. That distinction lets the hourly chart reconstruct the actual production period while the interface honestly shows that synchronisation occurred later.

11.3 VERIFICATION STATES PROTECT THE KPI FROM PREMATURE ACCOUNTING

A quality event needs a lifecycle that everyone can read. A model-generated detection begins as unverified. Human disposition can confirm it as a defect, reject it as a false alarm or hold it for further inspection. A confirmed defect then receives a physical disposition and, where action is required, a linked corrective-action record. Closure is a later state supported by process evidence. Collapsing these states into a single detected flag would inflate confirmed DHU, hide pending work and make false alarms look like production failures.

State transitions also protect offline operation. Each edge event needs a durable unique identifier assigned at capture, so store-and-forward synchronisation cannot insert it twice. If a verifier changes a disposition, the system should append a revision with author, time and reason rather than overwrite the original decision. The KPI service reads the latest valid revision, while auditors can reconstruct why the measure changed. This is especially important when several stations reconnect after a network interruption and historical hourly buckets must be recalculated.

The Quality Log should publish its completeness state beside its values. An hourly chart based on verified events alone may look reassuring while a large verification queue remains unseen. The supervisor therefore needs counts for captured, verified, rejected and pending events, plus any periods with missing denominators. Manage-

ment can then distinguish a genuine low-defect hour from an hour whose evidence is incomplete. This small governance feature prevents operational pressure from turning missing review into apparent quality.

11.4

Statistical control turns hourly counts into a disciplined andon

A live defect total is not yet a control signal. Production volume changes by hour, and several defect classes can rise together when a new style enters the line. The p-chart addresses one precise question: is the proportion of defective units reasonably consistent with the established process, or has something changed? NIST defines the proportion nonconforming as nonconforming items divided by all items inspected and grounds the p-chart in the binomial distribution. The chart is therefore built on defective units, not total defects. [S1101]

For subgroup i, the system calculates p_i = D_i / n_i, where D_i is verified defective units and n_i is units inspected. With a stable baseline proportion p-bar, the upper control limit is p-bar plus three times the square root of p-bar(1 - p-bar) / n_i. The lower limit uses the same term with subtraction and is floored at zero. When hourly volume changes, the limits change with n_i. This matters on a garment floor: fifty defective units has a different meaning in a small hour than in a large one, and the chart must retain that denominator. [S1101]

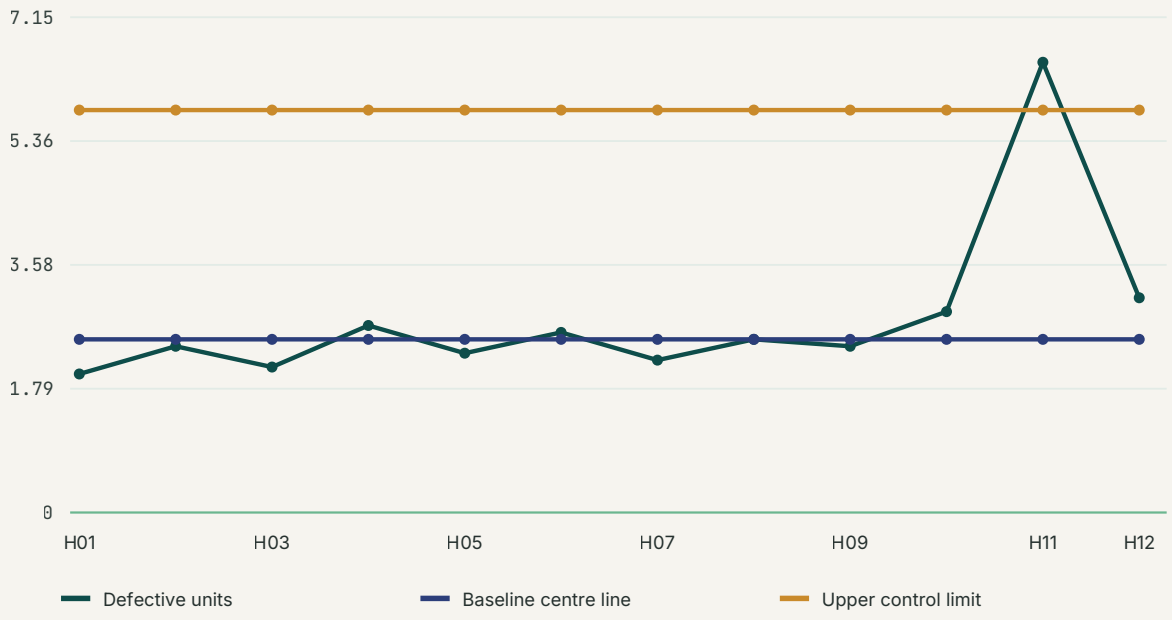
Control limits are not specification limits. A point inside the limits suggests a stable process; it does not prove that the defective proportion satisfies a buyer. Conversely, a point outside the limits calls for investigation even if the final sampled audit happens to pass. The andon reaction should therefore be triggered by process evidence, not postponed until an end-of-line acceptance decision. This separation also prevents a perverse response in which supervisors widen control limits merely to reduce alarms.

The baseline itself needs governance. It should come from a defined reference period with comparable products, verified dispositions and operating conditions. Known special causes, such as a failed needle or an uncalibrated light, should be investigated before their data are used to describe common-cause variation. After a confirmed process improvement, the team may recalculate the centre line under change control. It should never redraw limits continuously around a deteriorating process, because that would normalise the deterioration.

Exhibit 53, p. 171 illustrates the mechanism with a fixed hourly denominator and a baseline from an earlier stable period. Only the eleventh hour crosses the upper limit. The correct response is not to declare the whole shift defective. It is to contain suspect work, inspect the station and material context, and record whether the next verified observations return to control.

EXHIBIT 53 / PERCENT OF UNITS DEFECTIVE BY HOUR

The andon fires on a process signal rather than on every defect



ILLUSTRATIVE [S1101] Illustrative model. Each hour has n = 200 inspected units. The assumed stable baseline is p-bar = 0.025. UCL = 0.025 + 3 x square_root(0.025 × 0.975 / 200) = 0.0581, shown as 5.81 percent. LCL is floored at zero. Hour 11 is 6.5 percent and signals; the other constructed points remain below the limit.

WORKED EXAMPLE: ONE HOURLY ANDON AND THREE DISTINCT KPIS

An illustrative sewing station inspects 200 garments in an hour. Thirteen garments have at least one verified defect, and those garments carry 16 verified defects in total. The defective-unit proportion is 13 / 200 = 6.5 percent. DHU is 16 / 200 × 100 = 8.0. If all 13 garments require rework and no other garment does, RFT is (200 - 13) / 200 × 100 = 93.5 percent. These are constructed values used only to demonstrate the calculations.

Using the illustrative stable baseline of 2.5 percent and n = 200 from **Exhibit 53**, p. 171, the upper control limit is 5.81 percent under the NIST p-chart formula. The observed 6.5 percent exceeds it, so the andon opens an exception. The supervisor holds the suspect bundle, checks the needle, thread and tension settings, and assigns a mechanic. The line does not wait for the weekly Pareto review. The case is illustrative and does not represent a measured factory result. [S1101]

A useful andon policy also suppresses noise without hiding risk. One high-severity critical defect may require immediate containment regardless of statistical significance. Repeated minor defects may require a run-rule or a confirmed class-specific pattern before stopping the line. A high false-alarm rate may indicate an optical or model problem rather than a production shift. The response matrix should therefore combine severity, statistical evidence and verification state. It should not reduce every condition to the same red light.

The Indonesian SQC and FMEA case provides a relevant warning. Its p-charts still showed points outside the control limits for faded and misaligned prints, leading the authors to judge that the process was not fully under statistical control. The study then used risk priorities and work-standard recommendations. The operational lesson is the sequence from signal to cause to action, not the transfer of its defect rates to another factory. [S26]

11.5

Management cadences assign a decision to every timescale

Closed-loop control fails when every exception is sent to every meeting. Hourly, daily, weekly and monthly reviews have different decision rights. The hourly line conversation protects current production. The daily huddle removes recurring obstacles and confirms ownership. The weekly plant review tests causes across lines and authorises changes to maintenance, work methods or supplier handling. The monthly sales and operations planning review addresses capacity, purchasing and product-mix consequences. Escalation should move an unresolved decision upward, not merely repeat the same chart in a larger room.

The hourly review is deliberately narrow. The supervisor sees the current andon, the affected station, the verified thumbnail and the suspect bundle. The team chooses containment, immediate repair or continued observation under a defined rule. It records the action and expected verification time in the same event thread. If no one can act, the exception escalates with its evidence intact. The camera has then shortened the path from detection to responsibility without pretending to diagnose the cause itself.

The daily line huddle, planned as a 15-minute illustrative operating routine, reviews the prior period's largest verified classes, unresolved actions and denominator quality. Thumbnails help the team discuss the same physical condition, but the chair must ask what changed in the process. A rise in skipped stitches may prompt a needle and tension check. A rise in shade exceptions may prompt a roll-lot and illumination check. The huddle closes only actions with verification evidence, not actions with a completed meeting note.

The weekly review compares lines, styles, shifts and fabric lots under consistent definitions. It is the natural home for DMAIC selection and FMEA revision because those tasks require cross-functional evidence. A mechanic, quality lead, production supervisor and cutting representative can test whether the same defect follows a machine, material lot or method. The monthly review then uses confirmed trends for supplier discussions and capacity planning. It should not chase individual hourly points that the line has already resolved.

A simple action record prevents meeting theatre. It contains the originating quality-event identifiers, suspected cause, chosen countermeasure, accountable role, due time, containment status and verification result. Reopening is preferable to a falsely closed action. If the same class returns after closure, the system links the recurrence to the earlier action so the weekly review can reject the presumed cause rather than repeat it.

Decision rights should be written before the first alarm. Operators may contain and perform approved immediate corrections. Supervisors may authorise a line-specific response within the reaction plan. Maintenance changes to controlled settings require the designated technical role. Supplier claims, specification changes and production-plan trade-offs belong with the appropriate plant or commercial review. This allocation prevents a huddle from making decisions it cannot sustain and prevents urgent exceptions from waiting for a meeting with no operational authority.

Each cadence should also publish a small unresolved queue, ordered by risk and age rather than by the latest timestamp. The chair begins with actions whose verification failed or whose containment remains open. New Pareto entries come afterwards. That ordering stops a fresh stream of images from displacing yesterday's unfinished corrective work. It also creates an honest escalation route when the same constraint crosses several management levels.

11.6

Roll coordinates move corrective action upstream of the needle

Incoming fabric inspection creates value before sewing only when its coordinates survive the journey to the cutting table. A line-scan detection should be stored against roll identifier and web position, with cross-web and along-web coordinates under a declared measurement convention. The cutting system then translates that roll map into the current spread and marker. An operator can splice a severe fault, shift a marker where the rules permit, or place an acceptable minor flaw outside a critical panel zone. The decision is physical and immediate, not a supplier report read after the bundle has been sewn.

The interface must respect changing material states. A roll becomes a spread, a spread becomes plies, and cut panels become bundles. Each transformation needs a new identifier linked to its parent, because roll coordinates alone no longer locate a panel after cutting. Splice removal, end loss and marker movement alter the mapping. The practical design is therefore a chain of transformations with explicit losses, not an image viewer attached to a roll. If the chain is broken, the system should mark the affected panels for verification rather than invent coordinates.

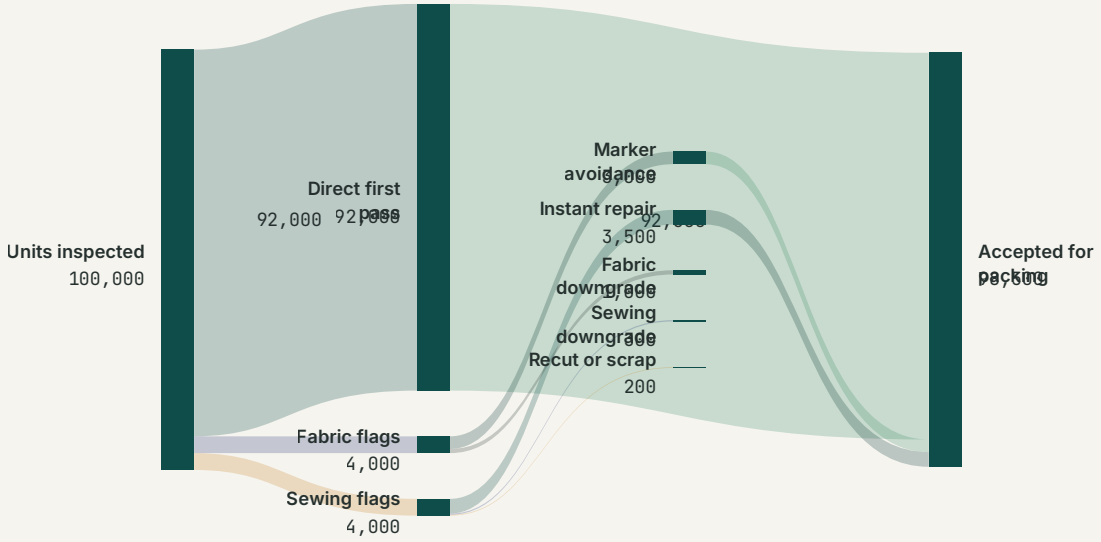
Severity rules should reflect product use. A small flaw that is acceptable in a hidden internal panel may be unacceptable on a visible uniform chest. The marker rule engine can encode those zones, but a cutting lead should approve exceptions during early deployment. This avoids two failure modes: discarding usable fabric through excessive caution and hiding known defects inside garments through aggressive nesting. Both outcomes would damage the economic case even if the detector remained accurate.

Upstream feedback also needs fair evidence. A supplier scorecard should separate confirmed fabric-origin faults from faults created during spreading, cutting or sewing. It should show inspected length, confirmed fault classes, severity and disposition before assigning a claim. Camera imagery strengthens the conversation, but it does not by itself prove when damage occurred. Joint review of roll condition, handling records and mapped coordinates is necessary where responsibility is disputed.

Exhibit 54, p. 174 shows an illustrative routing balance. It is not a forecast of savings. Its purpose is to force every inspected unit-equivalent into a declared destination and to show that detection, avoidance, repair and packing are different events. A real deployment would replace the constructed quantities with reconciled PMS records.

EXHIBIT 54 / ILLUSTRATIVE GARMENT-EQUIVALENT UNITS

Closed-loop routing separates avoidance and repair from final packing



ILLUSTRATIVE [R29] [S25] Illustrative conserved routing model using 100,000 garment-equivalent units to explain destinations, not an observed factory result. Input splits into 92,000 direct first-pass, 4,000 incoming-fabric flags and 4,000 inline-sewing flags. Marker avoidance returns 3,000 to packing, while 1,000 fabric units are downgraded. Instant repair returns 3,500 to packing, with 300 downgraded and 200 scrapped. Packing receives 98,500 units in total.

The same conservation discipline belongs in the daily report. Units inspected should reconcile to first-pass acceptance, verified exceptions and unresolved items. Exceptions should reconcile to repair, downgrade, recut, scrap and still-open dispositions. If those quantities do not balance, management has a data-integrity problem before it has a quality-improvement result. The system should surface the imbalance rather than force a tidy total.

11.7

DMAIC converts the Pareto into a tested cause and a held gain

The weekly Pareto is a selection tool, not a root cause. It tells the team where observed quality loss is concentrated under the current classification rules. Define then states the problem with a product, station, period and operational defect definition. Measure verifies the denominator, image review and measurement repeatability. Analyse tests possible causes. Improve changes a condition under controlled observation. Control holds the new setting through standard work, ownership and a reaction plan. Skipping from a camera Pareto directly to a machine adjustment would automate guessing.

The jacket DMAIC study illustrates why this discipline matters. Its Pareto stage identified broken threads and puckering as the dominant categories, and its cause analysis considered human, machine, material, method and environment factors before recommendations were made. Vision can improve the speed and consistency of the Measure phase. It cannot decide which of those five factor families explains a particular recurrence. [S25]

A practical root-cause board should therefore place evidence beside hypotheses. If skipped stitches cluster after needle changes, the team examines needle specification, installation and thread compatibility. If puckering follows one fabric lot, it examines feed, tension, shrinkage and material behaviour. If false alarms cluster near one luminaire or after lens cleaning, it examines the measurement system before touching sewing settings. Each hypothesis needs a predicted observable effect and a result that would disconfirm it.

FMEA adds priority where the Pareto alone is insufficient. Frequency matters, but so do severity and detectability. A rare critical seam failure may deserve action before a common cosmetic fault. The SQC and FMEA garment study used Risk Priority Numbers to rank faded print, inappropriate colour and misaligned print, then linked recommendations to work standards, equipment inspection and operator accuracy. The transferable lesson is to connect statistical evidence and risk ranking, not to copy its scores. [S26]

RPN should not become false precision. Teams should retain the individual severity, occurrence and detection assessments, the scoring scale and the rationale. A small change in one subjective score can reorder a ranking. Where the evidence is weak, the review should commission a check rather than hide uncertainty in multiplication. Vision may improve detection scoring for the classes it reliably sees, but it says little about defects outside its field of view.

Control is the stage that proves integration. The team states which setting changed, which chart should respond, how long the observation window runs and who may approve a return to normal operation. A short-lived improvement followed by recurrence is an open experiment, not a closed DMAIC project. The PMS should preserve both outcomes so later shifts do not repeat failed remedies.

11.8

OEE and incentives must reward prevention without punishing honest detection

Overall Equipment Effectiveness multiplies Availability, Performance and Quality. Vision primarily strengthens the information behind Quality through consistent RFT and verified defect records, but its interventions can also affect the other terms. A line stop lowers Availability while it occurs. Exception handling may reduce running speed. Those effects should remain visible. A programme that reports only quality improvement while hiding stopped time would overstate the production benefit.

The causal accounting needs care. An andon stop may reduce this hour's Availability while preventing a larger rework queue later. The PMS should therefore retain the stop event, its reason and the subsequent verified outcome. Management can then com-

pare immediate cost with avoided recurrence using actual production evidence. It should not assume that every camera stop was necessary or that every stop was harmful. The purpose is to improve the response policy over time.

Incentives create a second risk. Paying a line only for low recorded DHU can encourage reclassification, delayed verification or suppressed events. Paying only for output can encourage bypass of the andon. A fairer scorecard combines verified RFT, action closure, denominator completeness and agreed audit results. It distinguishes defects created on the line from incoming fabric faults and recognises prompt escalation of a genuine problem. The audit trail matters more than a single green dashboard tile.

