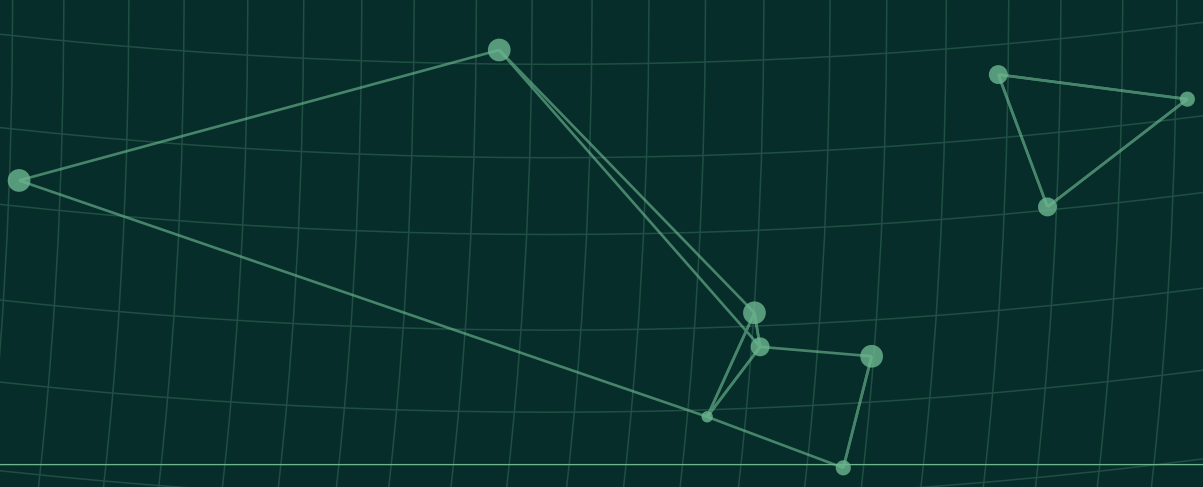
GREEN FUTURE TECH

LIMITED / HONG KONG SAR

GFT RESEARCH / WARP AND SIGNAL

From Mood Board to Specification

Generative AI in apparel design and what a greenfield maker should adopt



wp10
RESEARCH EDITION

EDITION 2026-10-30
HONG KONG SAR

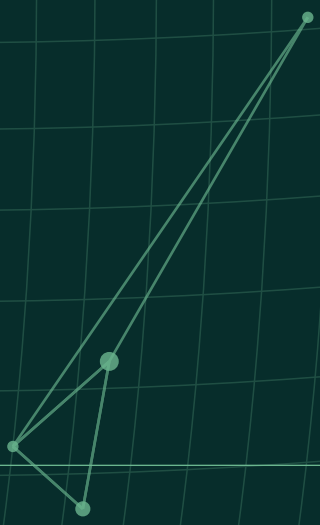


THE GFT THESIS

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	wp10 / From Mood Board to Specification
EDITION DATE	2026-10-30
AUDIENCE	Greenfield apparel sponsors and plant leadership, Designers, pattern technicians and product-development leads, Quality, procurement and digital-system teams, Technical partners assessing a bounded generative-design pilot
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: R03, R06, R07, R10, R12, R14, R15, R16, R17, R27, R28, R29, R30, R31. Temporal eligibility is unverified; these are inherited source-text checks, not independent source validation.

The operating case runs from fibre to feedback.

01 / A convincing image cannot authorise a cutting order 14

1.1 / Local value requires producible garments 15

1.3 / A design image omits the instructions a factory needs 17

1.5 / The greenfield programme defines the adoption sequence 19

1.7 / Research prototypes and released products answer different questions 20

1.9 / The paper owns the handoff from creative intent to approved specification 23

1.11 / The decision unit is an accepted style rather than a generated picture 24

1.13 / The evidence base supports conditional adoption 26

02 / Structured garment research predates the recent agent wave 30

2.1 / The historical window is a baseline rather than a market census 31

2.2 / Text guidance reached explicit sewing representations 33

2.3 / Multimodal models widened the interface to pattern systems 34

2.4 / Materials and fit became explicit research targets 36

2.5 / Structure and agentic refinement address different failure modes 38

2.6 / Commercial documentation establishes a narrower fact than factory success 39

2.7 / A claim ledger prevents research results becoming sales promises 41

03 / Controlled ideation starts with a manufacturable brief 44

3.1 / The brief must specify what can change 45

3.2 / References need permission and design rationale 48

3.3 / Silhouette and local detail require separate review 49

3.4 / Prompt iteration needs a record of decisions 51

3.5 / Material appearance is not material identity 52

3.6 / A designer must translate selected concepts into technical intent 54

3.7 / Exploration should be bounded by the cost of review 56

04 / Editable patterns require geometry and seam relationships 59

4.1 / Panel geometry is a different object from a garment image 60

4.4 / Programs and parameter files expose different edit controls 62

4.7 / Parameter constraints should be explicit and testable 65

4.10 / Template reuse and template-free research solve different needs 68

4.13 / Synthetic training data carries design assumptions 69

4.16 / Generated code needs a constrained execution path 72

4.19 / The output must remain editable by the technician 74

05 / A high seam score can coexist with failed garment assembly 77

5.1 / Seam correspondence determines how panels become a garment 78

5.2 / Edge-level F1 is not a whole-garment acceptance rate 79

5.3 / Many-to-one relationships require construction intent 81

5.4 / Geometric validity does not settle the sewing method 82

5.5 / Assembly review needs an explainable failure taxonomy 85

5.6 / Layering increases the need for inspection 87

5.7 / A seam gate should return evidence with the rejection 88

06 / Fabric physics must be checked against the fabric being bought 91

6.1 / A texture map does not describe mechanical behaviour 92

6.4 / Material inference is a hypothesis that needs validation 94

6.7 / Material assets need measurement provenance 95

6.10 / Direction and fabric variation belong in the model assumptions 97

6.13 / Simulation failures must remain visible 99

6.16 / Physical checks must calibrate the intended use 101

6.19 / A changed fabric can reopen a previously approved design 103

07 / A well-fitted image is not evidence that the garment fits 107

7.1 / Appearance-based try-on can hide a size mismatch 108

7.5 / A target size range is a product decision 110

7.9 / Ease should be a deliberate specification 111

7.13 / Grading requires checks beyond the base size 114

7.16 / Pose and intended activity change the acceptance question 116

7.19 / Wearer data needs its own permission boundary 117

7.22 / Physical fit approval remains a release condition 118

08 / Surface generation needs print geometry and cultural permission 121

8.1 / The word pattern describes two different objects 122

8.4 / Appearance control is useful within a defined artwork brief 124

8.7 / Texture coherence differs from manufacturing repeatability 125

8.10 / Colour decisions need a controlled physical comparison 128

8.13 / A cultural reference is not a rights-free style token 129

8.16 / AI asset provenance and cotton provenance need different records 131

8.19 / Publication and production require separate permissions 133

09 / A tech pack is complete only when the receiving team can use it 137

9.1 / The tech pack should be a controlled information contract 138

9.2 / Structured text can assist drafting but cannot invent missing inputs 140

9.3 / CAD exchange needs an explicit format and unit agreement 141

9.4 / Round-trip verification must inspect dimensions and annotations 142

9.5 / BOM links should resolve to approved physical inputs 144

9.6 / Change control must travel with the specification 145

9.7 / Release authority must be explicit 147

10 / Technicians must own the learning loop 150

10.1 / Skills must be described as tasks rather than software brands 151

10.4 / Recorded operations can make expert corrections teachable 153

10.7 / Human-AI ideation should preserve design authorship decisions 157

10.10 / A training exercise needs observable acceptance criteria 158

10.12 / The interface should show uncertainty and correction paths 160

10.15 / Role separation supports review without inventing headcount 161

10.18 / Local evaluation must test transfer rather than assume it 163

11 / Rights and model risks need versioned records 167

11.1 / A permission to use software differs from rights in an output 168

11.2 / Human contribution should be recorded without promising protection 170

11.3 / Sensitive data needs a declared destination 173

11.4 / Community engagement cannot be replaced by similarity screening 174

11.5 / Model changes require a regression decision 175
11.6 / Vendor due diligence must produce reviewable artefacts 177
11.7 / Incidents should lead to containment and correction 179

12 / Economics must count accepted styles and retained physical checks 182

12.1 / The counterfactual must describe the same accepted product 183
12.2 / A complete cost model includes rejected outputs 184
12.3 / A transparent example can test the arithmetic without predicting savings 187
12.4 / Breakeven depends on usable volume and actual avoidance 189
12.5 / A faster task is not automatically a shorter calendar 190
12.6 / Environmental claims need boundaries and physical data 191
12.7 / The investment decision should survive a downside case 193

13 / A local pilot must test the complete specification handoff 195

13.1 / The procurement unit should be a usable workflow 196
13.2 / The first pilot needs a deliberately narrow product envelope 197
13.3 / A matched baseline makes the test interpretable 199
13.4 / Hard release conditions must be set before results are seen 201
13.5 / Connectivity failures must preserve the approved work 203
13.6 / Buyers should retain exit and data-portability options 204
13.7 / The pilot report must distinguish observed results from pending work 206

14 / GFT should scale measured readiness rather than generated output 209

14.1 / The programme should retain its staged sequence 210
14.2 / Setup should build reusable records and decisions 212
14.3 / Launch should establish the operational baseline 213
14.4 / Scale should prepare a bounded partner evaluation 214
14.5 / Optimisation should expand only within the tested envelope 216
14.6 / GFT should measure approved outcomes and correction effort 218
14.7 / The evidence agenda must remain open 221

Executive summary 12
Evidence register 222
Glossary 239
Term index 242
Method and imprint 244

Every exhibit states a testable claim.