Quality bonuses also need a measurement-system check. Supervisors should periodically compare camera dispositions, human verification and an independent sample under the same defect definitions. Disagreement is a reason to examine taxonomy, optics, training or model drift, not automatically to blame an operator. The next chapter addresses those human workflows in detail; here, the requirement is narrower: no incentive should depend on an unverified model score.

The closed-loop acceptance test is consequently behavioural and numerical. A planted or naturally occurring exception must reach the Quality Log with the correct denominator, create the intended andon, appear in the line review, generate an accountable action, and retain verification after intervention. The report must also reconcile units and expose unresolved items. Only then can leadership treat an improvement in DHU, RFT or OEE as an operating result rather than a change in what the software chose to display.

Evaluation should compare equivalent work, not merely an earlier week with a later one. Style complexity, fabric lot, order size and staffing can all change the opportunity for defects. The review should retain those covariates and use matched runs or a controlled pilot where practical. It should also report verification coverage and false-alarm burden beside RFT. Otherwise a favourable trend may reflect an easier mix or incomplete review rather than the corrective action. A claimed payback belongs only after those operating conditions and measurement checks have been made visible.

GFT's recommendation is to commission the loop before expanding model classes or camera count. One station with a reliable event contract, a working reaction plan and verified closure teaches the factory more than several disconnected detectors. Scale follows when supervisors use the record without duplicate entry, actions close with evidence, and upstream teams can trace their decisions back to the original detection.

A camera finds the defect; the management system must end its recurrence.

CHAPTER SOURCE NOTES

S24 CEMERLANG: Jurnal Manajemen dan Ekonomi Bisnis. 2026-08-03.

S25 OPSI. 2024-06-30.

S26 Jurnal Teknik Industri (JATRI). 2026-09-25.

S07 arXiv. 2026-08-12.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S1101 National Institute of Standards and Technology. 2026-04-22.

12

Computer vision succeeds on the shop floor when designed as human-AI triage that elevates supervisors rather than alienates inspectors

A greenfield Cameroonian factory faces a thin supervisor pipeline; human-in-the-loop vision turns scarce quality leads into high-impact process coaches.

CHAPTER 12
DECISION TEST

A greenfield Cameroonian factory faces a thin supervisor pipeline; human-in-the-loop vision turns scarce quality leads into high-impact process coaches.

The correct deployment unit is not a camera that replaces an inspector. It is a supervised decision cell in which software clears validated nominal work, people resolve exceptions, and supervisors use the resulting evidence to remove causes. This division protects release decisions while the model encounters new cloth, new seams and early production instability. It also gives inspectors a route into higher-value technical work rather than asking them to compete with a machine on continuous watching. The human role must include authority to question the screen, stop the process and improve the workflow. Evidence from industrial anomaly inspection supports exception routing, while International Labour Organization research supports worker involvement in decisions about how AI affects skills and control. [S34] [S1201]

12.1

Why autonomous pass-or-fail without human triage fails in early factory deployment

An early garment line is an unstable inspection domain. Fabric construction changes between orders. Dye lots alter the background. Lint, folding and roller vibration change the image. A model can therefore produce a confident score outside the conditions in which its threshold was validated. Confidence is a model output, not proof that the physical item conforms. A greenfield deployment must treat that gap as an operating risk, not as a footnote to a benchmark.

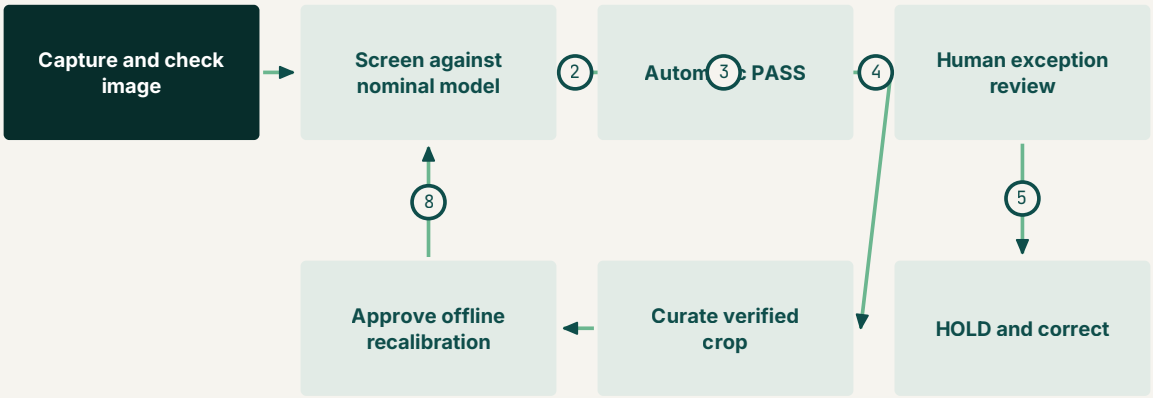
The manual alternative is also insufficient. The inherited textile-inspection evidence places manual defect detection at 60 to 70 percent accuracy, with 20 to 30 percent of defects missed. [R29] A metal-textile inspection study reports a wider manual detection range of 60 to 80 percent and explains why fatigue, repetition and individual variation prevent stable classification. [S34] These figures do not justify swapping one fallible decision-maker for another. They justify an architecture in which each handles the work it can perform most reliably.

That architecture has three decision lanes. Frames inside a validated nominal envelope may receive an automatic PASS. Flagged or ambiguous frames go to an inspector with the original image, the anomaly location and the relevant defect standard. A confirmed critical defect, an unresolved image or a failed station check goes to HOLD and escalates to the quality lead. Automatic rejection of a roll, bundle or garment lot should not be the default. Physical disposal and lot release carry commercial consequences that require an accountable person.

Exhibit 55, p. 180 makes the decision rights visible before any software is configured.

EXHIBIT 55 /

Human-AI triage reserves judgement for exceptions and corrective action



2. Validated nominal
3. Flag or ambiguity
4. Accept after review
5. Defect or unresolved
8. Controlled release

SOURCED [S34] [S07] Process synthesis of the STFPM human-in-the-loop inspection workflow and the edge cascade. Adapted as a GFT recommendation for fabric and garment inspection; it is not a claim that this exact workflow is already operating at a GFT facility.

The loop is deliberately asymmetric. Software may reduce the volume of routine screening, but it may not widen its own authority. A model update changes the screening function only after an approved test. A failed camera check routes work to review rather than allowing the last known model to continue silently. The inspector can disagree with an anomaly map, and that disagreement remains in the record. The supervisor can stop automatic PASS when the product, optics or process moves beyond the validated envelope.

This design also prevents a common interface error. Showing only a red or green result asks the inspector to obey an unexplained score. Showing the cropped region, the physical location and the applicable defect definition gives the inspector something to inspect. The metal-textile study uses anomaly heatmaps for precisely this guidance and reports that unfamiliar visual representations require careful introduction. [S34] Training must therefore cover the meaning and limits of the display, not just which button to press.

GFT NOTE

Commission the station in shadow mode, then assisted mode, before permitting automatic PASS. Advance only when the quality lead has reviewed defect escapes, exception burden and station failure drills on the actual product mix. Keep a timed manual fallback and record every bypass. The acceptance unit is the complete decision cell, including its people and escalation rules, not the model alone.

Human triage does not guarantee zero escapes. It creates a place to detect, contain and learn from them. Automatic PASS still needs a documented audit sample chosen independently of the model's score. If an audit finds an escape, the response is to hold affected work, review the operating envelope and test the threshold. The person who raises the alarm should be thanked for exposing a system problem, not measured as an obstacle to automation.

12.2

Active learning on the floor: every operator confirmation sharpens the local model

An exception decision becomes useful training material only when its context survives. A one-tap confirmation should store the crop, original frame reference, roll or bundle identifier, station, model version, anomaly score, defect class, severity and inspector disposition. Where the inspector cannot decide, the correct label is UNCERTAIN with escalation, not a forced guess. This discipline turns routine verification into a curated evidence set without pretending that every tap is infallible ground truth.

The first review belongs to the quality lead. Ambiguous crops, new defect forms and disagreements between inspectors should enter a consensus review before they influence a threshold. Duplicate views of the same physical flaw must be grouped so that repeated frames do not dominate the training set. Images from one roll or garment style should not appear in both training and held-out test partitions without that relationship being disclosed. These are basic defences against teaching the model the habits of a shift rather than the physics of the defect.

Active learning should remain offline. The station must not change its own weights after each confirmation. A technician prepares a candidate dataset, the quality lead checks label quality, and the model owner evaluates a candidate version against critical-defect recall, false alarms and subgroup performance. Release requires a named approver and a rollback path. This separation matters because a convenient touchscreen can otherwise become an uncontrolled model-training channel.

Label governance must also preserve the difference between NOT DEFECTIVE and NOT YET UNDERSTOOD. A harmless lint fleck can become a useful negative example only after the inspector establishes that it is removable contamination rather than a stain or foreign fibre. A novel texture should enter an investigation set until the quality lead decides whether it represents acceptable variation, a new defect class or an optics problem. Combining these cases into one nominal label would teach the model to ignore genuine process changes. The review queue therefore needs a small taxonomy of uncertainty as well as a taxonomy of defects.

The edge cascade study shows why better human labels matter operationally. On its disjoint benchmark, the recall-prioritised first stage forwarded 53.4 percent of frames, and 85 percent of those forwarded frames were false alarms. The authors estimate that tighter calibration could reduce inference work by 29 to 45 percent, while explicitly positioning the system as AI-assisted triage rather than autonomous acceptance. [S07] A high false-alarm rate is therefore not merely a computing issue. It determines how often a person is interrupted and how quickly alert fatigue can develop.

> **WORKED EXAMPLE: FEWER FALSE ALARMS MEAN FEWER HUMAN INTERRUPTIONS**
>
> Take an illustrative batch of 1,000 frames and apply the measured 53.4 percent forwarding rate from the edge study. The station sends 534 frames to review. If 85 percent are false alarms, approximately 454 are false alarms and 80 are other forwarded cases. Now assume, purely as a scenario, that calibration halves the false alarms while retaining every other forwarded case. Review falls to 80 + 227 = 307 frames. The reduction is (534 - 307) / 534 = 42.5 percent. This is GFT arithmetic from cited rates plus an explicit calibration assumption, not a measured factory result or a promise that recall can be preserved at a different threshold. [S07]

The worked example changes the management question. Asking how many frames the model classifies per second ignores the bottleneck experienced by the inspector. The better question is how many consequential interruptions the station creates, and whether each interruption adds information. Review queues should be measured by fabric type, station, model version and shift. A sudden increase may indicate lint, changed lighting or a new print rather than a genuine deterioration in production.

Inspectors should not be rewarded for agreeing with the model. Such an incentive teaches blind acceptance and contaminates the labels. Nor should the number of taps become a productivity target, because it rewards unnecessary forwarding and fast, careless review. Useful measures are audit agreement on independently reviewed samples, correct escalation of critical cases and the proportion of dispositions with complete context. The purpose of active learning is to make human judgement more informative, not to make it more frequent.

12.3

Ergonomics and occupational health: eliminating backlit perch-frame eye strain

The ergonomic case begins with the task, not with a claim that screens are inherently healthier than fabric. Manual inspection combines sustained attention, repeated eye movement and a body position fixed by the perch or table. A comprehensive fabric-inspection review identifies lack of concentration, human fatigue and time consumption as central drawbacks of manual detection. [S03] Replacing continuous web watching with a poorly designed alarm screen would merely exchange visual fatigue for cognitive overload.

The task should be redesigned around selective attention. An inspector reviews a stationary crop with a visible region of interest, compares it with the physical standard and records a disposition. The inspector no longer needs to search every centimetre

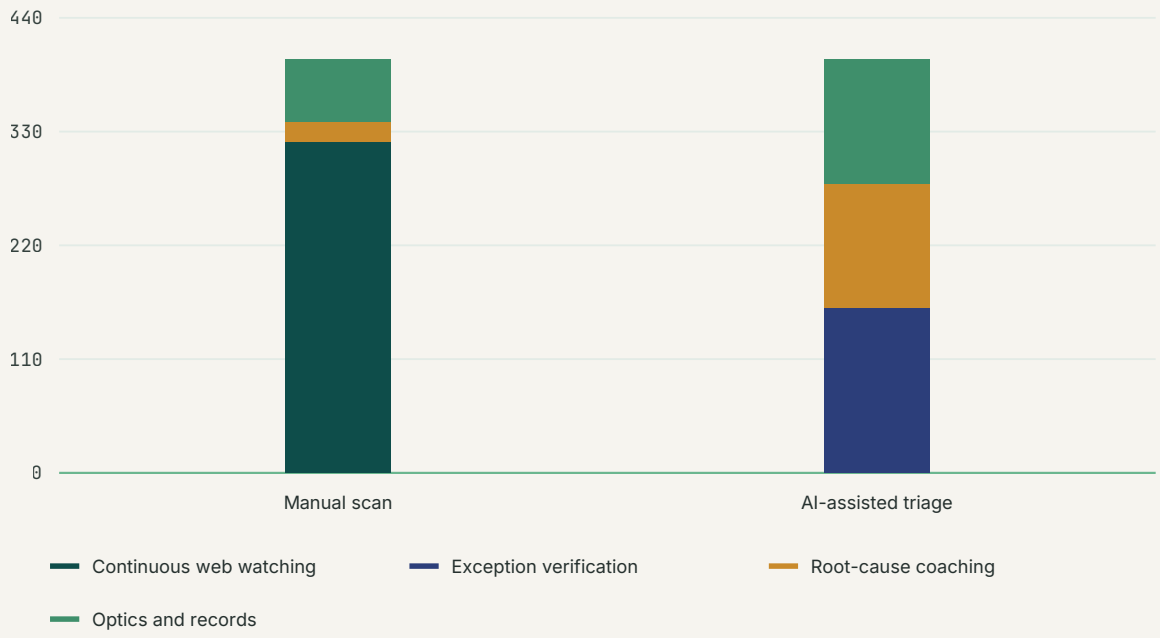
of moving cloth for a rare event. The metal-textile study reports an initial workload reduction above 50 percent, with a calculated 69.4 percent reduction when results were scaled, but its authors call for longer production trials and further investigation of worker relief. [S34] Those early results support testing the mechanism, not promising a health outcome or a fixed staffing reduction.

The shift envelope still needs explicit limits. The inherited Cameroon labour entry records a standard 40-hour week and a 48-hour maximum. [R16] Vision does not make sustained exception review harmless within those limits. Work design should include rotation between verification, physical confirmation and line coaching; recovery pauses away from the display; and a cap on repeated alarms that triggers supervisor intervention. The relevant control is the pattern of demands across the whole shift, including the late period when earlier fatigue becomes visible.

Exhibit 56, p. 183 shows an illustrative allocation of a 400-minute net work window. It is a design case, not a time-and-motion study.

EXHIBIT 56 / MINUTES PER 400-MINUTE NET WORK WINDOW

Illustrative triage moves inspector time from scanning to verification and coaching



ILLUSTRATIVE [S34] [S07] Illustrative task allocation, not measured GFT performance. Both cases total 400 minutes after breaks and handover. Manual scanning is assumed to fall from 320 to zero; exception review is set at 160 minutes, a 50 percent reduction of the former scanning load. The released 160 minutes are assigned to coaching and to optics and records. Actual allocation requires floor measurement.

In this illustrative allocation, the arithmetic is intentionally simple. The triage case reserves 160 minutes for exceptions, half of the former 320-minute scanning load, then assigns the same amount to coaching and to optics and records. This preserves total paid time and demonstrates why the proposal is a task redesign rather than an automatic headcount case. The actual queue may be larger or smaller. If exceptions consume the whole window, the correct response is to investigate forwarding and process conditions, not to accelerate clicking.

Physical layout determines whether the new task is sustainable. The monitor should sit at a comfortable viewing height and distance without forcing neck rotation away from the material. Controls should allow seated and standing work where the station permits. Ambient light should not wash out the anomaly display. The inspector needs direct access to the physical roll or garment, because some decisions depend on handling, transmitted light or comparison with a reference sample. Touch targets must remain usable with the gloves and hand positions required at that station.

Alarm design deserves the same engineering care as camera placement. Repeated audible alerts should be reserved for conditions that require immediate containment. Routine exceptions can enter a visible queue. Alerts should combine into one event when they refer to the same physical defect, and the interface should distinguish an image-quality failure from a suspected product defect. A supervisor needs to see queue age and station health, not just the number of red boxes.

Finally, the factory should test whether people are actually better off. Short confidential pulse checks, observation of posture, late-shift review quality and worker suggestions can reveal burdens that system logs miss. Retention and job satisfaction may be affected by many causes, so the deployment should not attribute changes to vision without a suitable evaluation. The defensible initial claim is narrower: triage can remove continuous searching and make workload visible, provided the factory acts on what it measures.

12.4

Bridging Cameroon's supervisor gap through visual quality coaching

A thin supervisor pipeline should be treated as a greenfield staffing assumption to be tested during recruitment, not as a quantified national fact. The inherited labour entry reports that about 90 percent of employment in Cameroon is informal, but that statistic does not count garment supervisors or establish their competence. [R16] It does show why a formal factory cannot assume that every applicant arrives with experience of structured quality routines, documented escalation or digital work instructions.

Vision can reduce the amount of abstract explanation a new supervisor must give. A skipped stitch, an open seam or a puckered seam can be shown as a cropped image beside an accepted example and the physical item. The anomaly map helps locate the region, while the defect dictionary supplies the class and severity. The conversation moves from a general accusation of poor workmanship to a specific question about thread tension, needle condition, fabric handling or method.

Good coaching follows a disciplined sequence. Confirm that the image and the physical item refer to the same unit. Ask the operator to describe what changed. Check the most plausible machine, material, method and environment causes. Agree one corrective action and its owner. Return to later evidence to see whether the action held. The camera supplies a shared object for discussion; it does not supply the root cause by itself.

This is a different supervisory job from watching a dashboard. The line quality lead must understand the process well enough to distinguish a recurring mechanism fault from an isolated handling error. The lead must also know when the image is misleading. A new print can trigger anomalies without any change in sewing. A tension problem can be real even when the model fails to flag it. Coaching therefore combines product knowledge, evidence interpretation and the authority to stop work.

Worker voice improves this learning loop. International Labour Organization case studies distinguish AI uses that complement skills from those that replace or deskill them, and identify worker oversight over how tools organise training and work as an important condition for greater discretion. [S1201] Supervisors should consequently review defect examples with operators before changing standards or incentives. An operator's explanation can reveal a material or method cause that no image contains.

The same principle protects dignity. A defect thumbnail is evidence about a product at a time and station, not a verdict on a person's effort or potential. Coaching should begin with process conditions and use individual attribution only where the evidence supports it. Repeat defects at one station call for examination of machine settings, work instructions, training coverage and incoming material before disciplinary action is considered.

A practical coaching board needs only a small set of views: the defect example, its class and severity, the agreed physical cause, the corrective action and the verification result. Keeping this view at line level prevents the quality lead from disappearing into central reporting. It also creates a teaching library in which new supervisors learn from resolved cases rather than relying solely on an experienced person's memory.

Supervisor assessment should reproduce that complete coaching cycle. An assessor selects a real or deliberately seeded case, asks the candidate to explain what the image can and cannot establish, and observes whether the operator is invited to contribute. The candidate must then choose a proportionate containment action and define the later evidence that would confirm recovery. A technically correct diagnosis is insufficient if the candidate humiliates the operator, hides uncertainty or closes the action without verification. This assessment makes respectful evidence handling a supervised skill rather than a statement of values.

12.5

Role redesign: from manual sorters to Quality Technicians and Vision Leads

The proposed career structure gives each new responsibility a name, an authorisation boundary and an assessed competence. It is a role architecture for recruitment and training, not a statement that these positions already exist at GFT or at a client factory. Progression depends on demonstrated judgement, not on willingness to accept automated decisions.

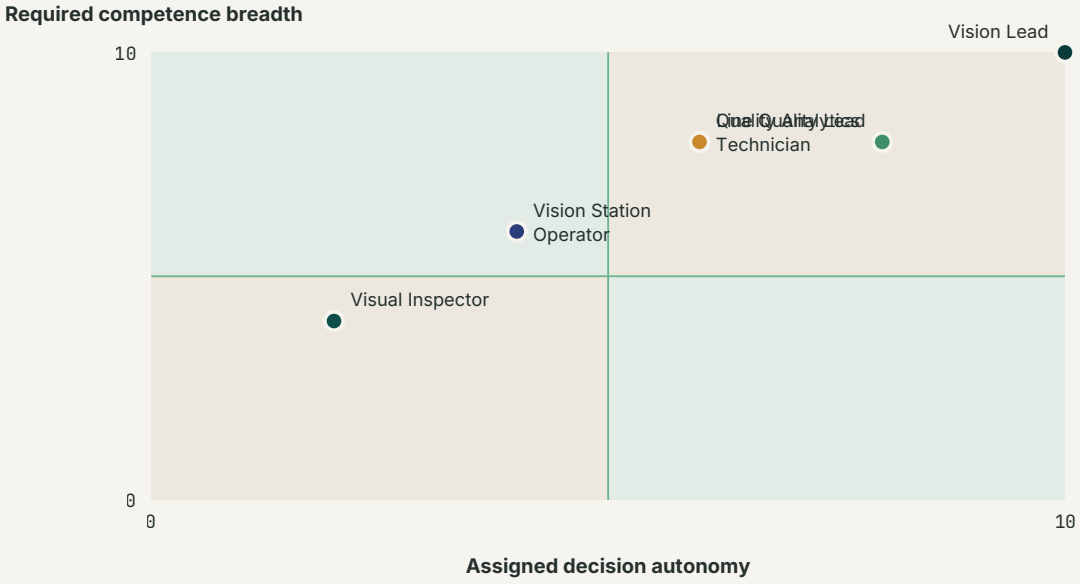
- A **Visual Inspector** applies the defect standard, confirms physical identity and escalates uncertainty. This role preserves craft knowledge and provides the first check on whether an image means what the interface suggests.

- A **Vision Station Operator** adds optics checks, queue control, safe fallback and complete disposition records. The operator can pause automatic PASS but cannot approve a model update.
- A **Quality Analytics Technician** curates datasets, reviews confusion matrices, compares performance across products and prepares drift investigations. This role connects floor evidence to the model owner without becoming a detached reporting function.
- A **Line Quality Lead** runs coaching, approves containment and verifies corrective action. The lead is accountable for release discipline and for ensuring that production pressure does not erase exceptions.
- A **Vision Lead** governs the operating envelope, approves threshold or model changes with the quality authority, coordinates supplier support and maintains the rollback decision. The role combines process authority with evidence governance.

Exhibit 57, p. 186 translates that architecture into an ordinal planning rubric. It does not represent observed skill scores, pay grades or a promised establishment.

EXHIBIT 57 / GFT ORDINAL ROLE-DESIGN SCORE

Greater decision authority requires broader demonstrated competence



DERIVED [S34] [S1201] [R16] Derived planning rubric, not measured attainment. Horizontal position = 2 x assigned autonomy level from 1 to 5. Vertical position = 2 x required competence-domain count from 1 to 5. Domains are product quality, station operation, data curation, coaching and model governance. Role assignments are GFT design assumptions informed by exception inspection, formal work organisation and worker-discretion principles; they require local validation.

The rubric makes an important restriction visible. A person may be an excellent garment inspector and still need preparation before governing a model. Promotion into analytics or vision leadership therefore requires evidence in the additional domains, not just seniority. Conversely, a technician who understands confusion matrices but cannot judge a seam should not hold final product-release authority. The two career paths can meet without collapsing into one generic digital role.

Role redesign should protect employment relationships during learning. The wage evidence argues against presenting the change as a simple labour-substitution exercise: the inherited entry records Cameroon's minimum wage at XAF 43,969 a month from 23 February 2024. [R14] The wider paper locates the economic case in avoided waste and rework. The workforce case is that experienced inspectors carry defect knowledge which the system needs, while new technical responsibilities create a route for that knowledge to be developed rather than discarded.

Workers and their representatives should help define the transition. Consultation can test whether the interface removes useful discretion, whether training is accessible and whether exception data could become covert individual surveillance. The International Labour Organization research links better outcomes to worker voice, skill investment and transparent routes to challenge technology-enabled decisions. [S1201] A formal appeal route is therefore part of the operating design: an inspector can request review of a disposition, and an operator can challenge an inference made about their work.

Recruitment should assess practical reasoning before software familiarity. Can the applicant compare a physical defect with a standard, explain uncertainty and stop a process when evidence conflicts? Those behaviours are harder to improvise after launch than basic touchscreen operation. Existing inspectors should receive the first opportunity to enter the new roles, with bridging instruction rather than an assumption that digital recruits are inherently superior.

12.6

Curriculum for a four-week vision operations bootcamp

The bootcamp should be a supervised floor programme, not a slide course followed by a login. An illustrative full-time format uses the standard 40-hour week recorded for Cameroon, giving 4 × 40 = 160 learning hours across four weeks. [R16] This duration is a planning assumption, not an externally validated certification standard. Shift cover, prior attainment and the availability of a competent instructor may justify a longer pathway.

Exhibit 58, p. 188 defines the proposed modules and observable competence gates.

EXHIBIT 58 / ILLUSTRATIVE FULL-TIME LEARNING PROGRAMME

A four-week bootcamp certifies observed decisions rather than screen familiarity

Module	Practical floor exercise	Observed competence gate
Week 1: Product quality and triage	Grade accepted and defective samples; route critical, uncertain and nominal cases	Correct containment and escalation on the assessor's complete case set
Week 2: Optics and station care	Clean optics; check reference tile, illumination, encoder and fallback	Checklist completed; failed reference check correctly stops automatic PASS
Week 3: Evidence and active learning	Disposition crops; resolve disagreements; read a confusion matrix	Every sampled crop carries valid class, severity, context and uncertainty
Week 4: Coaching and control	Lead a defect review; test a cause; run rollback and appeal drills	Quality lead accepts observed coaching, containment and review decisions

ILLUSTRATIVE [R16] [R29] [S34] [S1201] Illustrative GFT curriculum based on a 40-hour standard week and cited inspection and worker-voice principles. It is not an existing GFT course, an accredited qualification or evidence of staff attainment. A local quality authority must approve content, language, assessment and any use of the term certification.

The first module establishes the product standard before the model. Trainees learn the paper's defect dictionary, severity rules and the difference between a defect, an anomaly and an uncertain image. The inherited evidence describes computer vision across more than 40 defect types, but that breadth does not remove the need to define local acceptance criteria for each product. [R29] Assessment uses physical samples and deliberately includes borderline cases.

The second module makes station care a quality responsibility. Trainees practise lint removal, reference-tile checks, illumination inspection, encoder verification and controlled fallback. They learn to distinguish a dirty viewport from a process defect. They also learn what they may not do: cleaning must not become informal recalibration, and a supervisor must not silence a failed check to keep production moving.

The third module develops evidence judgement. Trainees review false alarms and escapes, complete context fields, separate duplicate defects and interpret a confusion matrix by class. Group review exposes disagreement and teaches when to seek consensus. The assessment rewards correct escalation of uncertainty, not speed or agreement with the model.

The fourth module moves the trainee into supervision. Each candidate leads a case review, asks for operator evidence, proposes a cause, agrees a corrective action and checks a later result. Rollback, manual fallback and appeal drills confirm that authority is real. A quality lead observes performance and records remaining development needs. Passing the course means demonstrated competence on the defined cases, not permanent mastery of every fabric and model.

Training continues after the bootcamp through short case reviews, cross-training and refreshers triggered by new products or material model changes. Each station needs more than one person able to perform safe fallback. Supervisors also need separate practice in interpreting queue and drift measures, because an operator's competence does not automatically prepare a lead to govern the whole decision cell.

12.7

Aligning people KPIs in the Performance-Management System

People measures should show whether the decision cell is becoming more capable, not whether workers are becoming more compliant. The Performance-Management System should connect station evidence to skill-matrix coverage, current training authorisation, review quality, coaching closure and safe escalation. Each measure needs a named denominator and a clear owner. A percentage without those definitions can make a weak training system appear strong.

Skill-matrix coverage counts people whose observed competence is current for a defined role and product family. It should distinguish course attendance from successful assessment and should show backup coverage for fallback. **Recheck agreement** compares a sample of dispositions with an independent quality review, including uncertainty and escalation. **Coaching closure** tracks whether agreed actions have an owner and a later verification, rather than counting meetings held.

Exception burden belongs beside those capability measures. Queue age, repeated alarms for one physical flaw and after-hours spillover reveal whether the system is creating hidden work. **Appeal outcomes** show whether challenges change a disposition or expose a defect in the standard. **Bypass and fallback** records show whether people can operate safely when the model or optics fail. These measures make resistance diagnostically useful instead of treating it as a behavioural problem.

A small measurement dictionary should resolve recurring disputes before the dashboard is built. Recheck agreement needs a sampling rule that includes accepted and rejected items, not only cases selected by the model. Coaching closure needs a due condition and evidence of effect, not a checkbox marked after discussion. Training currency needs an expiry rule linked to product or model change. Exception burden needs to separate product alarms from station-health alarms. Without these definitions, different shifts can improve their reported position simply by counting differently.

The quality authority should own recheck sampling, the people lead should own competence records, and the production lead should own removal of queue and workload constraints. Shared review matters because no single function can interpret the measures alone. A rise in disagreement may reflect weak training, a changed product standard or model drift. A fall in exceptions may reflect better screening, a failed camera or inspectors bypassing the queue. The KPI review must test those competing explanations against floor evidence.

The KPI design must also restrict misuse. Raw defect counts at an operator's station are not a fair individual score because machine condition, material, style difficulty and sampling influence them. Model agreement is not a fair inspector score because the model can be wrong. Individual images and event logs should not be repurposed for discipline without a defined rule, access control and an opportunity to contest. This follows the International Labour Organization's warning that algorithmic management can intensify monitoring and reduce worker control when oversight is weak. [S1201]

> **GFT NOTE**
>
> Put a short workforce charter beside every vision station and inside the Performance-Management System. State what the camera measures, who can see identifiable records, which decisions remain human, how to appeal and which measures will never determine discipline on their own. Review the charter with inspectors and operators before automatic PASS begins, then test it in the first drills and quality audits.

The management cadence should close with capability action. The daily huddle resolves exceptions and coaching needs. The weekly quality review examines recheck disagreement, queue burden and repeat causes. The monthly people review updates training authorisation, backup coverage and role progression. Where evidence shows overload or workarounds, management changes the system rather than asking inspectors to compensate indefinitely.

This is how vision elevates supervision. Inspectors become authors of verified evidence. Technicians become stewards of data quality. Supervisors become coaches with time and visual proof to improve the process. The model remains subordinate to a social and technical system that can question it, maintain it and stop it.

Give the machine the routine search, and give people the judgement, authority and time to improve the work.

CHAPTER SOURCE NOTES

S34 IntechOpen. 2025-01-03.

S1201 International Labour Organization. 2025-07.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S07 arXiv. 2026-08-12.

S03 Mathematical Problems in Engineering. 2020-11-16.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

13

Every defect frame logged at the edge builds an auditable quality and waste ledger for Digital Product Passports

When vision inspection is tied to cotton lot numbers, quality control generates the process and pre-consumer waste evidence that can later support buyer and Digital Product Passport requirements.

CHAPTER 13
DECISION TEST

When vision inspection is tied to cotton lot numbers, quality control generates the process and pre-consumer waste evidence that can later support buyer and Digital Product Passport requirements.

A defect frame becomes commercially useful beyond the inspection station only when it joins a conserved chain of identity, mass and decision. The factory should therefore treat every edge detection as an audit event linked to a fabric roll, a cut bundle and, where the chosen passport granularity requires it, an individual garment. That event can support a quality certificate, a pre-consumer waste balance and later Digital Product Passport (DPP) exports. It cannot establish fibre composition, chemical content, life-cycle impact or durability on its own. The correct architecture combines camera evidence with supplier declarations, physical weighing, laboratory tests and controlled corrections. This is an engineering and data-governance method, not legal advice or a claim that a passport system is already operating at a GFT facility. [S27] [S28] [S1301]

13.1

Shop-floor quality becomes traceable only when product genealogy survives every transformation

The inherited SODECOTON profile records more than 220,000 supported cotton producers and nine ginning plants in northern Cameroon. That establishes the relevance of lot-level provenance, not the completeness of any bale-to-garment record and not the quality of a particular fabric delivery. [R03] The factory must obtain the actual identifiers and documents from its suppliers. A bale or ginning-lot reference should connect to spinning and fabric-production batches, then to the received roll. Where upstream partners cannot provide that continuity, the record should state the earliest verified boundary rather than implying an unbroken chain back to cotton production.

A useful genealogy begins at goods receipt. The roll identifier carries supplier, received weight, declared composition, fabric construction, dye or print lot, document references and inspection status. The line-scan station adds its own event identifier, acquisition time, station, calibrated configuration, model version and defect coordinates. Deep-GD research shows that fabric and garment defects can be classified into categories including stains, double ends, snarls, dye spots and tiny holes, but such a classifier does not determine which upstream lot caused them. [S18] The roll genealogy makes a correlation investigable; it does not turn correlation into proof of causation.

The next transformations are where many traceability designs fail. One roll can supply several spreading lays. One lay can produce several bundles. A bundle can later split across sewing lines, and repaired panels can return to an earlier operation. Each event therefore needs parent and child identifiers. A roll-to-lay record preserves consumed length or mass. A lay-to-bundle record preserves the marker and ply relationship. A bundle-to-garment record preserves style, size and disposition. Without those links, an accurate defect coordinate on the incoming roll cannot reliably locate a garment or support a product-level claim.

Exhibit 59, p. 193 separates the physical, information and decision responsibilities. The architecture is deliberately more concerned with identity conservation than with a particular database brand.

EXHIBIT 59 /

A conserved garment genealogy carries provenance into quality and passport evidence



DERIVED [R03] [S28] [S1301] [S1302] Derived GFT architecture from inherited cotton provenance context, the cited textile tracking prototype and official DPP framework. It is a proposed design, not an installed GFT system. Final product-specific DPP fields and granularity must come from the applicable delegated act.

The physical identifier should remain readable at the point of work. A roll label can survive goods receipt and spreading; a bundle traveller can survive cutting and sewing; a garment carrier can be attached when item-level tracking is justified. Scanning should confirm the actual object rather than invite manual transcription. Where a label is damaged, the operator needs a controlled reissue process that references the old identifier and requires a second check. A convenient new label without that link silently breaks the audit trail.

The information model should distinguish three kinds of statement. An **observation** records what a sensor, scale or person measured. A **declaration** records what a supplier, laboratory or economic operator asserted. An **inference** records what a model calculated from those inputs. The DPP export must preserve that distinction. A vision class is an inference until a qualified disposition confirms it. A declared cotton origin remains a supplier declaration until the applicable chain-of-custody evidence verifies it. A waste figure derived from roll mass and packed mass remains arithmetic whose quality depends on the completeness of both measurements.

GFT NOTE

Start the traceability pilot with one fabric family and one garment style. Require a complete roll-to-lay-to-bundle genealogy before adding a customer-facing passport. GFT recommends a weekly exception review for unidentified rolls, orphaned bundles, duplicate scans and unreconciled mass. Do not print a QR code until the underlying record can answer which physical item it identifies, which evidence supports each field and who may correct it.

This narrow pilot also establishes the right granularity. Regulation (EU) 2024/1781 explains that a passport may relate to a model, batch or item, depending on product-specific rules and impact assessment. [S1301] Item-level capture can be valuable for critical uniforms or high-value garments, but it can be unnecessary for every basic style. The data model should permit all three levels without assuming that every defect thumbnail must be publicly attached to every garment. Batch-level disclosure can preserve useful evidence while limiting storage, cost and disclosure of sensitive production detail.

13.2

ESPR readiness means portable evidence, not a speculative passport template

The Ecodesign for Sustainable Products Regulation (ESPR), Regulation (EU) 2024/1781, establishes a framework for product-specific performance and information requirements. Its recitals state that information requirements should include a DPP and describe links to unique product, operator and facility identifiers. [S1301] The operative duties for garments depend on the relevant delegated act. A greenfield factory should consequently build portable evidence and identifier discipline now, while avoiding a claim that a self-designed passport already satisfies a future textile requirement.