EXHIBIT 01 / stacked
The cotton gap explains the industrial problem, not an AI return 16

EXHIBIT 02 / flow
Release requires evidence at every handoff 19

EXHIBIT 03 / matrix
Appearance and production readiness are separate axes 22

EXHIBIT 04 / table
The running example has assumptions rather than customers 26

EXHIBIT 05 / timeline
Explicit sewing representations appear throughout the research window 32

EXHIBIT 06 / table
A model comparison needs common tasks before common scores 35

EXHIBIT 07 / flow
Version dates change the meaning of a chronology 36

EXHIBIT 08 / matrix
Public availability and local suitability require separate checks 40

EXHIBIT 09 / table
A constrained brief is more useful than a list of style adjectives 47

EXHIBIT 10 / matrix
A visual review cannot clear a physical property 50

EXHIBIT 11 / table
The model study reports a task-specific recognition gap 53

EXHIBIT 12 / flow
Rejected concepts should return to the relevant decision owner 55

EXHIBIT 13 / table
A usable pattern record carries more than coordinates 62

EXHIBIT 14 / matrix
Representation choice should follow the required edit 64

EXHIBIT 15 / small-multiples
The same concept can conceal different geometry failures 67

EXHIBIT 16 / waterfall
Quality filtering in GarmentCodeData removes about 36 percent of simulated drapes per batch 70

EXHIBIT 17 / layered-architecture
A deterministic boundary belongs between generation and CAD 73

EXHIBIT 18 / keyfigures
AutoSew reports distinct relation and assembly metrics 80

EXHIBIT 19 / table
Identical relation F1 can produce different assembly outcomes 81

EXHIBIT 20 / table
Digital correspondence omits physical construction choices 84

EXHIBIT 21 / matrix
Assembly errors need categories before rates 86

EXHIBIT 22 / flow
A failed seam must return to an identified correction point 89

EXHIBIT 23 / layered-architecture
Physical and visual material records need a common identity 94

EXHIBIT 24 / table
Every simulation input should disclose its basis 97

EXHIBIT 25 / matrix
A material check needs both an avatar and a physical swatch 102

EXHIBIT 26 / flow
Material substitution is a controlled change 104

EXHIBIT 27 / matrix
Fit assessment needs matched and deliberately mismatched cases 109

EXHIBIT 28 / layered-architecture
Body measurements and ease connect intent to evaluation 113

EXHIBIT 29 / table
A grading review must survive export 115

EXHIBIT 30 / flow
A fit rejection should preserve the cause 119

EXHIBIT 31 / table
Surface artwork and sewing geometry need different fields 123

EXHIBIT 32 / small-multiples
A digital texture can fail at a physical seam 126

EXHIBIT 33 / flow
Cultural permission precedes commercial reuse 130

EXHIBIT 34 / matrix
Asset approvals must state the permitted destination 134

EXHIBIT 35 / table
A specification field must point to its source and owner 139

EXHIBIT 36 / flow
CAD readback closes the exchange loop 142

EXHIBIT 37 / table
Format support is narrower than tested interoperability 143

EXHIBIT 38 / matrix
A change can reopen several approvals 146

EXHIBIT 39 / matrix
Roles should follow the decision that can fail 153

EXHIBIT 40 / bar
Expert operation traces outperform static textbook rules on complex pattern revisions 155

EXHIBIT 41 / flow
An expert correction should become a reusable lesson 156

EXHIBIT 42 / table
A training assessment needs unseen tasks 159

EXHIBIT 43 / weave-matrix
Time and quality must be reviewed together 165

EXHIBIT 44 / matrix
A model contract answers only part of the rights question 169

EXHIBIT 45 / layered-architecture
The design record needs a provenance envelope 172

EXHIBIT 46 / flow
A model update should reopen the relevant test gate 177

EXHIBIT 47 / table
Procurement should demand an inspectable evidence pack 178

EXHIBIT 48 / table
An economic case needs measured denominators 186

EXHIBIT 49 / waterfall
Illustrative net saving is smaller than gross sample avoidance 188

EXHIBIT 50 / line
Illustrative breakeven requires enough accepted styles 189

EXHIBIT 51 / matrix
Environmental and financial claims need separate evidence 192

EXHIBIT 52 / matrix
A pilot needs both routine and held-out cases 198

EXHIBIT 53 / flow
Expansion requires every critical handoff to pass 202

EXHIBIT 54 / table
A vendor exit test starts with retained approved artefacts 205

EXHIBIT 55 / gantt
Pilot work should be sequenced by dependencies 207

EXHIBIT 56 / gantt
The roadmap is gated by readiness rather than a promised date 211

EXHIBIT 57 / matrix
A stop decision is a valid pilot outcome 217

EXHIBIT 58 / table
Every proposed KPI needs a numerator, denominator and owner 218

EXHIBIT 59 / flow
The approved style should feed the same measurement spine 220

EXECUTIVE SUMMARY

Build the factory and its evidence as one system.

Generative AI can broaden apparel design exploration, but a convincing image is not an instruction to cut fabric. This paper examines research published from 2023 to October 2026 and asks what a greenfield maker should adopt, under what conditions, and with whose approval. It distinguishes image ideation, structured sewing-pattern generation, material estimation, fit-aware evaluation and production release. Parametric garment research existed in 2023; later multimodal and agentic systems extend the available methods without establishing universal factory readiness. The paper argues for an accountable handoff from a constrained brief to editable geometry, measured material inputs, fit checks, a controlled tech pack and receiving-CAD readback. It separates research results from vendor documentation and from proposed GFT acceptance conditions. A running cotton-shirt example makes the decisions concrete without inventing an engagement, buyer or operating result. Rights review, cultural engagement, technician development and model-version control form part of the workflow rather than an afterthought. An explicitly illustrative sampling-cost model shows how review effort and retained physical checks can change the economic case. The closing roadmap aligns with GFT's Cameroon textile programme: establish records and operating measurement first, then qualify a bounded partnered generative-specification pilot. Adoption remains conditional on local evidence. Nothing in the paper constitutes investment, legal or tax advice.

01 The useful adoption unit is an accepted specification, not the number of generated images. [S09] [S22] [S19]

02 Structured sewing-pattern research already existed in 2023, so a simple image-only-to-pattern chronology is misleading. [S02] [S03]

03 Programs, pattern tokens, parameter JSON and schema-aware representations expose different technical control points. [S08] [S09] [S10] [S14] [S16]

04 A strong seam-relation metric is compatible with a substantially lower whole-garment digital assembly result. [S12]

05 Visual fabric generation and material-property inference do not replace physical material evidence. [S01] [S11] [S28] [S32]

06 Appearance-based virtual try-on cannot be assumed to establish size-correct fit. [S13] [S27]

07 Documented DXF support does not remove the need to verify units, grading, notches and allowances in the receiving workflow. [S22]

08 Commercial tool permission, human contribution and cultural permission answer distinct questions. [S31] [S21] [S23] [S24]

09 Economic and environmental benefits remain unproven until accepted outcomes and all changed costs or resource flows are measured. [S19] [S29] [S25]

10 GFT should preserve its staged programme and treat generative specification as a bounded partner pilot with stop conditions. [S19] [S20]

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

01

PART I: PRODUCTION READINESS IS THE TEST / THE PRODUCTION GAP

A convincing image cannot authorise a cutting order

GFT should buy creative assistance selectively and release garments only through an accountable specification process.

CHAPTER 01
DECISION TEST

GFT should buy creative assistance selectively and release garments only through an accountable specification process.

A greenfield apparel manufacturer should adopt generative artificial intelligence for design only where it accelerates the production of an inspected, technician-approved specification dossier, never as a direct instruction to cut cloth. [S09] Across research published between 2023 and October 2026, neural systems have progressed from unconditional appearance synthesis to controlled image generation, parametric pattern-making programs, multimodal tokenisers, stitch-prediction graphs and simulation-coupled refinement. [S02] [S12] [S26] Yet a realistic garment picture and a manufacturable cutting order remain different engineering objects. A raster image suggests silhouette, colour placement and surface mood, whereas a cutting room requires closed two-dimensional panel boundaries, explicit units, grainlines, graded size tables, seam allowances, notch coordinates, bill-of-materials links and named technical sign-off. [S09] [S22] Green Future Tech Limited (GFT), Hong Kong SAR, recommends a strict separation between upstream creative exploration and downstream release: visual tools may widen the candidate pool under a constrained brief, but every style must pass deterministic geometry checks, physical material verification, size-aware fit review and receiving-system readback before fabric is cut. [S19] [S22]

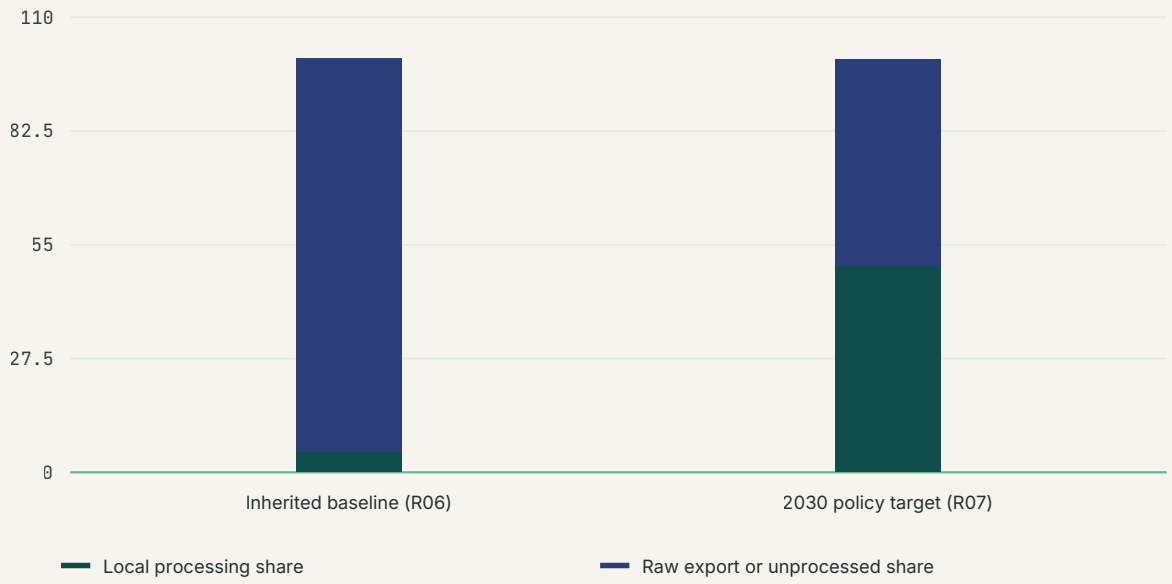
1.1

Local value requires producible garments

Industrialisation in a raw-material economy depends on converting agricultural fibre into repeatable, specification-compliant finished goods rather than exporting unprocessed lint. In the inherited evidence register for Cameroon, approximately 95 percent of domestic cotton fibre is exported raw while about 5 percent is processed locally, alongside a stated industrialisation benchmark of 50 percent domestic transformation. [R06] The national development strategy (SND30) restates that 50 percent local-processing target by 2030. [R07] Within this white paper and companion paper wp01, those proportions serve as sector context from the inherited register (where original publication dates were not recorded and verification rests on inherited source-text review) rather than independently audited national accounts or GFT operating measurements. [R06] [R07] Even with that qualification kept visible, the gap between raw fibre export and domestic garment manufacture defines the engineering problem a greenfield producer must solve: value is captured locally only when a plant translates cotton yarn and woven cloth into finished apparel meeting buyer tolerances on the first run. [R06] [S09]

EXHIBIT 01 / PERCENT OF COTTON FIBRE

The cotton gap explains the industrial problem, not an AI return



SOURCED [R06] [R07] Approximate inherited sector baseline from Cameroon Trade Hub (R06, publication date unrecorded) and stated 2030 national policy target from SND30 via Ecofin Agency (R07, publication date unrecorded). The 2030 column represents a national policy target rather than an achieved operating share or a forecast return from design software.

As shown in **Exhibit 1**, p. 16, closing the distance between a 5 percent local-processing baseline and a 50 percent national policy target is a physical manufacturing challenge before it is a software challenge. [R06] [R07] The exhibit establishes the scale of the stated industrial ambition in Cameroon, but it cannot establish a market demand forecast, a guaranteed return on digital investment, or proof that generative design tools increase factory utilisation. [R06] A garment plant that generates thousands of digital concept boards while failing to issue accurate cutting specifications still exports raw fibre or accumulates unsellable offcuts. Local value creation requires garments that can be cut, sewn, pressed, inspected and shipped to specification. [S09]

1.2 CREATIVE VOLUME HAS NO VALUE WITHOUT SPECIFICATION DISCIPLINE

In early commercial enthusiasm for generative tools, software vendors and sector commentaries measured success by ideation speed: how many mood boards, colourways or silhouette variations a designer could render in an afternoon. [R30] Consultancy estimates published in 2025 suggested that scaling artificial intelligence across fashion operations could reduce inventory costs by 35 to 40 percent, raise forecast accuracy by 20 to 30 percent, reduce lost sales by up to 65 percent and contribute US$150 to 275 billion in sector operating profit over three to five years, with 75 percent of surveyed executives prioritising deployment. [R27] Those broad sector figures aggregate merchandising, supply-chain forecasting, marketing and retail personalisation across established global brands; they do not isolate pattern engineering in a greenfield cut-make-trim plant, and GFT does not import them as an expected return. [R27] [S19]

GFT's position is stricter. In a greenfield apparel operation, artificial-intelligence-assisted concepts earn their place only if they reduce approved-specification errors, shorten the time required to reach a complete technical package, or reduce physical sampling rework without degrading fit or material compliance. [S19] [S20] This statement is a recommendation derived from GFT's Cameroon textile programme (feasibility study, performance-management system design and digital strategy, prepared 31 January 2024, with sources refreshed in 2026); it is not a claim that GFT operates a garment factory or has measured live production savings. [S19] Where an image generator multiplies unvetted proposals faster than pattern technicians can check geometry, it increases queue congestion rather than plant throughput. [S09] [S20]

1.3

A design image omits the instructions a factory needs

To understand why a rendered picture cannot authorise production, an engineer must compare what an appearance model controls with what a cutting table and sewing line consume. Fine-grained image synthesis frameworks such as IMAGGarment allow a user to condition two-dimensional garment generation on silhouette, colour palette and logo placement simultaneously. [S26] That capability is useful during early concept selection, because a design lead can inspect proportions and branding placement before commissioning pattern work. [S26] However, a pixel grid contains no explicit representation of flat panels, seam pairings, internal darts, ease allowances, fabric mass per unit area, or grading rules across sizes. [S09] [S26] [S27]

By contrast, structured garment-generation research explicitly models the intermediate geometry required for fabrication. GarmentCode formulates sewing patterns as modular, component-based parametric programs bridging construction goals and valid sewing patterns for simulation and fabrication. [S02] Design2GarmentCode uses large multimodal models to synthesise parametric pattern-making programs from images, text descriptions or sketches, treating sewing patterns as blueprints for cutting and tailoring. [S09] ChatGarment uses vision-language models to emit JSON descriptions and continuous numeric parameters driving a parametric sewing-pattern program capable of three-dimensional body draping. [S10] Architectures from April 2025, including GarmentDiffusion and GarmageNet, encode two-dimensional panel boundaries and sewing connections as compact edge tokens or structured geometry images to accelerate multimodal pattern synthesis. [S0101] [S0102]

1.4 THE MINIMUM RELEASE DOSSIER IS AN ENGINEERING CONTRACT

Even when a research model outputs vector panels or parametric code instead of pixels, a raw geometric prediction is only one component of an authorised factory instruction. [S12] [S22] Drawing on the distinction between visual ideation, parametric geometry and industrial CAD exchange, GFT proposes a seven-part minimum release dossier before any style moves to sample or bulk cutting: [S09] [S22]