The Commission's working plan for 2025 to 2030 places textiles and apparel among the new priority products and gives 2027 as the indicative adoption year for the relevant measure. That is a planning horizon, not a statement that all garments must carry a compliant passport from the beginning of that year. The same communication states that the information collected and made available in a DPP will be specified in product-specific delegated acts. [S1302] The compliance review must therefore follow the final act, its application date, its scope and its transition provisions rather than a vendor's generic checklist.

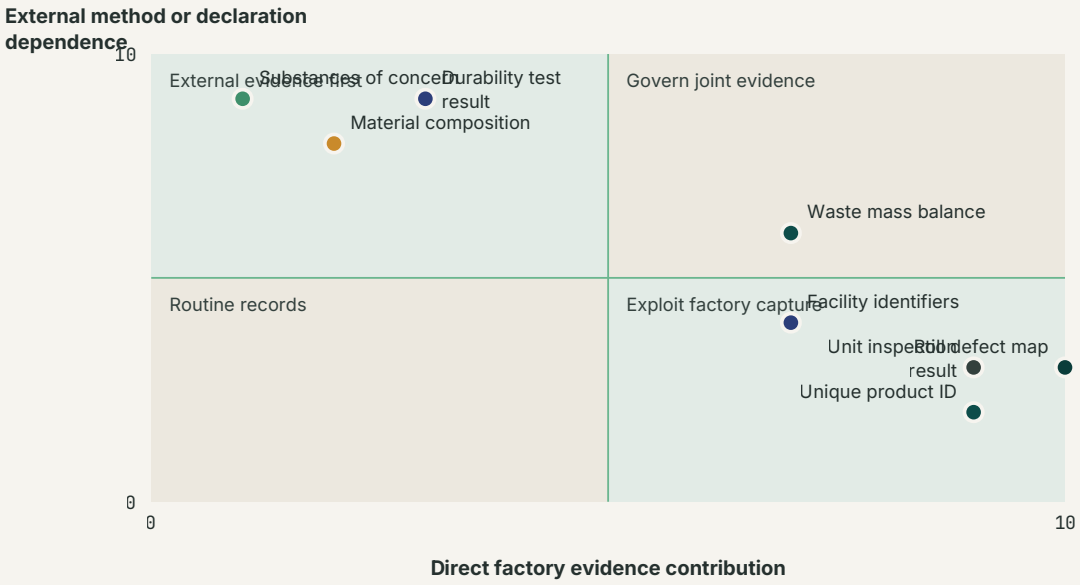
This distinction changes what the factory should buy. A rigid passport product with a fixed set of attractive fields may become obsolete when product rules arrive. A well-governed event store can map verified fields into a later schema without re-capturing production history. Procurement should test export in documented, machine-readable formats, stable identifiers, role-based access and the ability to retain evidence when a supplier changes. The cited tracking research demonstrates a permissioned Hyperledger Fabric prototype for textile attributes and traceability. It supports technical feasibility, not a regulatory requirement to use blockchain or that vendor's architecture. [S28]

Research with garment manufacturers supplying the European market identifies incomplete supplier information, verification difficulties, limited digital systems and insufficient environmental-data capacity as important readiness barriers. [S27] Automated inspection addresses only part of that problem. It can supply timestamped process and quality evidence at the factory boundary. It cannot fill an absent fibre declaration, validate a chemical inventory or perform a life-cycle assessment. Readiness work must therefore assign every required attribute an evidence owner and a method, rather than treating digitisation as a single software purchase.

Exhibit 60, p. 195 makes that division visible. Horizontal position represents how much of the evidence the proposed edge and Performance-Management System (PMS) architecture can generate directly. Vertical position represents dependence on an external method, declaration or test. These are ordinal design scores, not regulatory weights or measured readiness scores.

EXHIBIT 60 / DERIVED ORDINAL DESIGN SCORE FROM 0 TO 10

Vision covers process evidence while composition and durability still need external proof



DERIVED [S27] [S28] [S1301] [S1302] Derived mapping for planning, not a legal classification or measured score. X reflects the proposed share generated directly by factory edge and PMS records. Y reflects reliance on supplier evidence, laboratory methods or product-specific rules. Positions are GFT synthesis from the cited regulatory framework and garment data-readiness research.

The matrix prevents two opposite errors. The first is passport theatre: presenting a polished public record whose composition and environmental fields have no stronger evidence than an unverified spreadsheet. The second is scope paralysis: postponing all data capture because final textile rules are not yet available. Product identifiers, roll genealogy, inspection outcomes and mass balances have immediate operational value and are natural candidates for early capture. Composition, substances and durability require parallel supplier and laboratory work. The programme can advance both without confusing their evidence status.

Access design belongs in the same preparation. The ESPR framework provides for differentiated access to protect confidential business information while allowing relevant actors to perform their roles. [S1301] The factory should separate a public or buyer view from an internal evidence view. A buyer may need a signed inspection summary and its method. An auditor may need sampled event records and calibration evidence. A regulator may require fields defined by the final act. None of these audiences automatically needs every raw frame, operator identifier or supplier-commercial term.

Finally, the factory should version its regulatory mapping. Each candidate export field needs a status such as operationally observed, supplier declared, laboratory tested, derived, pending method or not applicable. When a delegated act or buyer specification changes, the compliance owner updates the mapping rather than rewriting history. This creates an honest readiness register: it shows what is available, what remains assumed and which evidence would close each gap.

13.3

Pre-consumer waste accounting must reconcile kilogrammes before claiming savings

A waste ledger starts with a physical boundary, not with defect counts. For a cut-make-trim plant, the practical opening quantity is fabric received or issued to spreading, expressed in kilogrammes. Closing quantities include packed garment fabric, reusable remnants, recyclable offcuts, downgraded material, rejected panels, irreparable garments and material awaiting disposition. The balance should also explain moisture or weighing differences where these are material. Counts of defective frames cannot substitute for this reconciliation because a tiny hole and a rejected full panel can carry the same event count but very different masses.

The reporting categories need stable definitions. Cutting offcuts are not automatically defective material. A mapped fabric fault removed before cutting is quality-driven waste, while an unavoidable marker edge is process offcut. A rejected garment that can be repaired is work in progress, not final waste. A second-quality garment remains a product but may have a different value and destination. UNEP's textile roadmap treats resource efficiency and circularity across the whole value chain, which reinforces the need to retain destinations rather than reporting one undifferentiated scrap figure. [S30] The factory should not claim waste prevention merely because material leaves the factory under a different label.

Vision contributes two specific measurements. At fabric entry, it supplies defect class, severity and coordinates against the roll. At garment inspection, it supplies a detected condition and the subsequent repair, downgrade or rejection disposition. The inherited inspection evidence places manual detection at 60 to 70 percent accuracy and automated systems above 90 percent in their cited contexts, but those figures do not quantify a factory's waste reduction. [R29] The saving must be measured through changed material disposition: defects avoided before cutting, panels recut rather than garments rejected, and sewing faults repaired before further value is added.

Defect surveys also show why count and mass must remain separate. In the Indonesian garment case, 107 defective units among 1,380 produced units comprised faded prints, misaligned prints and inappropriate colours. [S26] Those categories can support process investigation, but they do not reveal the fabric mass in each rejected item, whether repair was possible or what happened to the material. A sustainability ledger needs the later physical disposition before it can convert a quality event into waste accounting.

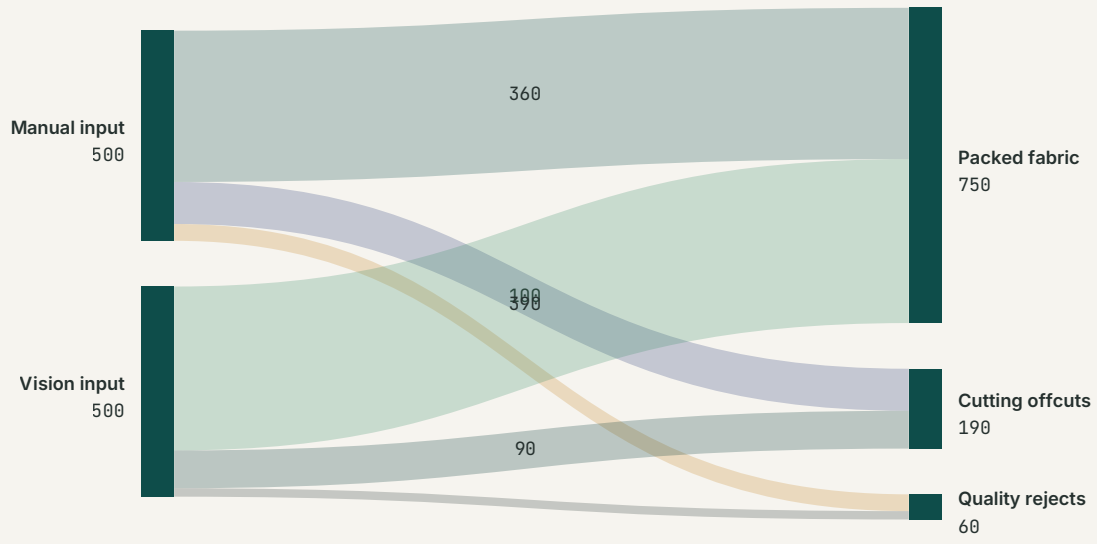
WORKED EXAMPLE: ONE INPUT, TWO EXPLICITLY ILLUSTRATIVE DISPOSITIONS

Take an illustrative 1,000 kg fabric input split into two hypothetical 500 kg scenarios solely to demonstrate the accounting method. In the manual scenario, assume 72 percent becomes packed garment fabric, 20 percent becomes cutting offcuts and 8 percent becomes quality-rejected material. The quantities are 360 kg, 100 kg and 40 kg. In the vision-assisted scenario, assume better flaw mapping and earlier repair change those shares to 78 percent, 18 percent and 4 percent. The quantities are 390 kg, 90 kg and 20 kg. Each scenario balances to 500 kg. Packed output differs by 30 kg and total offcut plus rejected material differs from 140 kg to 110 kg. These are assumed shares, not measured GFT performance and not a forecast of vision savings. The relevant empirical anchor is only that automated inspection can improve defect capture; local scales must establish the actual effect. [R29] [S32]

Exhibit 61, p. 197 draws exactly that worked example. The split input is a comparison device, not a claim that one physical batch was divided between two operating systems.

EXHIBIT 61 / KG OF FABRIC PER ILLUSTRATIVE SCENARIO

Two illustrative fabric inputs reconcile into packed, offcut and rejected mass



ILLUSTRATIVE [R29] [S30] [S32] Illustrative GFT arithmetic, not observed factory data. A 1,000 kg comparison input is split into two independent 500 kg scenarios. Manual: 500 × 72% = 360 packed, 500 × 20% = 100 offcuts, 500 × 8% = 40 quality rejects. Vision-assisted: 500 × 78% = 390 packed, 500 × 18% = 90 offcuts, 500 × 4% = 20 quality rejects. All shares are assumptions. Cited sources support defect detection and resource-efficiency context, not these quantities.

The example also exposes what a naive saving calculation misses. If packed fabric rises, the factory must check whether production mix, garment size or marker efficiency changed at the same time. If rejected mass falls, it must check whether defects escaped into packed goods or were reclassified as seconds. If offcuts fall, it must confirm that usable remnants were not simply moved into inventory. A fair comparison holds product and size mix sufficiently constant, uses calibrated scales and includes downstream audits. Otherwise the chart can reward weaker inspection rather than better material use.

A useful denominator is kilogrammes of final pre-consumer waste per hundred accepted garments, accompanied by total input mass and accepted units. The numerator should exclude material still awaiting a decision and should disclose reusable and recyclable destinations separately. A second measure, quality-rejected mass divided by fabric issued, isolates the fraction most directly connected to defect management. Neither measure should be compared across styles without explaining garment mass and construction. Denominator discipline matters more than an apparently precise percentage.

The monthly reconciliation should close in both directions. Roll issue records should account for all fabric entering spreading. Packed, stock, rework, remnant and waste records should account for where it went. The difference requires investigation, not an automatic adjustment to make the balance look complete. Waste contractors and internal bins can provide destination evidence, while stock counts test whether material has merely accumulated. Only after that closure should the PMS attribute a change to a vision intervention.

For passport readiness, the mass balance should remain an internal evidence product until applicable rules define what must be disclosed. The Commission states that product-specific acts will specify DPP information. [S1302] A factory can nevertheless retain period, boundary, method, destination and verifier for each waste figure now. That discipline makes a later disclosure auditable and gives domestic operations an immediate resource-efficiency measure, whether or not a particular field becomes mandatory.

13.4

A buyer quality certificate should expose its evidence and its limits

Domestic and regional buyers can use a quality certificate long before an EU passport applies. The useful object is not a decorative badge. It is a signed, batch-specific summary linking the buyer's acceptance standard to inspection coverage, observed defects, dispositions and unresolved exceptions. The certificate should identify the roll or production batch, garment style, inspection stations, method versions and the person or function authorising release. It should also state whether inspection covered every item or a defined sample.

The certificate's strength comes from traceable exceptions. A PASS summary should remain connected to the events that support it, including holds, repairs and later releases. Deleting failed intermediate states would make the final result look cleaner while destroying the evidence needed to investigate a dispute. A corrected record should reference the earlier event and explain the correction. This is more useful than calling the database immutable: authoritatively preserving corrections is the control, not preventing any new information from being added.

The textile tracking prototype cited here combines a permissioned blockchain with product tracking and reports an end-to-end organic-cotton garment experiment. [S28] That demonstrates one way to share signed records among parties. It does not establish that blockchain is necessary for a single factory certificate. A smaller plant

can begin with conventional relational records, cryptographic hashes of exported evidence and controlled signatures. The choice should follow the number of independent parties, their trust boundaries and their ability to operate the system.

Human verification remains part of the certificate method. The metal-textile inspection workflow routes doubtful cases to an inspector rather than asking the model to decide every case. [S34] The certificate should therefore distinguish automatic screening, inspector-confirmed disposition and later audit. If a buyer requires item-level assurance, unresolved critical exceptions block release. If the certificate covers a batch, it names the sampling rule and the lots excluded. The document must not imply that high model confidence equals product conformity.

A concise certificate can carry five evidence blocks: identity, acceptance standard, inspection coverage, exception closure and authorisation. Supporting thumbnails and raw events remain available under agreed access rather than crowding the buyer-facing page. Where a dimensional check, shade assessment or laboratory test falls outside the vision station's capability, the certificate marks it as separately evidenced or not performed. Empty fields are preferable to inferred compliance.

This approach has commercial value without inventing a market result. It gives a procurement team a reproducible way to compare deliveries and investigate failures. It also lets the factory answer a complaint by retrieving the exact production and inspection context. The certificate does not guarantee that no defect escaped, and it should say so. Its promise is narrower and stronger: the stated checks were performed on the identified scope, the exceptions were dispositioned, and the evidence can be inspected.

13.5

Lightweight MLOps keeps the ledger meaningful as optics and fabrics change

A traceable prediction is not merely an image plus a class. It also needs the model weights, threshold, preprocessing configuration, camera setup and illumination reference that produced it. Without those fields, the factory cannot reproduce why two stations treated similar fabric differently or why the same station changed behaviour after maintenance. Model governance is therefore part of product evidence, not a separate technical luxury.

Drift can begin outside the neural network. LED output can change, lint can obscure the viewport, roller vibration can alter focus, and a new dye lot can shift the background. A daily reference-image check should compare intensity and sharpness against an approved envelope before automatic PASS begins. A failed check should route work to human inspection and record an optics exception. Cleaning or replacing a light then requires a repeat reference capture, not an informal assumption that the station has recovered.

Model outputs need a separate audit. A weekly sample should include automatically cleared work as well as flagged frames, because reviewing only anomalies cannot estimate defect escapes. The edge-cascade study reports that 85 percent of forwarded frames were false alarms in its recall-prioritised benchmark and examines tighter

calibration as a way to reduce inference work. [S07] That result is specific to its data-set. It nevertheless shows why the factory should track false alarms, escapes and review burden together rather than optimising one detection metric.

Exhibit 62, p. 200 proposes a deliberately small operating rhythm. The sample sizes and intervals are illustrative starting points to be approved against production risk and volume.

EXHIBIT 62 / ILLUSTRATIVE FACTORY GOVERNANCE CADENCE

A light MLOps cadence preserves evidence before and after every model change

Cadence	Control and retained evidence	Hold or rollback condition
Each shift and daily	Reference image, lint and focus check; store configuration and result	Failed optical envelope stops automatic PASS
Weekly	Review 50 illustrative stratified frames, including auto-pass and flags	Critical escape or unresolved review backlog
Monthly	Curate labels; compare candidate and incumbent on held-out lots	Critical-class regression or broken subgroup performance
Quarterly and new fabric	Cross-fabric validation, calibration witness and export test	Unvalidated domain or irreproducible export
Every change	Version weights, threshold, data set, approver and rollback package	No release without signed evidence and restore path

ILLUSTRATIVE [S07] [S27] [S34] Illustrative GFT control plan informed by cited exception triage, calibration and data-readiness research. It is not an implemented GFT routine or a prescribed legal cadence. Local quality and production owners must set sample sizes and acceptance limits from risk, volume and defect severity.

The weekly sample in the table is not a statistical guarantee. Fifty illustrative frames may be enough to trigger an investigation but insufficient to estimate a low escape rate with useful precision. The quality lead should increase scrutiny after a product change, a critical complaint or a model update, and document the sampling rationale. Stratification should retain fabric family, colour, station, shift and defect severity so that an aggregate review does not hide a weak subgroup.

Candidate models should run beside the incumbent before release. The comparison uses held-out rolls or production periods and checks class-level recall, false alarms, review queue and processing time. Nearby frames from one physical defect must not leak between training and test partitions. The active-learning workflow can supply labels, but uncertain cases need consensus review rather than forced acceptance. Human exception review is valuable precisely because it preserves a route for disagreement with the model. [S34]

Release creates a new evidence version; it does not rewrite the old one. Events produced under the incumbent retain that version forever. Events produced under the candidate carry the new identifier only after approval. If performance deteriorates,

rollback restores the complete approved package, including threshold and preprocessing. A later auditor can then understand which decisions were made under which configuration and why the factory changed course.

The same discipline applies to exports. A monthly test should confirm that event identifiers, defect classes, timestamps and dispositions survive the journey from edge buffer to central store without duplication or silent type conversion. A quarterly recovery drill can restore a sampled batch and verify its hashes, references and images. These tests protect the quality ledger from a subtle failure: the model may still be accurate while the audit trail becomes incomplete.