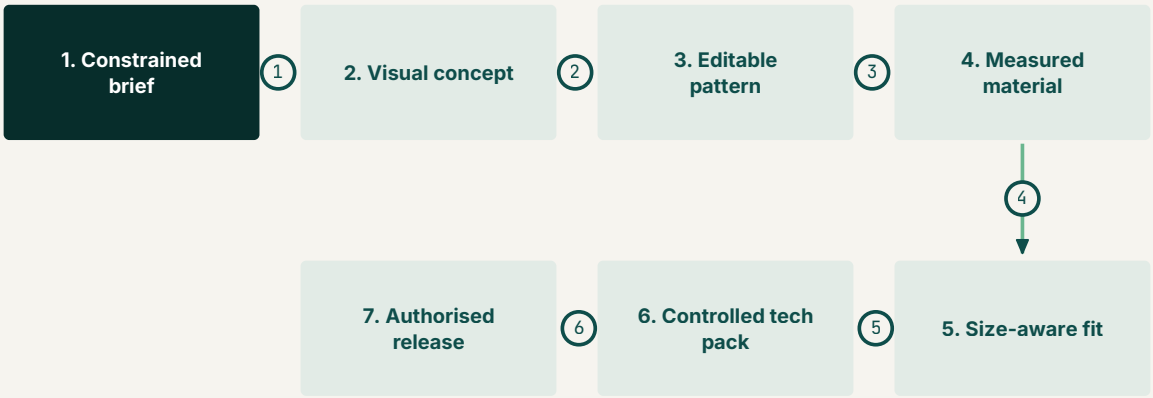
1. **Validated panel geometry:** Closed two-dimensional boundary curves for every component piece, with explicit units (millimetres or centimetres), cut quantities, mirror-pair rules, grainlines and internal drill locations. [S02] [S22]

2. **Graded size table and points of measure:** Explicit garment dimensions across all target sizes, specifying landmarks for each point of measure and separating body dimensions from ease allowances. [S13] [S22] [S27]
3. **Dimensional and construction tolerances:** Permitted plus-or-minus deviations at each point of measure after sewing and pressing, agreed with quality leadership before sample evaluation. [S19]
4. **Material and trim identifiers:** Traceable references linking each panel and fastener to a physical swatch record, measured fabric weight and thickness, weave direction and approved supplier lot. [S01] [S20]
5. **Seam and assembly instructions:** Edge-to-edge correspondence pairs, sewing direction, easing ratios, seam allowance widths, stitch class, needle and thread specifications, and reinforcement bartacks. [S12] [S22]
6. **Version and provenance log:** A revision history recording the originating brief, reference permissions, model version, human technician edits, and superseded specification identifiers. [S20]
7. **Named release sign-off:** Timestamped approvals from the pattern technician, quality reviewer and release authority, separating draft concepts from sample-authorised and bulk-authorised states. [S19] [S20]

> **GFT NOTE**
>
> GFT presents this seven-part release dossier as a proposed minimum governance contract for a greenfield maker, not as a universal statutory standard or a feature automated by current generative software. [S19] [S22] No candidate style should bypass any of the seven records merely because its concept render or three-dimensional drape looks convincing on screen. [S13] [S26]

EXHIBIT 02 / WORKFLOW GATES

Release requires evidence at every handoff



1. Scope and reference rights check
2. Closed panel and seam topology check
3. Physical swatch and property match
4. Explicit ease and sample fit check
5. DXF readback and BOM verification
6. Technician and quality release sign-off

ILLUSTRATIVE [S09] [S22] Illustrative GFT proposed workflow informed by programmatic pattern synthesis (S09), CAD exchange documentation (S22) and the internal digital strategy (prepared 31 January 2024). Any failed verification returns the style to the upstream owner rather than permitting fabric cutting.

The sequence in **Exhibit 2**, p. 19 illustrates how GFT structures the transition from creative intent to an authorised cutting instruction. [S09] [S22] Each connector is an inspectable evidence gate rather than an automatic file conversion. If a concept generated at Stage 2 implies an unsewable armhole or an unavailable weave, the workflow halts at the geometry or material gate and returns the candidate to the designer or procurement lead. Because **Exhibit 2**, p. 19 is an illustrative process design rather than an empirical log, it cannot establish gate failure frequencies or software selection; those parameters must be measured during a pilot. [S19] [S20]

1.5

The greenfield programme defines the adoption sequence

Technology selection in apparel manufacturing fails when advanced software is layered onto an unmeasured shop floor. GFT's Cameroon textile programme, prepared on 31 January 2024 with source refreshes retained in 2026, establishes an explicit operational and digital sequence across its feasibility study, performance-management system design and digital strategy. [R06] [S19] In the programme's base case, the initial footprint concentrates on garment manufacturing: cut, make, trim (CMT), quality assurance, pressing and packaging, while spinning, weaving, dyeing

and wet finishing remain optional later or partnered upstream stages once core assembly is stable. [R06] Nothing in this programme description implies that GFT has built a plant, installed a line or contracted a supplier; it is an internal engineering design basis that governs how digital tools are sequenced. [S19]

Under that staged architecture, generative design is excluded from the foundational launch phase. A new facility must first establish basic operational records: style identifiers, approved base blocks, physical swatch registries, operator task standards, shift output logs and defect tallies. [S19] [S20] In the performance-management design, every digital intervention relies on a common measurement spine recording what was planned, what material batch entered the line, what was produced and where rework occurred. [S19] Introducing neural pattern synthesis before the team maintains versioned base blocks and size tables would leave technicians without a baseline against which to judge generated outputs. [S02] [S30]

1.6 OPERATIONAL CAPTURE AND INSPECTION QUALIFICATION PRECEDE DESIGN PILOTS

The internal digital strategy sequences capabilities across four phases: record setup, operational launch and baseline capture, bounded scaling with qualified inspection and planning tools, and later optimisation through partnered pilots. [S19] [S20] Factory-floor data capture and local qualification of fabric and garment inspection take priority over generative design experiments. The inherited evidence register notes industry literature reporting manual textile defect inspection accuracy of 60 to 70 percent (missing 20 to 30 percent of defects), compared with reported computer-vision accuracy above 90 percent across more than 40 defect types at line-scan speeds up to 60 metres a minute, alongside a global fabric-defect-detection AI market projected from US$523.4 million in 2024 to US$6,522.1 million by 2033. [R28] [R29]

Even for fabric inspection, GFT rejects the unsourced inference that mounting a camera above a frame automatically delivers production-ready quality control. [R29] [S19] Lighting stability, local weave structures, cotton nep density, escalation rules and false-alarm rates must be qualified on actual factory cloth before any vision system governs lot acceptance, a subject examined in companion paper wp03. [R29] The same discipline applies to generative apparel design. Whereas the inherited register observed that 2023 to 2024 apparel AI activity centred mainly on image generation across style, silhouette, colour and surface pattern, moving from mood boards to cut-ready specifications touches every downstream department in a CMT plant. [R30] [S09] Consequently, GFT places generative specification assistance in Phase 3 of its roadmap as a bounded, partner-led pilot that proceeds only after foundational CAD records, swatch registries and manual specification baselines operate reliably. [S19] [S20]

1.7

Research prototypes and released products answer different questions

A recurring hazard in technical procurement is treating a conference paper or vendor help page as proof that a workflow will succeed on a specific factory floor. Between 2023 and October 2026, academic groups demonstrated substantial advances in translating text prompts, sketches and photographs into digital garment assets. [S09]

[S10] [S17] EasyFashion (submitted September 2026, revised October 2026) integrates structured design specifications, three-dimensional avatar evaluation and automated sewing-pattern generation in a human-AI co-creation workflow, and its abstract reports a real-world production case. [S17] That reported case shows researchers testing physical fabrication rather than stopping at renders, yet a single academic demonstration does not establish industrial repeatability across graded size runs, diverse cotton weaves or buyer audits. [S17] [S19]

Other benchmarks illustrate why partial metrics must not be confused with end-to-end production readiness. AutoSew (February 2026) applies graph neural networks to predict stitching relationships purely from two-dimensional panel geometry, reporting a 96 percent F1-score on edge-level stitch prediction alongside 73.3 percent error-free digital assembly across test garments. [S12] As Chapter 5 analyses, a 96 percent relation-level score still leaves more than one-quarter of test garments with at least one digital assembly error, and simulated assembly is narrower than physical sewing on an industrial line. [S12] Similarly, the FIT dataset paper (April 2026) shows that standard virtual try-on models default to synthesising visually well-fitted images regardless of garment or wearer dimensions, concealing size mismatches. [S13] And EASE (June 2026) shows that local material allowance relative to the body must be parameterized explicitly rather than inferred passively from mesh deformation. [S27]

Commercial software documentation answers a different, equally bounded question. Support documentation from CLO Virtual Fashion (updated July 2026) confirms that its two-dimensional pattern module imports and exports Standard DXF, AAMA and ASTM files, with explicit controls for units, grading rules, notches, seam allowances and cutting lines. [S22] That documentation confirms supported exchange options, not proof that an AI-generated pattern exported as AAMA DXF will reopen in a factory's marker-making software without spline distortion or lost notch orientation. [S22]