13.6

Schema continuity turns spreadsheets, edge buffers and central records into one ledger

The data spine should begin with a canonical event contract rather than a database decision. Each event needs a durable identifier, event type, product or lot reference, parent references, station, timestamp and time zone, evidence status, method version and disposition. Measured values carry units. Missing values carry a reason. The Excel Quality Log, edge buffer and central relational store then become different implementations of the same meanings. A factory can change technology without changing what a defect, repair, rejection or accepted unit means.

Schema control requires an explicit migration path. A new optional field can be added with a version note. A changed definition cannot quietly reuse an old field name. For example, if waste begins to include reusable remnants after previously excluding them, the report must show the definition break and restate the affected comparison. Controlled vocabularies for defect class, severity, disposition and waste destination prevent stations from inventing local labels that later cannot be reconciled.

Offline operation creates additional audit risks. An edge station may capture events while disconnected, then synchronise after another station has already written related records. Durable event identifiers and idempotent import rules prevent duplication. Original capture time must survive alongside receipt time. A late record should not overwrite a later human disposition without a visible conflict review. These controls let production continue without making the central ledger falsely appear complete at every moment.

Corrections should be append-only in effect. The original entry remains available, the correction references it, and the authoriser and reason are recorded. A role-based permission model restricts who can alter master identifiers, approve dispositions or export evidence. The textile tracking research addresses integrity across multiple supply-chain parties, but a small factory still needs these basic controls even without a distributed ledger. [S28] Technology cannot compensate for shared accounts, unlogged spreadsheet edits or uncontrolled label reissue.

Audit tests should be routine and short. Can a selected garment be traced to its bundle, lay and rolls? Do roll issue and disposition masses reconcile? Does every automatic PASS carry a valid model and station configuration? Are orphan identifiers, missing images and unresolved critical events reported? Can an exported batch be

restored and read independently? Each failed test receives an owner and containment action. A passport export drawn from a ledger that fails these checks would only make weak evidence easier to distribute.

Retention should follow evidence purpose and applicable obligations rather than an arbitrary promise to keep everything forever. Raw frames consume storage and may contain sensitive production detail. Event records, dispositions and calibration results may need a different retention period from thumbnails. The factory should define legal, buyer and operational holds, protect backups and test deletion as well as retrieval. Differentiated access, contemplated by the ESPR framework, reinforces the case for separating internal evidence from external disclosure. [S1301]

13.7

Measured quality supports circular claims only when its evidentiary boundary stays visible

The anti-greenwashing rule is simple: do not let a measurement say more than its method permits. A seam-puckering study shows that computer vision can make grading more consistent than a subjective visual method. [S21] That supports objective inspection evidence. It does not prove how the garment will perform after repeated washing, wearing or repair. Durability claims need an appropriate test protocol and sample, while life-cycle claims need their own boundaries and data.

The same caution applies to provenance and waste. A supplier lot reference is not by itself proof of farm origin. A lower waste balance is not evidence of lower total environmental impact if production, energy or product lifetime moved adversely. UNEP's roadmap presents circularity as a value-chain programme rather than a single factory metric. [S30] The factory should therefore publish each claim beside its denominator, period, evidence owner and unresolved exclusions. Vision becomes one credible strand in a wider evidence system, not a substitute for it.

A practical release review can ask five questions. Is the identified product the same one described by the claim? Was the method valid when the evidence was captured? Are exceptions and corrections visible? Can another reviewer reproduce the calculation? Does the wording distinguish observation, declaration, inference and estimate? If any answer is no, the claim should be narrowed or held. This discipline protects buyers from overstatement and protects the factory from building a polished passport on unstable evidence.

The resulting ledger is valuable even before external passport obligations settle. It connects defect prevention to physical yield, quality release to retrievable evidence and supplier dialogue to identified lots. It also shows where cameras stop and other

methods must begin. That is the correct preparation for Digital Product Passports: not a promise of automatic compliance, but an inspectable record capable of supporting the fields that final product rules and buyers actually require. [S27] [S1302]

A passport is only as credible as the genealogy, mass balance and evidence behind each field.

CHAPTER SOURCE NOTES

S27 Sustainability (MDPI). 2026-07-13.

S28 Textile & Leather Review. 2026-04-28.

S1301 European Parliament and Council of the European Union. 2024-06-13.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

S18 International Journal of Electrical and Electronics Research. 2024-02-05.

S1302 European Commission. 2025-04-16.

S30 United Nations Environment Programme (UNEP). 2023-05.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S26 Jurnal Teknik Industri (JATRI). 2026-09-25.

S32 Security and Communication Networks. 2021-05-10.

S34 IntechOpen. 2025-01-03.

S07 arXiv. 2026-08-12.

S21 Textile Research Journal. 2016-10-14.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

14

A four-phase crawl-walk-run roadmap deploys computer vision in step with factory ramp and proves payback inside the PMS

Synchronising vision deployment with the eighteen-month Performance-Management System rollout takes a greenfield line from manual logs to closed-loop quality without operational shock.

CHAPTER 14
DECISION TEST

Synchronising vision deployment with the eighteen-month Performance-Management System rollout takes a greenfield line from manual logs to closed-loop quality without operational shock.

A greenfield garment factory should release computer vision through four condition-gated phases, not through a single automation purchase. The illustrative roadmap runs for 24 months. Phase 0, Months 1 to 3, establishes the manual Quality Log, representative test sets and vendor acceptance. Phase 1, Months 4 to 9, stabilises incoming-fabric and finished-garment inspection with human review. Phase 2, Months 10 to 15, connects verified roll maps to cutting and tunes models on local evidence. Phase 3, Months 16 to 24, adds critical-seam inspection and traceability links after the earlier stations have proved their operating value. [R29] [S04] [S07] [S19]

The answer is therefore crawl, walk, run, with a gate at every change in scope. A calendar date may trigger a review, but it cannot authorise the next expenditure. The Performance-Management System (PMS) must show stable denominators, acceptable defect escapes, manageable false alarms, completed corrective actions and reconciled savings before the factory adds cameras or model classes. This chapter is a proposed deployment and value-realisation method, not a report of an installed GFT system or a guaranteed investment return. Nothing here constitutes investment, legal or tax advice.

14.1

One gated roadmap aligns the factory ramp, the PMS and digital adoption

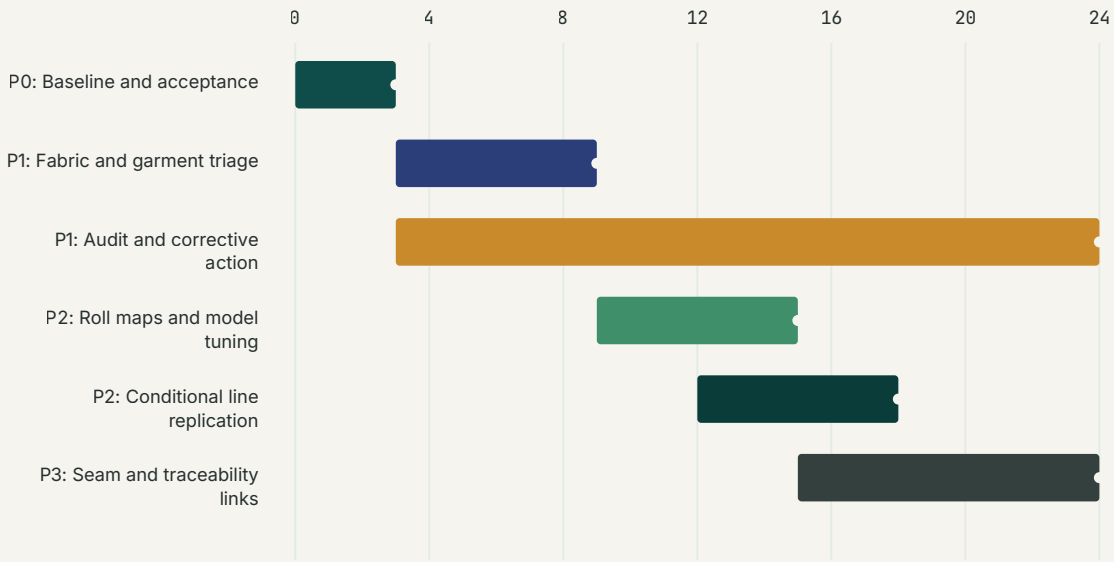
GFT's Cameroon textile programme places feasibility, performance management and digital strategy on one phase spine. That design matters because each workstream removes a different constraint. Feasibility defines the operating boundary and capital sequence. The PMS defines the measures, responsibilities and review cadence. The digital strategy supplies tools only when their inputs and decision rights exist. A camera installed before those three elements agree would create a technical station without an operating purpose. [R27] [R29]

The sequence also respects data dependency. Pre-trained and one-class vision can begin with images of current production rather than a sales ledger. PatchCore explicitly addresses the cold-start problem by fitting a model with nominal images only. Demand forecasting occupies a different position: the programme defers local forecasting until useful sales history exists, even though fashion-sector research reports substantial potential gains from artificial intelligence at scale. Generative design remains a later partner-led capability rather than a prerequisite for inspection. [R27] [R29] [R31] [S04]

Exhibit 63, p. 206 makes the proposed timing inspectable. Its bars are planning allocations, not commitments or achieved milestones. The phases overlap because preparation, operation and review must continue together. Manual logging does not stop when the first camera arrives, and vendor acceptance does not end when a model is fine-tuned.

EXHIBIT 63 / ILLUSTRATIVE MONTHS FROM PROGRAMME APPROVAL

Four conditional phases sequence inspection capability with factory and PMS maturity



ILLUSTRATIVE [R27] [R29] [R31] [S04] [S07] [S19] Illustrative GFT deployment plan, not a record of implementation. Phase 0 prepares the measurement system and acceptance evidence; Phase 1 stabilises two inspection boundaries; Phase 2 adds cutting feedback and local model adaptation; Phase 3 adds selected inline seams and traceability. Every transition depends on evidence, not elapsed time alone.

The ramp changes the economics as well as the hardware count. During early production, a station may have too little volume to recover its fixed cost. During later expansion, an unproven station can multiply false alarms and review queues across several lines. The inherited record of CICAM utilisation between 10.83 and 48.9 percent in July 2023 is a warning about underused industrial capacity, not a benchmark that a new vision facility will reproduce. Capital release must therefore follow measured throughput and quality at the actual ramp, not nameplate output. [R12]

Three review rhythms keep the phase spine coherent. The line review protects current production and approves containment. The weekly plant review examines measurement quality, recurring causes and maintenance. The monthly investment review compares verified value with the next phase's full cost. Each review receives the same event identifiers and definitions, so an engineering improvement cannot appear as a financial gain while the production record shows unresolved losses.

14.2

Phase 0 proves the measurement system before approving the vision station

Phase 0 begins with a manual record that can survive automation. The proposed Quality Log captures units inspected, defect count, defect class, severity, roll or bundle, station, style, disposition and verifier. The defect dictionary proposed for this paper contains 42 classes, but a code is useful only when two reviewers apply it to the same physical condition in the same way. Training must therefore use real specimens, borderline cases and explicit severity rules, not merely a list of names. [R29]

The manual baseline needs a defined opportunity and denominator. Fabric inspection may use linear metres, rolls or garment-equivalent panels. Garment inspection may use units, seams or defects per hundred units. These measures should not be merged into one percentage. A baseline also needs a time window and product context: fabric construction, colour, print, style, line, shift and inspector can all alter the difficulty of detection. Without those fields, a later improvement may reflect an easier production mix rather than a better inspection system.

An independent reference sample tests the manual record before it becomes the comparator. A quality lead and a second qualified reviewer independently grade a representative set, resolve disagreements and retain the reasons. The factory then measures agreement by class and severity. This does not make the reference infallible. It creates an auditable basis for estimating false negatives, false positives and pending decisions. The inherited 60 to 70 percent manual-inspection range is context for that test, not a substitute for factory-specific measurement. [R29]

14.3 REPRESENTATIVE LOCAL SPECIMENS DECIDE VENDOR ACCEPTANCE

The acceptance set should cover the materials and defect opportunities the factory actually expects. Incoming trials can include traceable cotton cloth associated with the northern supply base, while recognising that SODECOTON's inherited profile of more than 220,000 producers and nine ginning plants says nothing about the quality of any particular roll. Each specimen still needs its own supplier, lot, construction and condition record. [R03]

A useful trial separates the set used to configure the station from the set used to accept it. The first supports lighting, focus, thresholds and model selection. The second remains untouched until the formal run. It includes nominal production, rare but consequential defects, difficult colours and prints, and samples close to the acceptance boundary. The factory should report class-level recall and precision, review burden per inspected length or unit, and the number of unreviewed cases. One aggregate accuracy figure would hide the errors that determine safety and workload.

The edge-cascade study shows why this granularity matters. On its disjoint 249-image benchmark, a recall-prioritised first stage found all 20 defective images but also flagged 113 of 229 non-defective images. Its authors report that 85 percent of forwarded frames were false alarms and position the system as assisted triage rather than autonomous acceptance. Those results come from four knit-fabric classes and one small benchmark, not a Cameroonian factory. They nevertheless demonstrate why a perfect recall headline can still produce an unworkable review queue. [S07]

Optical and timing tests belong in the same acceptance protocol. A line-scan study demonstrates high-speed acquisition, triggering, illumination and preprocessing in a PCB application, which supports the engineering method but does not establish textile-defect performance. The fabric station must therefore prove resolution, encoder synchronisation, contrast and processing time on its own moving web. The inherited reference of up to 60 metres a minute is an upper contextual benchmark, not a default factory requirement. [R29] [S14]

Finally, the setup phase tests failure recovery before purchase. The station disconnects from the network, loses and regains power under the approved protection arrangement, and restores its local event buffer without duplicate records. The buyer

also verifies export of images, labels, model identifiers and quality events in documented formats. Data ownership, offline inference and an exit path are acceptance criteria, not contractual afterthoughts.

> **GFT NOTE**
>
> Do not accept a vendor demonstration as the factory baseline. GFT recommends a locally witnessed acceptance run on sealed specimens, followed by a shadow-production trial. The contract should preserve factory ownership of captured images and verified labels, permit local inference and data export, and state the evidence required for final acceptance. A supplier's benchmark may shortlist a system; only representative local evidence should release production capital.

14.4

Phase 1 stabilises two inspection boundaries through human triage

In the illustrative roadmap, Phase 1 starts where a detection can prevent a different downstream decision. At fabric entry, the decision is whether to accept, query, splice or map a roll before cutting. At final garment inspection, the decision is whether to pack, repair, downgrade or hold an item. These boundaries have different images, denominators and failure costs. They should be commissioned as two operating stations with a shared event schema, not as two installations of the same generic model.

The fabric station runs in shadow mode before it influences material disposition. The team compares its output with the manual perch and the independent reference sample, then investigates disagreement by defect class and fabric context. A missed hole has a different consequence from an over-flagged cosmetic variation. The operating threshold must reflect that asymmetry while keeping the human review queue within the available inspection capacity. Once the station enters assisted operation, severe or ambiguous cases still require an authorised disposition.

Cold-start models help during this stage, but their benchmark language must remain precise. PatchCore reports image-level anomaly detection AUROC of up to 99.6 percent on MVTec AD, alongside the ability to learn from nominal examples. AUROC on that benchmark is not equivalent to defect recall on Cameroonian cotton, nor does a nominal memory bank remove the need for an acceptance set. The practical benefit is a starting mechanism for unfamiliar fabric; the factory still has to choose and validate its operating threshold. [S04]

The finished-garment study provides a different and equally bounded reference. Its Addis Ababa dataset contains 1,000 trouser images. The reported system achieved 86.7 percent overall accuracy and processed a sample in 0.45 seconds, while class results ranged from 80.0 percent for seam slippage to 96.7 percent for untrimmed threads. Those findings support a cycle-time and workflow hypothesis. They do not justify autonomous release or a universal accuracy target for every garment style. [S19]

Human triage joins these stations to the PMS. In the metal-textile workflow described by Elsen and co-authors, the model handles clear pass cases and sends doubtful cases to a human inspector. The garment adaptation is to route clearly acceptable frames under an approved rule, route anomalies and low-confidence decisions to review, and record the verifier's conclusion alongside the model output. This structure creates useful labels while preserving human responsibility for consequential dispositions. [S34]

14.5 A MANAGEABLE EXCEPTION QUEUE IS AN ACCEPTANCE RESULT

The daily station review tracks more than detections. It records inspected volume, auto-passed volume, flagged volume, verified defects, rejected alarms, pending cases and downstream escapes. These quantities reveal whether the station is operating within its design. A rising pending count can indicate insufficient review capacity. A rising false-alarm count can indicate lint, lighting drift, a new fabric or a threshold problem. A falling defect count accompanied by rising customer escapes can indicate dangerous under-detection rather than improvement.

The edge study's measured throughput also deserves careful interpretation. Its parallel pipeline reached 13.45 frames a second against 9.86 frames a second for a sequential detector loop, but the authors attribute 91 percent of the speed gain to overlapping JPEG decode with inference rather than to the cascade itself. A factory acceptance test should therefore measure the complete data path, including acquisition, decode, inference, display and event writing. Model latency alone cannot prove line capability. [S07]

Corrective action begins during Phase 1, not after model tuning. Statistical Quality Control and Failure Mode and Effect Analysis provide a disciplined route from verified defect patterns to containment and investigation. The Indonesian garment case used p-charts and risk priorities to identify an incompletely controlled print process. The transferable lesson is to connect the signal to a reaction plan, not to transfer its defect rates or risk scores. [S26]

The launch phase can close only when the operating record is complete enough to support comparison. Representative production runs have known denominators, unresolved critical cases are contained, the review queue is staffed, and downstream audits reconcile with station decisions. The team must also demonstrate that a supervisor can open an exception, assign an action and verify closure without re-keying the event. If those conditions fail, the remedy is to stabilise the station, not to add more stations.

14.6

Phase 2 converts verified detections into cutting decisions and local model evidence

In the illustrative roadmap, Phase 2 begins only after Phase 1 has produced a usable record of local operating conditions. Its first purpose is to improve material yield. A roll detection carries cross-web and along-web coordinates against a roll identifier. Spreading and cutting then transform that roll into plies, panels and bundles. Each

transformation needs a linked identifier and an explicit record of splices, end loss and marker movement. Without that genealogy, an accurate roll coordinate cannot reliably locate a cut panel.