1.8 FOUR EVIDENCE LABELS GOVERN EVERY CLAIM IN THIS PAPER

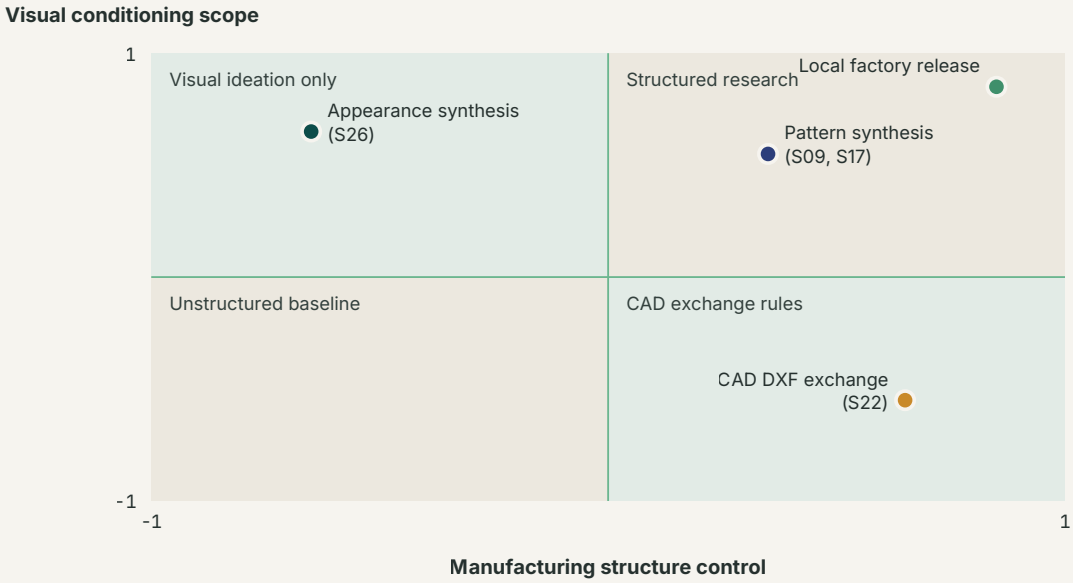
To prevent category errors between academic benchmarks, vendor manuals, internal designs and verified factory outcomes, this paper applies the contextual evaluation logic of the NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0) and its Generative AI Profile. [S19] [S20] Every statement across the fourteen chapters carries one of four explicit evidence labels: [S19] [S20]

- **Research result:** A mechanism, dataset property or benchmark reported in an opened academic paper or preprint (S01 to S18, S26 to S28, S30, S0101 and S0102), valid only for the authors' task, dataset, model version and protocol. [S09] [S12] [S30]
- **Vendor-documented capability:** A format, setting, commercial-use condition or asset requirement stated in official product documentation (such as S22), establishing that an option exists in the documented release without certifying plant fitness. [S22]
- **Illustrative GFT proposal:** An engineering workflow, release gate, responsibility matrix, key performance indicator or scenario calculation designed by GFT to govern adoption, marked as a proposal or illustrative model rather than an observed result. [S19] [S20]

- **Locally accepted result:** A specification workflow, CAD round trip, physical sample outcome or cost saving verified on the maker's own fabrics, blocks, machinery and buyer tolerances. As of 30 October 2026, this fourth category has zero demonstrated examples in the research base. [S19] [S20]

EXHIBIT 03 / CONCEPTUAL READINESS QUADRANT

Appearance and production readiness are separate axes



ILLUSTRATIVE [S09] [S17] [S22] [S26] Illustrative conceptual placement on two categorical axes; coordinates (-1 to 1) are layout positions to separate the four evidence categories rather than numerical maturity scores. Categories contrast controllable appearance synthesis (S26), programmatic pattern research (S09) and co-creation case studies (S17), documented CAD exchange options (S22), and locally accepted factory release (zero demonstrated examples in the research base).

Exhibit 3, p. 22 separates visual conditioning scope on the vertical axis from explicit manufacturing structure on the horizontal axis. [S09] [S22] [S26] Coordinates in the chart are illustrative layout positions within the four quadrants rather than numeric maturity scores. Controllable appearance models such as IMAGGarment occupy the upper-left quadrant because they offer visual conditioning over silhouette, colour and logo placement while omitting pattern geometry. [S26] Programmatic and co-creation systems such as Design2GarmentCode and EasyFashion move into the upper-right quadrant by linking multimodal prompts to parametric panels and seams, yet they remain research prototypes rather than qualified factory systems. [S09] [S17] Documented DXF exchange in commercial CAD occupies the lower-right quadrant because it governs deterministic geometry transfer without generating concepts. [S22] The top-right position representing a locally accepted generative specification workflow remains an unoccupied target awaiting pilot evidence. [S19]

1.9

The paper owns the handoff from creative intent to approved specification

Within GFT's twenty-paper research library, each white paper owns a distinct system boundary so that engineering controls are defined once and referenced consistently. This tenth paper (wp10) owns the technical and governance handoff from an initial creative brief to an approved, CAD-verified specification ready for sample or bulk release. [S09] [S22] It defines the information records and rejection return paths connecting product development to adjacent disciplines without duplicating their models. [S19] [S20]

Specifically, demand sensing and style-count projections belong to companion paper wp04; wp10 consumes a bounded style brief from merchandising without treating a forecast as permission to release unverified patterns. [R27] [S19] Optical inspection on fabric and sewing lines belongs to wp03; wp10 addresses pre-production digital assembly and specification verification rather than inline camera hardware. [R29] Physical cotton chain of custody belongs to wp05, whereas wp10 defines provenance logs for reference images, artwork and pattern revisions. [S20] Wet-processing resource intensity belongs to wp06, and plant power and network resilience belong to wp07; wp10 requires that approved patterns and tech packs remain locally accessible when cloud endpoints are unreachable. [S20] [S22]

Similarly, plant-wide workforce strategy belongs to wp09, while Chapter 10 specifies the pattern-engineering competencies needed to audit AI-assisted outputs. [S09] Waste valorisation belongs to wp11, capital structuring to wp13, enterprise data governance and model operations to wp14, product carbon accounting to wp16, regulatory and trade duties to wp17, plant-wide connected architecture to wp18, and three-dimensional digital sampling and print execution to wp20. [S19] [S20] Across these interfaces, wp10 supplies the acceptance gate: a digital sample in wp20 or a unit-cost estimate in wp13 counts as valid only when the specification has passed the geometry, seam, material, fit, rights and readback checks defined here. [S09] [S22]

1.10 EXPLICIT RETURN PATHS PREVENT DEFECTIVE SPECIFICATIONS ESCAPING DOWNSTREAM

A handoff contract is incomplete if it only describes the forward path when every check passes. In a development studio, generative outputs fail in specific ways: a prompt yields a collar that cannot attach to the neckband, a generated motif lacks community permission or repeat continuity, a predicted fabric preset stretches more in simulation than warehouse cotton cloth, or an exported DXF file loses its seam-allowance layer in the cutting room's CAD software. [S01] [S12] [S22]

To keep failures contained, GFT defines explicit return paths at every stage of the wp10 workflow. An ambiguous silhouette or unapproved reference returns from the concept gate to the brief owner (Chapter 3). [S26] Open panel contours, self-intersecting splines or wrong unit scales return from the geometry validator to pattern representation (Chapter 4). [S02] [S30] Missing stitch pairs or reversed sewing directions return from the assembly gate to topology editing (Chapter 5). [S12] Unmeasured fabric assumptions return to swatch testing and procurement (Chapter 6), while size-blind virtual try-on renders or insufficient movement ease return to grading review

(Chapter 7). [S13] [S27] Print scale errors or unresolved cultural permissions return to artwork governance (Chapter 8), and failed DXF readbacks return to tech-pack compilation before any cutting ticket is printed (Chapter 9). [S22] Planning tools in the GFT suite, including `use-case-portfolio`, `risk-matrix` and `sampling-cost`, help teams structure these gates, though neither those tools nor the proposed workflow have been field-proven in an operating factory. [S19]

1.11

The decision unit is an accepted style rather than a generated picture

When management evaluates whether creative assistance improves product development, the denominator determines whether the evaluation is honest. Counting generated images, prompt completions or raw meshes rewards volume regardless of manufacturability. [S19] [S20] A model that produces 500 garment pictures in an hour while requiring 40 hours of senior technician rework to rescue one cuttable style degrades plant productivity. Following GFT's performance-management design and the empirical evaluation guidance of NIST AI RMF 1.0, GFT defines the primary unit of account as an **accepted style**: a design whose complete specification dossier has passed technical review, physical material linkage, size-aware fit verification, rights clearance and receiving-CAD readback. [S19] [S20] [S22]

Around that unit of account, GFT proposes four evaluation metrics whose numerators and denominators must be fixed before any pilot begins: [S19] [S20]