The cutting lead approves the disposition rules during early operation. A severe fault may require removal. A minor fault may be acceptable outside a visible panel zone. A questionable mapping may require physical verification rather than automatic nesting. These rules protect two forms of value at once: they avoid sewing known defects into garments, and they avoid discarding usable cloth through excessive caution. The PMS records the resulting recut, splice, downgrade or accepted use so that material savings can be reconciled rather than inferred from model detections.

14.7 LOCAL ADAPTATION FOLLOWS A CONTROLLED COMPARISON

The second purpose of the scale phase is to reduce avoidable review work and improve class-specific detection on local materials. In this illustrative design, verified crops from Phase 1 become training and validation candidates, but they require deduplication, label review and separation by roll, lot or production run. Nearby frames from one defect must not leak into both training and test sets. Rare critical classes should remain visible in evaluation even when that makes the reported result less flattering.

The outline's planning milestone of 10,000 verified local crops is illustrative, not a statistical guarantee of readiness. A smaller, well-represented and independently reviewed set may be more useful than a larger collection of repetitive or inconsistent labels. The release decision should depend on class coverage, label quality, the proposed operating threshold and performance on a held-out production period. Every new model version also needs a rollback path and a record of the data and configuration that produced it.

Patterned fabrics require particular caution. The reverse-distillation fabric study addresses diverse colours, textures and defect types across MVTec AD, AITEX, TILDA and a textile-facility dataset. That breadth supports testing distillation as a candidate for domain variation, not assuming that one nominal model will cover every print. A new motif or dye lot should trigger a defined review and, where necessary, re-enrolment before the station returns to routine operation. [S12]

A candidate model should run beside the incumbent on comparable production before replacement. The review compares critical-class recall, false alarms per inspected unit or length, pending volume, processing time and downstream escapes. It also examines whether the proposed gain persists across lines, shifts and material lots. A model that improves the aggregate score while weakening a critical class does not pass. Approval covers a named version at a named threshold, not an indefinite permission to update weights.

Phase 2 is also the earliest sensible point for local demand forecasting in the programme sequence, because the factory now has some sales and order history. The fashion-sector figures in the shared register describe potential from artificial intelligence at scale, not a guaranteed result for this plant. Forecasting should remain a separate acceptance exercise with its own baseline and error measures. It must not borrow credibility from the vision station or divert attention from unresolved quality losses. [R27]

Replication follows the same evidence rule as model updates. A second line receives the proven station design only after its own lighting, mechanics, product mix and reference sample have been checked. The original station's event schema and reaction plan can transfer. Its thresholds and performance cannot be assumed to transfer without verification.

14.8

Phase 3 adds critical-seam control and traceability without destabilising the line

In the illustrative roadmap, Phase 3 addresses defects created between the cut bundle and the finished garment. The factory selects critical seams from verified Pareto evidence, buyer requirements and failure severity. It does not place a camera at every needle by default. A station earns its position when earlier detection can contain a suspect bundle, guide a needle or tension check, and prevent repeated rework without creating an unmanageable stop pattern.

The sewing-model evidence supports a measured trial rather than a broad autonomy claim. In one study, adding Coordinate Attention to the authors' YOLOv8 configuration increased mAP@0.5 by 3.7 percentage points. The complete improved configuration also reduced speed from 48 to 43 frames a second on the reported experimental setup. Its authors acknowledge limited generalisation to complex textured backgrounds. A Cameroonian acceptance run must therefore test both detection and production consequences on its own garments. [S17]

Inline reactions should be proportional to risk. A critical seam failure can require immediate containment. Repeated minor faults may justify a run-rule, a supervisor review or a maintenance check rather than an automatic stop. Every intervention records stopped time, affected bundles and the subsequent verified result. This prevents the quality dashboard from claiming a gain while the production record hides lost availability.

Traceability integration follows the same condition gate. Research on Digital Product Passport readiness describes fragmented upstream data, limited digital systems and methodological complexity as barriers for garment manufacturers. It also notes that textile-specific requirements continue to be defined through later acts and product rules. Phase 3 should therefore make quality and material records exportable and traceable, without claiming that a particular passport format is already sufficient for future compliance. [S27]

Partner-led design experiments may begin only after the production data spine is stable. The shared register describes multi-model pipelines spanning fabric construction, shade, weave and silhouette, but those capabilities do not repair an unreliable inspection record. They remain a separate Tier 3 pilot with their own acceptance criteria. The run phase is complete when added digital scope no longer degrades containment, review workload or data integrity. [R31]

14.9

Condition gates and paired controls decide whether the programme advances

Exhibit 64, p. 212 converts the roadmap into an evidence contract. The thresholds are deliberately expressed as approved local criteria rather than universal numbers. Before each trial, the sponsor, production lead and quality lead must convert them into measurable values suitable for the product and buyer. They must also state who signs the gate and what evidence would block it.

EXHIBIT 64 / ILLUSTRATIVE GFT PHASE-GATE CONTRACT

Each phase releases capital only after operational and economic evidence agrees

Gate	Required evidence	Capital or scope released	Hold condition
P0 to P1	Stable manual schema; sealed local acceptance; offline recovery	Two assisted inspection boundaries	Unresolved labels, optics or export failure
P1 to P2	Complete denominators; controlled escapes; staffed review; closed actions	Roll-to-marker link and candidate model tuning	Pending critical cases or unreconciled downstream audits
P2 to P3	Matched-run model gain; conserved roll genealogy; verified material value	Selected critical seams and further lines	Critical-class regression or broken coordinate mapping
P3 to run	Stable added stations; traceable exports; cumulative cash reconciliation	Routine operation and controlled renewal	Rising stops, incomplete records or unverified savings

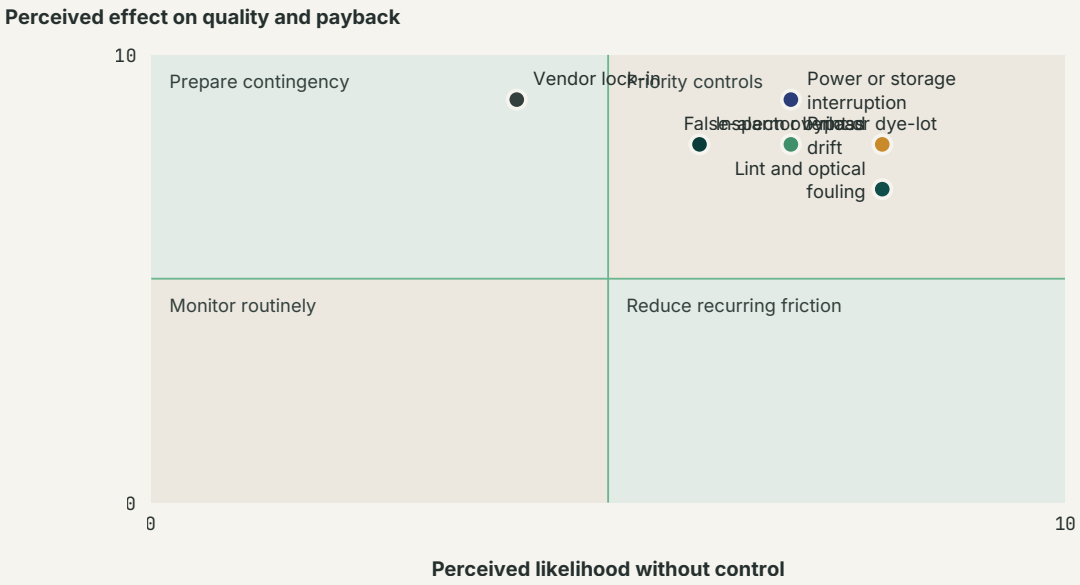
ILLUSTRATIVE [R29] [S04] [S07] [S19] [S24] [S26] Illustrative governance design. Cited studies support inspection, triage and management mechanisms, not these exact gates. Local owners must set numeric acceptance limits before each trial and retain the signed evidence.

The gate is not a claim that management integration guarantees quality. In the Central Java garment study, digital quality control and production automation jointly explained 64.5 percent of product-quality variance through a survey-based structural model. That association supports combining technology with management routines, but it does not provide a causal effect or a phase threshold for another factory. Local gate evidence remains necessary. [S24]

Exhibit 65, p. 213 positions the main deployment risks on two illustrative scales. The coordinates are workshop planning scores, not observed frequencies. Their purpose is to pair every material risk with an engineering control and a management owner before capital is released.

EXHIBIT 65 / ILLUSTRATIVE SCORE FROM 0 TO 10

Six deployment risks require both a technical control and a management response



ILLUSTRATIVE [S07] [S12] [S26] [S34] Illustrative workshop scores, not measured likelihoods or severities. Higher x means greater perceived operational likelihood without controls; higher y means greater perceived effect on quality and payback. Teams must rescore with local FMEA evidence.

Lint and optical fouling require protected housings, cleaning standard work and image-quality checks. Power and storage interruption require conditioned power, clean shutdown, durable local event identifiers and duplicate-free recovery. Print or dye-lot drift requires a reference check, controlled re-enrolment and a route back to the prior approved configuration. These are station-level controls, but the PMS must record completion and exceptions if they are to remain effective.

False-alarm overload requires class-specific threshold review and enough verifier capacity. Inspector bypass requires usable screens, training, measured agreement and a response policy that does not punish honest escalation. Vendor lock-in requires portable data, documented export, local inference and contractual exit rights. The SQC and FMEA method is useful here because it links risk priority to concrete work standards and equipment checks rather than leaving risk as a colour on a matrix. [S26]

Every gate also needs explicit stop conditions. A critical defect escape, a broken roll-to-bundle chain, a rising unreviewed queue, unexplained denominator loss or an irreconcilable savings claim should hold expansion. The response may be to return to shadow mode, restore the previous model, retrain staff or repair the measurement system. A phase gate is credible only when the factory is willing to use it.

14.10

The PMS proves payback only when physical and financial records reconcile

The value-realisation scorecard begins with external reference points, then replaces them with local measurements. **Exhibit 66**, p. 214 keeps the cited benchmarks visible without presenting them as GFT results. Their scopes differ: classification accuracy, web speed, frame throughput and return estimates cannot be combined into one performance score.

EXHIBIT 66 / REPORTED RESEARCH AND INHERITED REFERENCE VALUES

External benchmarks guide acceptance targets but do not prove local value

60 to 70%

Manual inspection reference

Inherited manual defect-inspection accuracy range, R29

>90%

Automated inspection reference

Inherited computer-vision summary across more than 40 defect types, R29

Up to 60 m/min

Fabric web speed reference

Contextual line-scan reference, not a site requirement, R29

13.45 FPS

Edge pipeline throughput

Parallel Jetson Nano pipeline in S07

+3.7 points

Sewing-model change

mAP@0.5 change with Coordinate Attention in S17

12 to 18 months

Research payback estimate

Estimated for medium-sized Ethiopian factories in S19

SOURCED [R29] [S07] [S17] [S19] Sourced reference values from different studies and datasets. R29 is the inherited inspection summary. S07 concerns a four-class knit-fabric edge cascade. S17 concerns a sewing-model ablation. S19 reports an Ethiopian garment study and an estimated return, not an achieved GFT payback.

The internal scorecard has five linked tests. Detection performance measures verified false negatives and false positives by class. Flow performance measures Right First Time, Defects Per Hundred Units, pending volume and review time. Material performance measures fabric accepted, spliced, recut, downgraded and scrapped under a conserved balance. Workforce performance measures agreement, training and action closure. Financial performance measures realised savings and the full cost of owning and operating each station.

Attribution requires a baseline and a counterfactual. The team compares equivalent products and operating conditions, retains changes in style and material mix, and distinguishes the effect of the camera from maintenance, training or supplier action. Where several interventions occur together, the review should report the combined effect rather than assign the same saving to each intervention. Detection counts alone are not benefits.

METHOD

Monthly value and payback boundary

For benefit bucket b, realised monthly value equals the verified physical change multiplied by an approved unit value. Net monthly value is the sum of buckets less incremental licence, maintenance, power, review and support costs. Cumulative net cash value is the sum of monthly net values less capital expenditure already incurred. Payback occurs in the first month when cumulative net cash value becomes positive, subject to finance review. The model must not count avoided fabric loss and avoided garment cost on the same item, treat detection as a completed saving, or annualise a pilot month as if the whole factory had reached stable output.

WORKED EXAMPLE: A REPRESENTATIVE STABLE MONTH

This illustrative case examines 100,000 garment-equivalent units inspected in one stable month. It is not a forecast for the factory. Manual incoming-fabric detection is assumed at 65 percent, the midpoint of the inherited 60 to 70 percent reference, and assisted vision detection at 90 percent. The assumed consequential incoming-fabric defect rate is 2.0 percent. Vision therefore intercepts an additional 100,000 × 0.02 x (0.90 - 0.65) = 500 garment equivalents before the assumed downstream conversion loss. [R29]

Each additional interception is assigned an illustrative avoided conversion cost of XAF 1,200, giving XAF 600,000 a month. This bucket excludes the fabric itself and any supplier claim, which would require separate evidence. It also excludes any item already counted in sewing seconds.

An illustrative reduction of 600 reworked garments at 12 minutes each saves 120 direct labour hours. Using the inherited XAF 43,969 monthly wage floor and a 40-hour week gives an illustrative hourly reference of XAF 43,969 / (40 × 52 / 12) = XAF 253.67. The direct labour saving is only about XAF 30,440 a month. This deliberately small figure shows why labour substitution is not the main value mechanism. Actual costing must use the appropriate employment and overhead basis rather than the statutory floor alone. [R14] [R16]

A separate illustrative reduction of 400 downgraded sewing seconds at XAF 900 contribution loss per item adds XAF 360,000. Gross monthly benefit is therefore XAF 990,440. Subtracting illustrative incremental monthly operating cost of XAF 150,000 leaves XAF 840,440. Against illustrative installed capital of XAF 12,000,000, simple stable-rate payback is XAF 12,000,000 / XAF 840,440 = 14.3 months.

Exhibit 67, p. 217 reconciles those buckets. If the fabric-conversion benefit were absent while all other illustrative assumptions remained unchanged, net monthly value would fall to XAF 240,440 and simple payback would extend to 49.9 months. The scenario is a sensitivity test, not a confidence interval. It demonstrates why the investment case depends on material and disposition evidence rather than on inspection labour savings alone.

EXHIBIT 67 / XAF PER REPRESENTATIVE STABLE MONTH

Material avoidance, not inspection labour, carries the illustrative monthly value case



ILLUSTRATIVE [R14] [R16] [R29] Illustrative arithmetic from the worked example, not observed factory value. Additional fabric interceptions: 100,000 × 0.02 x (0.90 - 0.65) = 500; x XAF 1,200 = XAF 600,000. Rework: 600 × 12 / 60 = 120 hours; x XAF 253.667 = XAF 30,440. Seconds: 400 x XAF 900 = XAF 360,000. Net = 600,000 + 30,440 + 360,000 - 150,000 = XAF 840,440. Simple payback on illustrative XAF 12,000,000 capital = 14.3 months.

The worked example also explains the reporting sequence. First, the Quality Log verifies the additional interception and the garment disposition. Next, the material and line records prove that the physical loss did not occur elsewhere. Finance then approves the unit value and operating cost. Finally, the investment review updates cumulative cash using actual ramp output. A stable-rate calculation can screen a proposal, but only the cumulative record can declare capital recovered.

The Ethiopian garment study's 12 to 18 month return estimate provides a useful comparison for the illustrative 14.3-month result, not independent confirmation. Geography, product, defect prevalence, equipment scope and cost assumptions differ. The correct conclusion is narrower: a phased vision programme can test payback inside the PMS, and the test is credible only when physical reconciliation and full operating cost remain visible. [S19]

GFT's final recommendation is to fund the next phase from evidence produced by the current one. Start with a trustworthy manual record. Stabilise two high-decision-value inspection boundaries. Add cutting feedback and model adaptation after veri-

fied local data exist. Add inline seams and traceability only when the factory can absorb them. Throughout, keep production as the warp and measurement as the weft, so every expansion carries its own proof of quality and value.

Scale should follow the evidence of value, never the schedule of procurement.

CHAPTER SOURCE NOTES

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S04 arXiv / IEEE CVPR. 2021-06-15.

S07 arXiv. 2026-08-12.

S19 Textile Research Journal. 2026-09-06.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

R31 arXiv 2509.04705. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

S14 Applied Sciences (MDPI). 2024-12-23.

S34 IntechOpen. 2025-01-03.

S26 Jurnal Teknik Industri (JATRI). 2026-09-25.

S12 arXiv. 2024-01-04.

S17 Journal of Measurements in Engineering. 2024-09-09.

S27 Sustainability (MDPI). 2026-07-13.

S24 CEMERLANG: Jurnal Manajemen dan Ekonomi Bisnis. 2026-08-03.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

Provenance travels with the claim.

Available inherited records and all chapter-local source records are graded below. Identical source identities are consolidated without losing citation aliases. Missing annex references are excluded. Unrecorded dates are disclosed rather than invented.

R01 Inherited evidence R01 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Cameroon seed-cotton production 2022-2024: 330,000 t (2022), 314,455 t (2023), about 340,000 t (2024).

Inherited source-text check; original source not independently re-opened for this edition.

R02 Inherited evidence R02 GRADE C

Fibre2Fashion; Business in Cameroon. Publication date unrecorded. fibre2fashion.com; businessincameroon.com

2025 seed-cotton projection above 350,000 t.

Inherited source-text check; original source not independently re-opened for this edition.

R03 Inherited evidence R03 GRADE C

SODECOTON. Publication date unrecorded. sodecoton.cm

Company profile: more than 220,000 cotton producers supported in northern Cameroon; nine ginning plants; two oil mills.

Inherited source-text check; original source not independently re-opened for this edition.

R04 Inherited evidence R04 GRADE C

Financial Fortune Media. Publication date unrecorded. financialfortunemedia.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R05 Inherited evidence R05 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R06 Inherited evidence R06 GRADE C

Cameroon Trade Hub. Publication date unrecorded. cameroontradehub.cm

About 95 percent of cotton fibre exported raw against 5 percent processed locally; industrialisation target of 50 percent local processing.

Inherited source-text check; original source not independently re-opened for this edition.

R07 Inherited evidence R07 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

SND30 restates the 50 percent local-processing target by 2030.

Inherited source-text check; original source not independently re-opened for this edition.

R08 Inherited evidence R08 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R09 Inherited evidence R09 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA integrated cotton-to-textile project, US$300 million (180 billion XAF); Panafritex and Arise with SODECOTON, CNPS and Marlo; 12 November 2025.

Inherited source-text check; original source not independently re-opened for this edition.

R10 Inherited evidence R10 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA: 12,000 t a year processed; about 12,000 jobs at the Dibamba hub and 3,000 at Garoua; return on investment near 11 years; optimal capacity around 2033.

Inherited source-text check; original source not independently re-opened for this edition.

R11 Inherited evidence R11 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

CICAM market share from about 85 percent three decades ago to about 5 percent today.

Inherited source-text check; original source not independently re-opened for this edition.

R12 Inherited evidence R12 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

CICAM capacity utilisation, July 2023 report: Garoua 33.7 percent, CICAM 1 48.9 percent, CICAM 2 10.83 percent; underutilisation cost about CFAF 2 billion in 2020.

Inherited source-text check; original source not independently re-opened for this edition.

R13 Inherited evidence R13 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R14 Inherited evidence R14 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Cameroon minimum wage (SMIG) XAF 41,875 a month in 2023, raised to XAF 43,969 a month from 23 February 2024, about US$76, a rise of 5 percent.

Inherited source-text check; original source not independently re-opened for this edition.

R15 Inherited evidence R15 GRADE C

wage.is. Publication date unrecorded. wage.is/cameroon

Average salary across sectors in Cameroon about US$359 a month.

Inherited source-text check; original source not independently re-opened for this edition.

R16 Inherited evidence R16 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Informal share of employment about 90 percent; standard 40-hour week, 48-hour maximum, overtime at 1.2 times.

Inherited source-text check; original source not independently re-opened for this edition.

R17 Inherited evidence R17 GRADE C

just-style. Publication date unrecorded. just-style.com

Cambodia garment-sector minimum wage US$204 a month for 2024 (labour-cost benchmark).

Inherited source-text check; original source not independently re-opened for this edition.

R18 Inherited evidence R18 GRADE B

Congressional Research Service IF10149; AGOA (Wikipedia). Publication date unrecorded. congress.gov; en.wikipedia.org

AGOA: Cameroon terminated effective 1 January 2020 for gross human-rights violations; remained ineligible; listed ineligible for calendar year 2025.

Inherited source-text check; original source not independently re-opened for this edition.

R23 Inherited evidence R23 GRADE C

MacMap; AfCFTA. Publication date unrecorded. macmap.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R24 Inherited evidence R24 GRADE C

Mordor Intelligence. Publication date unrecorded. mordorintelligence.com

Africa textile market about US$37.3 billion in 2025, projected US$49.41 billion by 2030 at 4.73 percent a year.

Inherited source-text check; original source not independently re-opened for this edition.

R25 Inherited evidence R25 GRADE C

Market.us. Publication date unrecorded. market.us

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R26 Inherited evidence R26 GRADE C

Statista. Publication date unrecorded. statista.com

Cameroon clothing and footwear spend about US$94.70 a person in 2024 and about 8.0 pieces a person; apparel eCommerce revenue US$20.32 million (2023) to US$33.61 million (2027) at 13.41 percent a year; user penetration 18.0 percent (2023) to 22.7 percent (2027).

Inherited source-text check; original source not independently re-opened for this edition.

R27 Inherited evidence R27 GRADE B

McKinsey, The State of Fashion 2025. Publication date unrecorded. mckinsey.com

AI at scale in fashion: inventory cost down 35 to 40 percent; forecast accuracy up 20 to 30 percent; revenue per customer up 15 to 25 percent; forecast errors cut 20 to 50 percent; lost sales down up to 65 percent; 75 percent of executives prioritising; sector operating-profit potential US$150 to 275 billion over three to five years.

Inherited source-text check; original source not independently re-opened for this edition.

R28 Inherited evidence R28 GRADE C

DataIntelo. Publication date unrecorded. dataintelo.com

Fabric-defect-detection AI market US$523.4 million (2024) to US$6,522.1 million (2033) at 33.8 percent a year; Asia-Pacific about 46 percent in 2024.

Inherited source-text check; original source not independently re-opened for this edition.

R29 Inherited evidence R29 GRADE C

MDPI; industry inspection literature. Publication date unrecorded. mdpi.com

Manual defect inspection 60 to 70 percent accurate, missing 20 to 30 percent of defects; computer vision above 90 percent across more than 40 defect types; line-scan inspection up to 60 metres a minute.

Inherited source-text check; original source not independently re-opened for this edition.

R30 Inherited evidence R30 GRADE C

Design Journal (Taylor and Francis). Publication date unrecorded. tandfonline.com

Generative AI in apparel 2023-2024 mainly design-side image generation (Midjourney, Stable Diffusion) across style, silhouette, colour and pattern.

Inherited source-text check; original source not independently re-opened for this edition.

R31 Inherited evidence R31 GRADE C

arXiv 2509.04705. Publication date unrecorded. arxiv.org/abs/2509.04705

By 2025-2026 the field moved to multi-model garment-design pipelines (fabric construction, shade, weave, silhouette) and physics-constrained generation coupling neural generation with simulation validation.

Inherited source-text check; original source not independently re-opened for this edition.

R32 Inherited evidence R32 GRADE C

Wise; XE. Publication date unrecorded. wise.com; xe.com

CFA franc to US dollar reference rates; franc pegged to the euro at 655.957. The GFT study's method note places the franc near 590 to 610 per US dollar across the 2023 study window and near 563 to 577 more recently, and treats currency conversions as estimates.

Inherited source-text check; original source not independently re-opened for this edition.

c01-S01 YOLO Object Detection for Real-Time Fabric Defect Inspection in the Textile Industry: A Review of YOLOv1 to YOLOv11 GRADE B

Sensors (MDPI). 2025-04-03. https://doi.org/10.3390/s25072270

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S04 Towards Total Recall in Industrial Anomaly Detection GRADE B

arXiv / IEEE CVPR. 2021-06-15. https://arxiv.org/abs/2106.08265

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S05 PaDiM: a Patch Distribution Modeling Framework for Anomaly Detection and Localization GRADE B

arXiv / ICPR. 2020-11-17. https://arxiv.org/abs/2011.08785

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S07 **Low Cost Two-Stage Fabric Defect Detection at the Edge** GRADE B

arXiv. 2026-08-12. https://arxiv.org/abs/2608.14727

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S18 **Deep-GD: Deep Learning based Automatic Garment Defect Detection and Type Classification** GRADE B

International Journal of Electrical and Electronics Research. 2024-02-05. https://doi.org/10.37391/ijeer.120107

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S19 **Development of an automatic garment defect recognition system using a machine learning approach** GRADE B

Textile Research Journal. 2026-09-06. https://doi.org/10.1177/00405175261484722

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S31 **Transforming Fashion with AI: A Comparative Study of Large Language Models in Apparel Design** GRADE B

arXiv. 2025-09-04. https://arxiv.org/abs/2509.04705

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S36 **World Development Indicators: Manufacturing, value added (% of GDP) - Cameroon (NV.IND.MANF.ZS)** GRADE A

World Bank. 2026-10-08. https://api.worldbank.org/v2/country/CMR/indicator/NV.IND.MANF.ZS?date=2020:2024&format=json

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S01 **YOLO Object Detection for Real-Time Fabric Defect Inspection in the Textile Industry: A Review of YOLOv1 to YOLOv11** GRADE B

Sensors (MDPI). 2025-04-03. https://www.mdpi.com/1424-8220/25/7/2270

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S03 **Fabric Defect Detection Using Computer Vision Techniques: A Comprehensive Review** GRADE B

Mathematical Problems in Engineering (Wiley). 2020-11-16. https://onlinelibrary.wiley.com/doi/10.1155/2020/8189403

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S08 **Fabric Defect Detection Using Deep Learning** GRADE C

International Journal for Research in Applied Science and Engineering Technology. 2023-05-31. https://www.ijraset.com/research-paper/fabric-defect-detection-using-deep-learning

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S19 **Development of an automatic garment defect recognition system using a machine learning approach** GRADE B

Textile Research Journal (SAGE Publications). 2026-09-06. https://api.crossref.org/works/10.1177/00405175261484722

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S21 **Image analysis for seam-puckering evaluation** GRADE B

Textile Research Journal (SAGE Publications). 2016-10-14. https://api.crossref.org/works/10.1177/0040517516673330

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S0202 **Comparative Study of the Inspection Parameters for AQL 2.5% and 1.5% in Garments Manufacturing Process** GRADE B

Journal of Textile Science & Fashion Technology. 2019-05-16. https://doi.org/10.33552/jtsft.2019.02.000544

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S0203 **An approach to automatic fault detection in four-point system for knitted fabric with our benchmark dataset Isl-Knit** GRADE B

Heliyon (Elsevier). 2024-08-24. https://pmc.ncbi.nlm.nih.gov/articles/PMC11639326/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S0204 **DEVELOPING AN INTELLIGENT AUTOMATION AND REPORTING SYSTEM FOR FABRIC INSPECTION MACHINES** GRADE B

TEKSTIL VE KONFEKSIYON. 2019-03-26. https://dergipark.org.tr/en/pub/tekstilvekonfeksiyon/article/491362

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S0205 **Visual Inspection: A Review of the Literature** GRADE A

Sandia National Laboratories, United States Department of Energy. 2012-10-01. https://digital.library.unt.edu/ark:/67531/metadc835891/m2/1/high_res_d/1055636.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S15 **Online Fabric Defect Inspection Using Smart Visual Sensors** GRADE B

Sensors (MDPI). 2013-04-09. https://doi.org/10.3390/s130404659

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S18 Same source record as c01-S18.

c03-S21 **Image analysis for seam-puckering evaluation** GRADE B

Textile Research Journal. 2016-10-14. https://doi.org/10.1177/0040517516673330

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S22 **Defect Detection Techniques for Airbag Production Sewing Stages** GRADE B

arXiv. 2015-10-25. https://arxiv.org/abs/1510.07905

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S25 **Quality control of garment product using DMAIC Six Sigma** GRADE B

OPSI. 2024-06-30. https://doi.org/10.31315/opsi.v17i1.7107

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S0301 **Standard Terminology Relating to Fabric Defects, ASTM D3990-12(2020), preview** GRADE A

ASTM International. 2020-07. https://cdn.standards.iteh.ai/samples/106846/be97e8fde51b41fda578dbdcdf7cb076/ASTM-D3990-12-2020-.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c03-S0302 **Surface Ornamentation Techniques (Embroidery), Trade Theory, revised syllabus 2022, first edition February 2023** GRADE A

National Instructional Media Institute, Directorate General of Training, Government of India. 2023-02. https://bharatskills.gov.in/pdf/E_Books/CTS/129/English/TT/SOT%20-%20TT%20%282022%29.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S01 Same source record as c01-S01.

---

c04-S02 **Deep Convolutional Neural Network Optimization for Defect Detection in Fabric Inspection** GRADE B

Sensors (MDPI). 2021-10-25. https://doi.org/10.3390/s21217074

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S03 **Fabric Defect Detection Using Computer Vision Techniques: A Comprehensive Review** GRADE B

Mathematical Problems in Engineering. 2020-11-16. https://doi.org/10.1155/2020/8189403

Source metadata retained for traceability. This entry is not cited in the chapter text.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S04 Same source record as c01-S04.

c04-S06 **One-Class Model for Fabric Defect Detection** GRADE B

arXiv / Computer Science and Information Technology. 2022-04-20. https://arxiv.org/abs/2204.09648

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S07 Same source record as c01-S07.

---

c04-S08 **Fabric Defect Detection Using Deep Learning** GRADE C

International Journal for Research in Applied Science and Engineering Technology. 2023-05-31. https://doi.org/10.22214/ijraset.2023.53502

Source metadata retained for traceability. This entry is not cited in the chapter text.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S11 **SPFFNet: Strip Perception and Feature Fusion Spatial Pyramid Pooling for Fabric Defect Detection** GRADE B

arXiv. 2025-02-03. https://arxiv.org/abs/2502.01445

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S12 **Distillation-based fabric anomaly detection** GRADE B

arXiv. 2024-01-04. https://arxiv.org/abs/2401.02287

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S14 **Design and Development of a Precision Defect Detection System Based on a Line Scan Camera Using Deep Learning** GRADE B

Applied Sciences (MDPI). 2024-12-23. https://doi.org/10.3390/app142412054

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S15 Same source record as c03-S15.

---

c04-S16 **DF-YOLO: An attempt on enhancing generalization in fabric defect detection based on YOLO network** GRADE B

Textile Research Journal. 2024-09-24. https://doi.org/10.1177/00405175241269163

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S17 **Improved YOLOv8 garment sewing defect detection method based on attention mechanism** GRADE B

Journal of Measurements in Engineering. 2024-09-09. https://doi.org/10.21595/jme.2024.24283

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S18 Same source record as c01-S18.

c04-S19 Same source record as c01-S19.

c04-S20 Intelligent Detection Algorithm for Garment Pattern Defects Based on Computer Vision and Deep Learning GRADE B

International Journal of e-Collaboration. 2026-10-26. https://doi.org/10.4018/ijec.426013

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S21 Same source record as c03-S21.

c04-S22 Same source record as c03-S22.

c04-S23 On-Loom Fabric Defect Inspection Using Contact Image Sensors and Activation Layer Embedded Convolutional Neural Network GRADE B

University of North Texas Libraries. 2024-05-01. https://doi.org/10.12794/metadc1404537

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S24 Digitalisasi Quality Control dan Otomatisasi Produksi: Analisis Pengaruhnya Terhadap Kualitas Produk Garment di PT X, Perusahaan Garment di Jawa Tengah GRADE B

CEMERLANG: Jurnal Manajemen dan Ekonomi Bisnis. 2026-08-03. https://doi.org/10.55606/cemerlang.v6i3.9950

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S25 Same source record as c03-S25.

c04-S26 Analysis of Quality Control and Product Failure Risks in Garment Manufacturing Using Statistical Quality Control (SQC) and Failure Mode and Effect Analysis (FMEA) Case Study: CV Muharram Garment Indonesia GRADE B

Jurnal Teknik Industri (JATRI). 2026-09-25. https://doi.org/10.30872/jatri424894

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S32 Fabric Defect Detection in Textile Manufacturing: A Survey of the State of the Art GRADE B

Security and Communication Networks. 2021-05-10. https://doi.org/10.1155/2021/9948808

Source metadata retained for traceability. This entry is not cited in the chapter text.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0401 Search algorithm for determining the illumination frequency in a line scan camera-based measurement system GRADE B

Production Systems and Information Engineering. 2025-12-09. https://doi.org/10.32968/psaie.2025.1.4

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0402 **Fabric Defect Detection Using Vision-Based Tactile Sensor** GRADE B

arXiv. 2020-03-02. https://arxiv.org/abs/2003.00839

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0403 **LUSITANOv2: A Real-World Dataset for Fabric Defect Detection in Active Textile Production** GRADE B

Electronics (MDPI). 2026-09-24. https://doi.org/10.3390/electronics15194403

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0404 **Novel online fabric color difference detection system based on machine vision** GRADE B

Scientific Reports (Nature Portfolio). 2026-06-22. https://doi.org/10.1038/s41598-026-58724-3

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0405 **Industrial Artificial Intelligence Methods for Fabric Color-Defect Detection in Textile Manufacturing** GRADE B

International Journal Science and Technology. 2026-03-01. https://doi.org/10.56127/ijst.v5i1.2600

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0406 **A Dataset of Raw Fabric Grayscale Images for Defect Detection** GRADE B

Data (MDPI). 2026-05-12. https://doi.org/10.3390/data11050116

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S01 Same source record as c01-S01.

c05-S02 Same source record as c04-S02.

c05-S04 Same source record as c01-S04.

c05-S06 Same source record as c04-S06.

c05-S09 **High-Precision Fabric Defect Detection via Adaptive Shape Convolutions and Large Kernel Spatial Modeling** GRADE B

arXiv. 2025-01-24. https://arxiv.org/abs/2501.14190

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S10 **Fab-ME: A Vision State-Space and Attention-Enhanced Framework for Fabric Defect Detection** GRADE B

arXiv. 2024-12-04. https://arxiv.org/abs/2412.03200

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c05-S11 Same source record as c04-S11.

---

c05-S12 Same source record as c04-S12.

---

c05-S16 Same source record as c04-S16.

---

c05-S32 **Fabric Defect Detection in Textile Manufacturing: A Survey of the State of the Art** GRADE B

Security and Communication Networks. 2021-05-10. https://onlinelibrary.wiley.com/doi/10.1155/2021/9948808

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c05-S33 **An Automated Framework for Fabric Defect Detection in Textile Inspection** GRADE B

Journal of Information Systems Engineering and Management. 2025-05-12. https://doi.org/10.52783/jisem.v10i46s.8815

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c05-S35 **Differentiable NMS via Sinkhorn Matching for End-to-End Fabric Defect Detection** GRADE B

arXiv. 2025-05-11. https://arxiv.org/abs/2505.07040

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c06-S01 Same source record as c01-S01.

---

c06-S02 Same source record as c04-S02.

---

c06-S04 Same source record as c01-S04.

---

c06-S05 Same source record as c01-S05.

---

c06-S06 Same source record as c04-S06.

---

c06-S07 Same source record as c01-S07.

---

c06-S08 Same source record as c04-S08.