- **Cost per accepted style:** Total product-development expenditure during the window (designer and technician hours, software subscriptions, compute charges, sample fabric, trim consumption and failed-candidate rework) divided by the count of styles achieving formal specification release. [S19]
- **Time to approved tech pack:** Elapsed calendar time and active touch hours from sign-off of a constrained brief to final release approval of the technical package, tracked separately for queue waiting and active engineering. [S19] [S20]
- **CAD round-trip pass rate:** Exported pattern packages that reopen in the receiving CAD system with all panels, units, grainlines, notches, seam allowances and grading tables intact on the first attempt, divided by total exported candidate packages. [S22]
- **Fit rejection rate:** Physical or calibrated fit evaluations that fail agreed dimensional or movement tolerances across the target size range, divided by total fit evaluations conducted for that style cohort. [S13] [S27]

METHOD

System boundary and denominator lock for specification trials

To compare manual development against an AI-assisted specification workflow without survivorship bias, the boundary must begin when a constrained brief is issued and end only when receiving-CAD readback and physical fit approval are signed off. [S19] [S22] Every generated candidate, rejected prompt branch, failed geometry script and aborted sample round remains in the cost and time numerator, while the denominator counts only styles passing all seven elements of the release dossier. Locking these definitions before testing prevents a vendor or team from redefining a partial mesh as a completed deliverable midway through a trial. [S19] [S20]

WORKED EXAMPLE

Illustrative comparison: cost per generated image versus cost per accepted style

Consider an explicitly illustrative product-development scenario constructed to demonstrate denominator selection (all inputs are hypothetical teaching assumptions, not observed GFT costs or market quotes). Suppose a studio evaluates an assisted workflow over one quarter to deliver 10 woven shirt styles, generating 400 concept images and 40 raw parametric pattern drafts. Quarterly software and compute charges are assumed at XAF 600,000; designer prompting takes 50 hours at an illustrative rate of XAF 6,000 per hour (XAF 300,000); technician inspection, geometry repair and DXF readback across the 40 drafts take 160 hours at an illustrative rate of XAF 7,500 per hour (XAF 1,200,000); and physical fit sampling consumes XAF 900,000 in cloth, trims and sample-room labour.

Total illustrative expenditure across the window is `XAF 600,000 + XAF 300,000 + XAF 1,200,000 + XAF 900,000 = XAF 3,000,000`. Dividing only software and prompting outlay (`XAF 900,000`) by 400 generated images yields an apparent cost of `XAF 2,250` per concept, and dividing total outlay (`XAF 3,000,000`) by 40 pattern drafts yields `XAF 75,000` per draft. However, if only 8 of the 10 target styles pass CAD readback and physical fit inspection while 2 fail armhole ease checks, GFT's required denominator of 8 accepted styles gives an illustrative cost of `XAF 3,000,000 / 8 accepted styles = XAF 375,000 per accepted style`. The 166.7-fold gap between `XAF 2,250` per picture and `XAF 375,000` per accepted style shows why image-level metrics conceal technician repair effort and sample rejections. This worked example is an illustrative arithmetic model, not a forecast.

1.12 A SINGLE ILLUSTRATIVE COTTON SHIRT ANCHORS THE FOURTEEN CHAPTERS

To make representation and governance rules inspectable across the paper, Chapters 1 through 14 carry one explicitly illustrative teaching example: a short-sleeved woven cotton shirt for a greenfield CMT line in Cameroon. [R06] [S09] The shirt uses a two-piece collar with stand, a two-piece back with shoulder yoke, left and right button-front panels, a left-chest patch pocket, two short sleeves with turned hems, and an optional surface print motif on the yoke or pocket facing. [S02] [S09] Across later chapters, this garment tests controlled block modifications (collar and pocket variants), curved armhole seam correspondence, woven cotton shrinkage and drape, size-aware chest and armhole ease, repeat-tile alignment, cultural motif permissions, and AAMA DXF readback. [S12] [S22] [S27]

This woven cotton shirt is an analytical teaching device grounded in the internal feasibility study (prepared 31 January 2024). It has no buyer, no purchase order, no wear-trial dataset and no claimed GFT operating result. [S19] Where live manufacture requires specific dimensions, tolerances or fabric test results, the teaching brief marks those fields as assumptions or awaiting local evidence.

EXHIBIT 04 / TEACHING SPECIFICATION FIELDS

The running example has assumptions rather than customers

Specification domain	Illustrative baseline assumption	Permitted variant or missing input	Accountable review owner
Base garment block	Woven cotton short-sleeved shirt, 11 pattern pieces	Collar point and chest pocket shape may vary; body block fixed	Senior pattern technician
Target fabric input	Local plain-weave shirting cotton (context R06)	Awaiting measured mass (g/m2), thickness, shrinkage and batch ID	Materials and quality lead
Size range and ease	Illustrative 5-size adult grading run	Awaiting local wearer dimensions and buyer chest/armhole ease table	Pattern and fit technician
Surface artwork	Solid dyed base or bounded yoke accent motif	Awaiting repeat-tile dimensions, dye proof and cultural permission	Design and rights reviewer
CAD exchange contract	2D vector panels with seam allowances and notches (S22)	Awaiting receiving-system DXF convention (AAMA/ASTM) and readback	CAD systems technician

ILLUSTRATIVE [R06] [S09] [S22] Illustrative teaching brief for a woven cotton short-sleeved shirt carried across WP10, aligned with the cut-make-trim product scope of the GFT Cameroon textile programme (feasibility study prepared 31 January 2024) and structured pattern requirements (S09, S22). All product parameters are teaching assumptions; no commercial buyer or production order exists.

Exhibit 4, p. 26 summarises the initial state of the running cotton-shirt example before it enters Chapter 3's brief-constraint rules. [R06] [S09] [S22] By separating fixed base-block assumptions (11 pattern pieces) from missing empirical inputs (unmeasured shrinkage, uncollected local anthropometric distributions, unverified dye fastness and untested receiving-CAD conventions), the table shows what must be collected before any generated variant is cut. [S09] [S22] Because **Exhibit 4**, p. 26 is an illustrative construct, it is not a factory cutting sheet or commercial specification.

1.13

The evidence base supports conditional adoption

The research record examined in this white paper closes on 30 October 2026. [S19] Every local source entry (S01 through S32, plus chapter additions S0101 and S0102) was opened and inspected directly during preparation. Even so, a candid engineering assessment must state what the opened literature can and cannot support. For most preprints published between 2023 and October 2026, verification rests on the mech-

anisms, task definitions, datasets and headline metrics explicitly stated in their published abstracts, supplemented by full-text inspection of foundational datasets such as GarmentCodeData and governance frameworks from NIST. [S19] [S20] [S30]

Abstract-level evidence establishes the existence of a technical representation (for example, that Design2GarmentCode synthesises parametric pattern-making programs or that GarmentDiffusion encodes three-dimensional sewing patterns as compact edge tokens) and identifies the evaluation denominator reported by the authors. [S09] [S0101] It cannot verify hidden failure modes, hyperparameter sensitivity, licence restrictions, installation stability, or generalisation to woven cotton styles outside synthetic training distributions. [S20] [S30] Furthermore, synthetic datasets such as GarmentCodeData explicitly implement quality-control filters to exclude bad samples during generation, meaning that benchmark distributions reflect curated parametric variations rather than the edge cases of a commercial sample room. [S30]

1.14 CORRECTING OVERBROAD CLAIMS IN THE INHERITED ANNEX

Source hygiene also requires correcting places where an earlier secondary summary overstated what an underlying paper proved. In the inherited annex, entry R31 cited arXiv preprint 2509.04705 to support a broad statement that by 2025 to 2026 the apparel AI field had moved to multi-model garment-design pipelines and physics-constrained generation coupling neural synthesis with simulation validation. [R31] Direct inspection of arXiv:2509.04705 (S01, *Transforming Fashion with AI: A Comparative Study of Large Language Models in Apparel Design*, submitted 4 September 2025) shows a much narrower scope. [S01]

Rather than building a physics-constrained pattern-and-simulation pipeline, S01 evaluates three general-purpose large language models (OpenAI, GeminiAI and Deepseek) across apparel design tasks, reporting approximate fabric-classification accuracy of 80 percent on base construction versus 55 percent on detailed construction on its custom dataset, and noting that Deepseek faced significant challenges in generating and recognising fabric images. [S01] While other opened papers in our register investigate simulation-ready garment generation, stitch prediction and visual agents (such as S09, S12, S17 and S0102), S01 itself does not substantiate the broad pipeline claim attached to R31. [S01] [S09] [S12] [S17] [S0102] This paper therefore relies on the specific findings of S01 and the individual pattern-generation studies rather than repeating the summary in R31. [S01]

UNRESOLVED OPERATIONAL CONSTRAINTS AT THE PUBLIC-SOURCE CUTOFF

Four constraints remain unresolved in the public evidence base as of 2026-10-30 and require local verification before capital commitment: [S19] 1. **Abstract-level research coverage:** Most 2024 to 2026 pattern-generation papers (S09, S10, S12, S17, S0101, S0102) were reviewed at abstract level; runnable code, weights and commercial licences have not been audited on local hardware. [S09] [S20] 2. **Unmeasured local materials and bodies:** No mechanical test dataset for Cameroonian woven cotton lots and no consented local anthropometric size survey are included in this register. [R06] [S30] 3. **Untested CAD round-trip execution:** Vendor documentation confirms DXF export options (S22), but no live export-and-readback trial has been executed between a generative model and a factory marker system. [S22] 4. **Unquoted commercial economics:** Software fees, cloud compute bills, technician training hours and sample avoidance rates remain unmeasured until a bounded Phase 3 pilot is conducted. [S19] [S20]

1.15 CONDITIONAL ADOPTION PROTECTS THE FACTORY FLOOR

The evidence assembled in this opening chapter supports conditional, gated evaluation rather than wholesale rejection or uncritical deployment. A greenfield apparel maker in Cameroon should track the shift from pixel-only diffusion models to structured, programmatic and multimodal sewing-pattern representations. [S02] [S09] [S0101] Equally, plant leadership must recognise that no published model guarantees automatic manufacture, higher operating profit or lower environmental impact merely because its renders look persuasive. [R27] [S19] [S20] Every candidate tool must be tested inside a technician-led workflow that preserves physical material identity, explicit ease and grading control, rights permissions, and verifiable CAD readback. [S13] [S22] [S27] Nothing in this chapter or white paper constitutes investment, legal or tax advice; all commercial, technical and compliance decisions require project-specific validation. [S19]

An image can propose a garment, but only an accountable specification may command the blade.

CHAPTER SOURCE NOTES

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S12 Authors via arXiv; WACV 2026. 2026-02-25.

S26 Authors via arXiv. 2025-04-17.

S22 CLO Virtual Fashion. 2026-07-15.

S19 National Institute of Standards and Technology. 2023-01-26.

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

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

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

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

S20 National Institute of Standards and Technology. 2024-07-26.

S27 Authors via arXiv; Computers & Graphics. 2026-06-28.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S0101 Authors via arXiv. 2025-04-30.

S0102 Authors via arXiv. 2025-04-02.

S13 Authors via arXiv; SIGGRAPH 2026. 2026-04-09.

S01 Authors via arXiv. 2025-09-04.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

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

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

S17 Authors via arXiv. 2026-09-16.

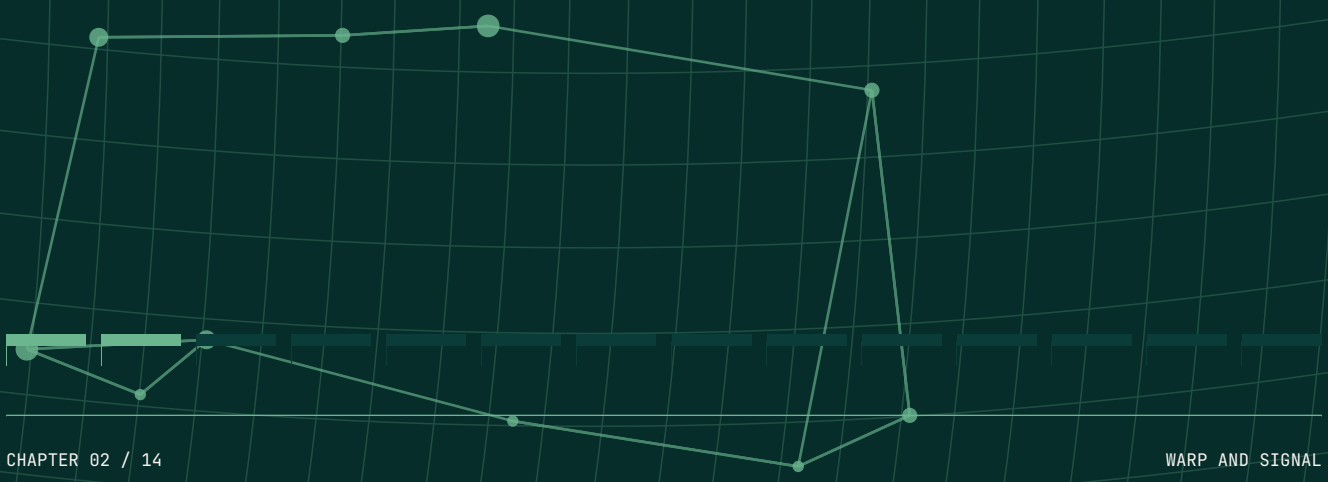
R31 arXiv 2509.04705. Publication date unrecorded; inherited source-text check.

02

PART I: PRODUCTION READINESS IS THE TEST / THE EVIDENCE TIMELINE

Structured garment research predates the recent agent wave

The 2023 to 2026 record shows multiple technical paths advancing together, with no basis here for an industry adoption rate. [S02] [S18]



CHAPTER 02
DECISION TEST

The 2023 to 2026 record shows multiple technical paths advancing together, with no basis here for an industry adoption rate. [S02] [S18]

The research record supports a precise but limited conclusion: structured garment generation was already active before the recent interest in agents, and later systems added interfaces, representations and refinement loops rather than following a single march from pictures to production. GarmentCode addressed parametric sewing-pattern construction in June 2023, while Sewformer addressed image-to-pattern reconstruction in November 2023. [S02] [S03] Text-guided sewing patterns, multimodal pattern tokens, generated programs, parameter files, material inference, fit-aware data and schema-aware synthesis then developed along partly independent paths. [S04] [S08] [S09] [S10] [S11] [S16] None of these publication events measures how many apparel firms use a method, whether a factory accepts its output or whether a buyer approves the resulting garment. GFT should therefore read the chronology as a map of testable techniques, not as evidence of market adoption or production readiness.

2.1

The historical window is a baseline rather than a market census

The inherited programme baseline describes apparel generative AI in 2023 and 2024 as mainly design-side image generation across style, silhouette, colour and surface pattern. [R30] That description is useful because it records what the source programme emphasised when framing its digital strategy. It is not a count of tools, firms, users or factory installations. The inherited entry also lacks a complete journal citation, so it cannot carry a stronger claim about the whole research field or any national market. [R30] The correct use of R30 is therefore narrow: retain the programme's image-led starting point, then test that starting point against opened research records.

The opened record immediately complicates an image-only origin story. GarmentCode's first submission introduced a domain-specific language for hierarchical, component-oriented sewing-pattern construction, with meaningful design parameters and body measurements separated from low-level geometry. [S02] Sewformer's first submission then addressed recovery of sewing patterns from daily photographs, using synthetic image-and-pattern pairs for training and evaluation. [S03] Both first submissions predate the 2024 text-guided and multimodal systems reviewed below. [S04] [S08] A chronology that begins with images and discovers patterns only after agents arrive would erase these earlier structural contributions.

This distinction changes the adoption question. The useful question is not when AI became capable of making a garment, because that phrase hides unlike tasks. It is which representation a particular method receives, which artefact it emits, what its authors evaluated and which decisions remain unresolved. [S02] [S09] [S20] A parametric construction language, an image-reconstruction network and an appearance generator can all be described as garment AI while addressing different engineering boundaries. Counting them as one market would produce a denominator with no coherent meaning.

The chronology also has a selection boundary. This chapter reviews a purposeful set of technical and governance sources, not a systematic census of publications, products or deployments. First-submission dates show when a research claim entered the public record; they do not show when code became usable, when a revision changed the method or when any manufacturer adopted it. [S07] [S20] Publication density cannot substitute for an adoption survey, and a source's presence in this paper cannot substitute for local qualification.

Exhibit 5, p. 32 places the first submissions used in this chapter on one time axis. Its purpose is to make coexistence visible: parametric structure and image-to-pattern reconstruction appear early, while material, fit, schema and agentic methods join later. The axis records publication events only.

EXHIBIT 05 / DECIMAL YEAR OF FIRST ARXIV SUBMISSION

Explicit sewing representations appear throughout the research window

2.02e+03

GarmentCode

2023-06-06: parametric sewing-pattern programs (S02).

2.02e+03

DressCode

2024-01-29: text-guided sewing patterns and textures (S04).

2.03e+03

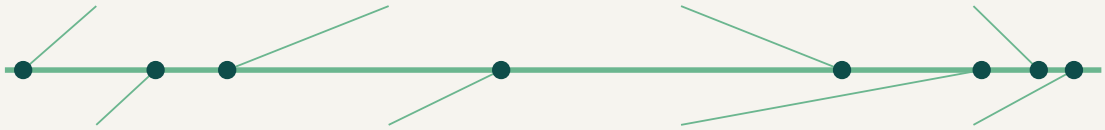
Image2Garment

2026-01-14: geometry with material and physics estimation (S11).

2.03e+03

GarmentWeaver

2026-08-31: schema-aware pattern synthesis (S16).



2.02e+03

Sewformer

2023-11-07: image-to-pattern reconstruction (S03).