---

c06-S09 Same source record as c05-S09.

c06-S11 Same source record as c04-S11.

c06-S12 Same source record as c04-S12.

c06-S13 A Unified Anomaly Synthesis Strategy with Gradient Ascent for Industrial Anomaly Detection and Localization

GRADE B

arXiv. 2024-07-12. https://arxiv.org/abs/2407.09359

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S14 Same source record as c04-S14.

c06-S15 Same source record as c03-S15.

c06-S18 Same source record as c01-S18.

c06-S32 Same source record as c04-S32.

c06-S34 Anomaly Detection in Metal-Textile Industries

GRADE B

IntechOpen. 2025-01-03. https://doi.org/10.5772/intechopen.1008411

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S35 Same source record as c05-S35.

c06-S0601 Self-supervised fabric defect segmentation via texture regularity learning

GRADE B

Elsevier BV (SSRN). 2026-10-09. https://doi.org/10.2139/ssrn.7589368

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0602 Unsupervised textile defect detection using convolutional neural networks

GRADE B

arXiv. 2023-11-30. https://arxiv.org/abs/2312.00224

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0603 The MVTec Anomaly Detection Dataset: A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection

GRADE B

International Journal of Computer Vision (Springer). 2021-01-06. https://doi.org/10.1007/s11263-020-01400-4

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0604 **Anomaly Detection via Reverse Distillation from One-Class Embedding** GRADE B

arXiv / IEEE CVPR. 2022-01-26. https://arxiv.org/abs/2201.10703

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0605 **SimpleNet: A Simple Network for Image Anomaly Detection and Localization** GRADE B

arXiv / IEEE CVPR. 2023-03-27. https://arxiv.org/abs/2303.15140

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S17 **Improved YOLOv8 garment sewing defect detection method based on attention mechanism** GRADE A

JVE International, Journal of Measurements in Engineering. 2024-09-09. https://www.extrica.com/article/24283

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S18 **Deep-GD: Deep Learning based Automatic Garment Defect Detection and Type Classification** GRADE A

FOREX Publication, International Journal of Electrical and Electronics Research. 2024-02-05. https://ijeer.forexjournal.co.in/archive/volume-12/ijeer-120107.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S19 **Development of an automatic garment defect recognition system using a machine learning approach** GRADE A

SAGE, Textile Research Journal. 2026-09-06. https://journals.sagepub.com/doi/10.1177/00405175261484722

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S21 **Image analysis for seam-puckering evaluation** GRADE A

SAGE, Textile Research Journal. 2016-10-14. https://journals.sagepub.com/doi/abs/10.1177/0040517516673330

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S22 Same source record as c03-S22.

c07-S0701 **An Intelligent Solution for Automatic Garment Measurement Using Image Recognition Technologies** GRADE A

MDPI, Applied Sciences; author record at Kaunas University of Technology. 2022-04-28. https://www.mdpi.com/2076-3417/12/9/4470

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S01 Same source record as c01-S01.

c08-S02 Same source record as c04-S02.

c08-S04 Same source record as c01-S04.

c08-S05 Same source record as c01-S05.

c08-S06 Same source record as c04-S06.

c08-S07 Same source record as c01-S07.

c08-S14 Same source record as c04-S14.

c08-S15 Same source record as c03-S15.

c08-S17 Same source record as c04-S17.

c08-S18 Same source record as c01-S18.

c08-S19 Same source record as c01-S19.

c08-S20 Same source record as c04-S20.

c08-S23 On-Loom Fabric Defect Inspection Using Contact Image Sensors and Activation Layer Embedded Convolutional Neural Network GRADE B

University of North Texas Libraries. 2024. https://doi.org/10.12794/metadc1404537

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S24 Same source record as c04-S24.

c08-S26 Same source record as c04-S26.

c08-S34 Same source record as c06-S34.

c08-S36 Same source record as c01-S36.

c08-S0801 World Development Indicators: Electric power transmission and distribution losses (% of output) - Cameroon (EG.ELC.LOSS.ZS) GRADE A

World Bank. 2026-10-08. https://api.worldbank.org/v2/country/CMR/indicator/EG.ELC.LOSS.ZS?date=2018:2024&format=json

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S0802 World Development Indicators: Fixed broadband subscriptions (per 100 people) - Cameroon (IT.NET.BBND.P2) GRADE A

World Bank. 2026-10-08. https://api.worldbank.org/v2/country/CMR/indicator/IT.NET.BBND.P2?date=2018:2024&format=json

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S0803 World Development Indicators: Access to electricity (% of population) - Cameroon (EG.ELC.ACCS.ZS) GRADE A

World Bank. 2026-10-08. https://api.worldbank.org/v2/country/CMR/indicator/EG.ELC.ACCS.ZS?date=2018:2024&format=json

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S02 Deep Convolutional Neural Network Optimization for Defect Detection in Fabric Inspection GRADE B

Sensors (MDPI). 2021-10-25. https://www.mdpi.com/1424-8220/21/21/7074

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S04 Same source record as c01-S04.

c09-S05 Same source record as c01-S05.

c09-S07 Same source record as c01-S07.

c09-S12 Same source record as c04-S12.

c09-S14 Design and Development of a Precision Defect Detection System Based on a Line Scan Camera Using Deep Learning GRADE B

Applied Sciences (MDPI). 2024-12-23. https://www.mdpi.com/2076-3417/14/24/12054

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c09-S19 Same source record as c02-S19.

c10-S07 Same source record as c01-S07.

c10-S19 Same source record as c02-S19.

c10-S24 Same source record as c04-S24.

c10-S25 Quality control of garment product using DMAIC Six Sigma GRADE B

OPSI, metadata and abstract through GARUDA. 2024-06-30. https://garuda.kemdiktisaintek.go.id/documents/detail/5650084

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S26 Same source record as c04-S26.

c10-S1001 Critical Analysis of Material Consumption and Cost Reduction Techniques for the Apparel Cutting Processes GRADE A

Tanzania Journal of Science, University of Dar es Salaam. 2021-12. https://www.ajol.info/index.php/tjs/article/view/217586/205207

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1002 Décret n° 2023/00338/PM du 21 mars 2023 fixant le Salaire Minimum Interprofessionnel Garanti (SMIG) GRADE A

Services du Premier Ministre, decree text reproduced by Maathis. 2023-03-21. https://app.maathis.com/droit/cameroun/decret/2023-00338-decret-n-202300338pm-21-mars-2023-fixant-salaire-minimum-interprofess

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S07 Same source record as c01-S07.

c11-S24 Same source record as c04-S24.

c11-S25 Same source record as c10-S25.

c11-S26 Same source record as c04-S26.

c11-S1101 NIST/SEMATECH e-Handbook of Statistical Methods: Proportions Control Charts GRADE A

National Institute of Standards and Technology. 2026-04-22. https://www.itl.nist.gov/div898/handbook/pmc/section3/pmc332.htm

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S03 Fabric Defect Detection Using Computer Vision Techniques: A Comprehensive Review GRADE B

Mathematical Problems in Engineering. 2020-11-16. https://ideas.repec.org/a/hin/jnlmpe/8189403.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S07 Same source record as c01-S07.

c12-S34 Anomaly Detection in Metal-Textile Industries GRADE B

IntechOpen. 2025-01-03. https://www.intechopen.com/chapters/1206817

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S1201 Global case studies of social dialogue on AI and algorithmic management GRADE A

International Labour Organization. 2025-07. https://www.ilo.org/sites/default/files/2025-07/wp144_web.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S07 Same source record as c01-S07.

c13-S18 Same source record as c01-S18.

c13-S21 Same source record as c03-S21.

c13-S26 Same source record as c04-S26.

c13-S27 From Policy to Practice: Challenges of Environmental Data Generation and Digital Product Passport Readiness for Circular Garments Under EU Circular Textile Regulation GRADE B

Sustainability (MDPI). 2026-07-13. https://doi.org/10.3390/su18147127

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S28 A Sustainable Product Tracking System under the Combination of Textile Label Regulations and Digital Product Passport GRADE B

Textile & Leather Review. 2026-04-28. https://doi.org/10.31881/tlr.2026.1234

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S30 Sustainability and Circularity in the Textile Value Chain - A Global Roadmap GRADE A

United Nations Environment Programme (UNEP). 2023-05. https://doi.org/10.59117/20.500.11822/42580

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S32 Same source record as c04-S32.

c13-S34 Same source record as c06-S34.

c13-S1301 Regulation (EU) 2024/1781 establishing a framework for the setting of ecodesign requirements for sustainable products GRADE A

European Parliament and Council of the European Union. 2024-06-13. https://eur-lex.europa.eu/eli/reg/2024/1781/oj/eng

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c13-S1302 Ecodesign for Sustainable Products and Energy Labelling Working Plan 2025-2030, COM(2025) 187 final GRADE A

European Commission. 2025-04-16. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0187

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c14-S04 Same source record as c01-S04.

c14-S07 Same source record as c01-S07.

c14-S12 Same source record as c04-S12.

c14-S14 Same source record as c04-S14.

c14-S17 Same source record as c04-S17.

c14-S19 Same source record as c01-S19.

c14-S24 Same source record as c04-S24.

c14-S26 Same source record as c04-S26.

c14-S27 From Policy to Practice: Challenges of Environmental Data Generation and Digital Product Passport Readiness for Circular Garments Under EU Circular Textile Regulations GRADE B

Sustainability (MDPI). 2026-07-13. https://doi.org/10.3390/su18147127

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c14-S34 Same source record as c06-S34.

Stable definitions make comparisons useful.

Acceptable Quality Limit (AQL)

Defined under ISO 2859-1 as the worst tolerable process average in statistical batch sampling, commonly set at 0.0 for critical defects, 1.5 or 2.5 for major defects and 4.0 for minor defects in garment audits.

Active learning

A semi-supervised machine-learning workflow in which ambiguous or high-uncertainty camera frames are routed to a human inspector whose one-tap confirmation labels the sample for subsequent model retraining.

Adaptive Shape Convolution Module (ASCM)

A feature-fusion neck module used in YOLOv8-based fabric detectors (such as Fab-ASLKS) to adapt receptive fields to irregular and elongated textile defects.

AITEX Fabric Image Database

A public benchmark dataset developed by the Textile Research Institute in Spain containing high-resolution industrial fabric images across multiple weave defect classes.

Andon

A visual and audible shop-floor control signal that alerts line supervisors and mechanics immediately when a process or vision station crosses a defect threshold.

Area-scan camera

A two-dimensional matrix camera that captures a complete rectangular frame in a single exposure, suited to stationary flat-lay garment inspection and sewing-needle zone monitoring.

ASTM D5430 (Four-Point System)

The standard textile inspection method that assigns 1, 2, 3 or 4 penalty points according to defect length or hole presence and normalises total points per 100 square yards of fabric.

AUROC (Area Under the Receiver Operating Characteristic)

A threshold-independent performance metric measuring a model's ability to discriminate between defect-free and anomalous images or pixels from 0.0 to 1.0 (or 100 percent).

Brightfield illumination

Front-mounted coaxial or diffuse lighting directed at the fabric surface to reveal colour variations, oil stains, soil marks and print misregistration.

CIELAB Delta-E

A quantitative metric defined by the International Commission on Illumination measuring perceptual colour distance between a target textile shade standard and a sampled fabric roll.

Cold-start problem (industrial vision)

The operational constraint on a greenfield production line or new fabric style where zero annotated local defect images exist at launch, requiring pre-trained or one-class nominal models.

Contact Image Sensor (CIS)

A compact, 1:1 magnification linear optical sensor array combining LED illumination, rod-lens optics and CMOS detectors for sub-millimetre on-loom or short-working-distance web scanning.

Convolutional Autoencoder (CAE)

An unsupervised neural network trained to compress and reconstruct defect-free fabric textures so that unseen defects produce high pixel-level reconstruction error.

Coordinate Attention (CA)

An attention mechanism that embeds positional horizontal and vertical coordinate information into channel attention maps, improving small-target sewing and weave defect localisation.

Coreset subsampling

A greedy geometric reduction algorithm used in PatchCore to shrink a nominal patch-feature memory bank while preserving coverage of normal texture variations for fast nearest-neighbour search.

Cost of Poor Quality (COPQ)

The total financial loss incurred from internal failures (fabric scrap, unpicking, re-sewing, downgraded seconds, OEE loss) and external failures (AQL lot rejections, chargebacks, returns).

Cut-Make-Trim (CMT)

The core garment manufacturing stage encompassing marker spreading, panel cutting, sewing assembly, trimming, pressing and packing.

Darkfield grazing illumination

Low-angle lighting directed nearly parallel to the fabric plane so raised surface anomalies such as slubs, knots, nep clusters and 3D seam puckering cast high-contrast shadows.

Defects Per Hundred Units (DHU)

A primary garment-floor quality KPI calculated as total defects found divided by total units inspected multiplied by 100, capturing multiple defects occurring on a single garment.

Digital Product Passport (DPP)

A structured, machine-readable digital record mandated under Regulation (EU) 2024/1781 (ESPR) that links a physical product via QR, NFC or RFID to verifiable provenance, durability and sustainability data.

DMAIC (Define, Measure, Analyze, Improve, Control)

The five-step Six Sigma continuous-improvement cycle used with Pareto charts and fishbone diagrams to eliminate root causes of recurring sewing and fabric defects.

Edge inference

Executing trained neural-network forward passes locally on industrial compute modules attached directly to factory cameras rather than streaming raw video to a remote cloud server.

Failure Mode and Effect Analysis (FMEA)

A preventive quality engineering method that ranks potential process failures by multiplying Severity, Occurrence and Detection scores into a Risk Priority Number (RPN).

False negative (defect escape)

An inspection error in which a genuinely defective fabric segment or garment is misclassified as acceptable and passes downstream to cutting or the customer.

False positive (false alarm)

An inspection error in which a nominal, acceptable fabric region (or benign lint speck) is flagged as a defect, triggering unnecessary human triage or line stops.

First-Pass Yield (FPY) / Right First Time (RFT)

The percentage of garments or fabric rolls that pass through production and inspection complete and defect-free on the first attempt without rework or repair.

Fourier transform (2D FFT)

A spectral decomposition technique that converts spatial fabric images into the frequency domain, isolating periodic warp-and-weft weave structures from localized non-periodic defects.

Gabor filter bank

A set of orientation- and frequency-tuned linear filters inspired by the human visual cortex, widely used to extract multi-scale directional texture features from woven and knitted fabrics.

Greige fabric

Unbleached, undyed raw woven or knitted cloth taken directly off the loom or knitting machine prior to wet processing.

INT8 / FP16 quantisation

Model-compression techniques that convert 32-bit floating-point neural weights and activations into 16-bit floating-point or 8-bit integer representations to accelerate edge GPU/NPU inference.

Line-scan camera

An industrial camera containing a single row (or trilinear RGB rows) of pixel sensors that builds a continuous two-dimensional image line by line as fabric moves past on rollers.

Mean Average Precision (mAP)

The standard object-detection benchmark metric computing the mean of average precision integrals across all defect classes at a specified Intersection-over-Union threshold (such as mAP@0.5).

MVTec Anomaly Detection (MVTec AD)

The primary international computer-vision benchmark dataset for unsupervised industrial anomaly detection across 15 texture and object categories including carpet, grid and leather.

Non-Maximum Suppression (NMS)

A post-processing step in object detectors that eliminates redundant overlapping bounding boxes around a single detected defect.

Overall Equipment Effectiveness (OEE)

The foundational manufacturing productivity metric calculated as the product of Availability, Performance and Quality (Right First Time).

PatchCore / PaDiM

Leading one-class industrial anomaly detection frameworks that model normal texture patches using pre-trained CNN embeddings and memory banks or multivariate Gaussian distributions.

Rotary shaft encoder

An electro-mechanical sensor mounted on a fabric inspection roller that emits pulses proportional to linear web travel, synchronising line-scan camera acquisition to exact millimetre increments.

Seam puckering

An unacceptable wavy or wrinkled distortion along a sewn garment seam caused by thread tension imbalance, structural jamming, differential feed or fabric shrinkage.

Tianchi Fabric Defect Dataset

A large-scale annotated industrial textile dataset released through Alibaba Cloud Tianchi competitions, widely used to benchmark YOLO and R-CNN architectures across 20-plus defect types.

TILDA Textile Texture Database

A foundational benchmark dataset developed at the University of Freiburg containing representative textile weave groups and defined defect classes under controlled illumination.

Follow the terms back to their context.

Page references point to occurrences in chapter prose. When a term does not occur in a chapter, its glossary definition is indexed.

Acceptable Quality Limit (AQL) 237

Active learning 105 181 182

Adaptive Shape Convolution Module (ASCM) 237

AITEX Fabric Image Database 237

Andon 98 132 135 170 171 172 175 176

Area-scan camera 59 72

ASTM D5430 (Four-Point System) 237

AUROC (Area Under the Receiver Operating Characteristic) 237

Brightfield illumination 237

CIELAB Delta-E 96 130

Cold-start problem (industrial vision) 237

Contact Image Sensor (CIS) 237

Convolutional Autoencoder (CAE) 238

Coordinate Attention (CA) 238

Coreset subsampling 91 101

Cost of Poor Quality (COPQ) 238

Cut-Make-Trim (CMT) 238

Darkfield grazing illumination 238

Defects Per Hundred Units (DHU) 238

Digital Product Passport (DPP) 238

DMAIC (Define, Measure, Analyze, Improve, Control) 238

Edge inference 121

Failure Mode and Effect Analysis (FMEA) 238

False negative (defect escape) 238

False positive (false alarm) 238

First-Pass Yield (FPY) / Right First Time (RFT) 238

Fourier transform (2D FFT) 239

Gabor filter bank 76 93

Greige fabric 239

INT8 / FP16 quantisation 239

Line-scan camera 59 62 130 135

Mean Average Precision (mAP) 239

MVTec Anomaly Detection (MVTec AD) 239

Non-Maximum Suppression (NMS) 239

Overall Equipment Effectiveness (OEE) 239

PatchCore / PaDiM 239

Rotary shaft encoder 59 99

Seam puckering 33 51 65 66

Tianchi Fabric Defect Dataset 80

TILDA Textile Texture Database 239

A reproducible edition keeps its limits visible.

Production is the warp. Measurement is the weft.

Green Future Tech prepared this white paper from published sources opened during research, each recorded with its URL, publisher, date and a quoted sentence in the evidence register, and from GFT's Cameroon textile programme evidence register (R entries, transcribed from the programme's data annex and flagged where refreshed in 2026). Derived figures show their inputs; illustrative models are labelled with their basis. No figure is presented as audited, and nothing here is legal, tax or investment advice.

Source treatment preserves uncertainty.

Inherited records reproduce the supplied finding and attribution only where cited. Their source-text status is preserved. Original publication dates and exact source URLs are incomplete in that register. These records require independent validation before external reliance. New chapter-local sources carry a publication date, complete URL, publisher, exact supporting quotation and access date.

The exports share a content model.

The PDF and web edition use the same parsed chapter blocks and SVG exhibits. The build records input hashes, artifact hashes, layout measurements and elapsed time. Cross-references, source notes and the term index are generated from identifiers rather than manually typed page numbers.

IMPRINT

Issued by Green Future Tech Limited, Hong Kong SAR.
Engineering the sustainable production of tomorrow.

GFT Research. Edition 2026-10-30. Document wp03. Locally generated proof. No named author, office, certification or client endorsement is implied.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

The factory that makes the garment should make the evidence.

Engineering the sustainable production of tomorrow.