2.02e+03

Multimodal cluster

December 2024: Alpparel, Design2GarmentCode and ChatGarment (S08, S09, S10).

2.03e+03

Structure and ease

June 2026: PatternGSL and EASE (S14, S27).

2.03e+03

DressAgent

2026-10-11: closed-loop inverse garment design (S18).

DERIVED [S02] [S03] [S04] [S08] [S09] [S10] [S11] [S14] [S16] [S18] [S27] Derived visual positions from cited first-submission dates: decimal year = calendar year + elapsed days since 1 January / days in that calendar year. Cluster positions are display locations within the named month, not additional publication dates; the source register preserves every exact date. These are publication events, not adoption measurements or capability scores.

The timeline shows accumulation, not a relay race in which each approach replaces the last. Image2Garment still estimates geometry, but adds material attributes and physical-parameter prediction; EASE makes local material allowance an explicit design variable; DressAgent edits patterns through repeated simulation and judgement. [S11] [S18] [S27] These are different responses to different missing variables. The chart cannot establish that later methods are better at a shared task, because their inputs, outputs and evaluation settings differ. Nor can the spacing between publication events reveal commercial diffusion.

For a greenfield maker, the practical baseline is consequently wider than image ideation but weaker than a deployment claim. The research menu already includes inspectable structure, so a buyer need not treat a rendered garment as the only available interface. [S02] [S09] At the same time, the menu contains research artefacts whose local code, data, licences and operating behaviour still need inspection. The chronology justifies a structured evaluation, not a purchase forecast.

2.2

Text guidance reached explicit sewing representations

DressCode joined natural-language guidance to two distinct outputs: SewingGPT generates sewing patterns with text guidance, while a separate Stable Diffusion branch generates tile-based physically based rendering textures. [S04] Its stated target is CG-friendly garments through natural-language interaction, with pattern completion and texture editing supporting the design process. [S04] That combination matters because a sewing representation and a surface appearance can be produced within one research framework without becoming the same artefact. A texture can describe the rendered surface while leaving panel dimensions, seam allowances and production tolerances unresolved.

The paper's production language therefore needs an artefact boundary. A text prompt expresses intent. A generated sewing representation encodes candidate construction information. A rendered result presents an appearance for review. [S04] [S09] The three can be linked in a workflow, but approval of one does not automatically approve the others. A designer may accept the rendered collar proportion while a technician rejects the corresponding panel relationship. Conversely, a structurally valid pattern may produce an appearance that the designer did not intend. The research output becomes useful only when the review route preserves those separate decisions.

Garment3DGen illustrates a different route. It starts from a base mesh and an input image, uses image-to-three-dimensional generation as pseudo ground truth, then deforms the template mesh while preserving mesh quality and topology for draping and simulation. [S05] This is not DressCode's explicit text-to-sewing-pattern route. A simulatable mesh can be valuable for visualisation without supplying an editable flat pattern or a factory tech pack. [S04] [S05] [S22] The word garment in both titles does not make their outputs interchangeable.

Three input-output cards make the distinction operational. DressCode receives text and emits sewing patterns plus rendering textures, leaving factory specification outside its demonstrated CG-friendly objective. [S04] Garment3DGen receives an image and base mesh and emits a textured simulation asset, leaving explicit sewing-pattern export outside the abstract's claim. [S05] Sewformer receives a daily photograph and predicts a sewing pattern after training on synthetic pairs, leaving the transfer from its evaluated setting to a maker's product and material records unresolved. [S03] Each card identifies a research contribution and a boundary; none supplies a comparative production score.

This is why the chapter does not rank the methods by a single maturity number. A common score would require a common task, denominator and acceptance rule. The opened abstracts do not provide that shared protocol across text-guided pattern generation, template-mesh deformation and photo-based reconstruction. [S03] [S04] [S05] GFT can legitimately ask which artefact is needed next, then test the candidate for that artefact. It cannot infer factory readiness from the richness of a rendered result.

2.3

Multimodal models widened the interface to pattern systems

The December 2024 cluster did not converge on one interface. AIpparel introduced custom tokenisation for sewing patterns and fine-tuned large multimodal models on text, images and pattern annotations for generation and editing. [S08] Design2GarmentCode translated images, textual descriptions, sketches or combinations into parametric pattern-making programs. [S09] ChatGarment fine-tuned a vision-language model to emit JSON containing garment descriptions and continuous numerical attributes, then used a programming parametric model to create sewing patterns. [S10] All three widen the input interface, but their intermediate artefacts expose different inspection points.

That difference determines where deterministic processing begins. In Design2GarmentCode, the learned model emits a program in a structured pattern-making representation; the program, rather than a raw image, becomes the object that can be inspected and executed. [S09] In ChatGarment, the learned model emits a JSON record, and the downstream parametric model constructs the pattern from its fields and values. [S10] AIpparel's abstract establishes prediction and editing of a custom pattern token sequence, but does not by itself establish the same downstream program-execution boundary. [S08] A procurement review should therefore name the generated representation and the decoder, compiler or interpreter that follows it.

Exhibit 6, p. 35 extends this input-output reading across the principal methods in the chapter. It deliberately omits a common performance column. The table is a classification of documented artefacts, not a league table assembled from incompatible experiments.

EXHIBIT 06 / DOCUMENTED INPUT, OUTPUT AND REPRESENTATION

A model comparison needs common tasks before common scores

Method	Documented input	Output or intermediate artefact	Boundary retained for local review
DressCode (S04)	Text guidance	Sewing patterns and PBR textures	CG-friendly output is not a released factory specification
AIpparel (S08)	Text, images and sewing patterns	Custom sewing-pattern tokens for generation and editing	Decoder behaviour and local output validity need testing
Design2GarmentCode (S09)	Images, text, sketches or combinations	Parametric pattern-making programs	Program execution and garment acceptance remain separate
ChatGarment (S10)	Images, sketches, text and editing instructions	JSON descriptions and continuous parameters	Schema validity does not prove garment validity
Image2Garment (S11)	A single image	Geometry, material attributes and physical-parameter estimates	Estimated material inputs require physical validation
PatternGSL (S14)	A single image	Template-free specification with panels, seams and stitch topology	Deterministic decoding is not buyer or production approval
GarmentWeaver (S16)	Sketches and textual descriptions	Hierarchical, schema-aware sewing-pattern targets	Executability claims remain benchmark-specific
DressAgent (S18)	A single image of a layered outfit	Layered patterns refined through editing, simulation and judgement	Agent refinement does not confer release authority

SOURCED [S04] [S08] [S09] [S10] [S11] [S14] [S16] [S18] Sourced from the opened abstracts. Each row describes the authors' stated task and mechanism. No cross-paper accuracy, speed or factory-readiness ranking is inferred because the tasks, datasets and evaluation denominators differ.

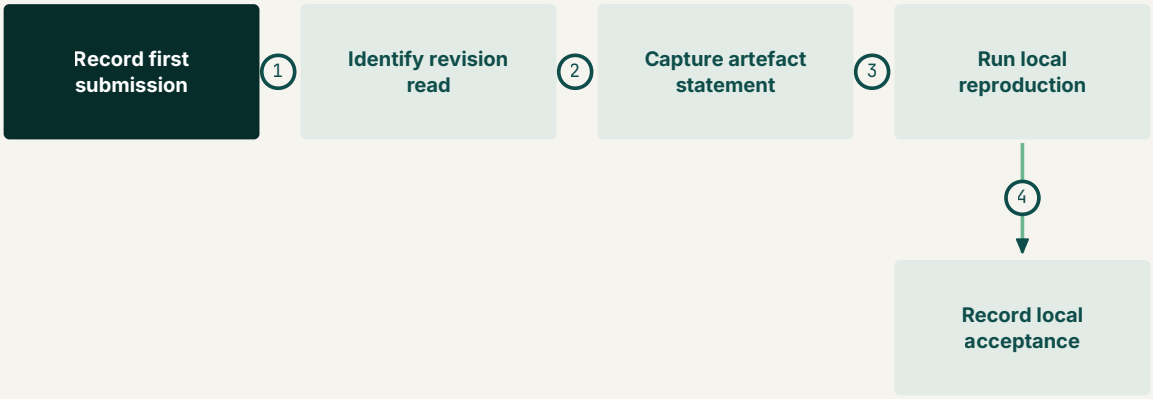
The comparison shows why model names are poor units of procurement. One team may need an editable parameter file, another a pattern token for research, and another a complete refinement workflow.**Exhibit 6**, p. 35 makes those choices visible without pretending that a texture, token, JSON object, program and agent trace can be ranked on one unexamined scale. The next step is a matched local task with the same brief, material and acceptance conditions, not a synthesis of unrelated headline results. [S20]

Dates need equal care. AIpparel was first submitted on 5 December 2024 and the version read was revised on 5 April 2025; Design2GarmentCode was first submitted on 11 December 2024 and revised on 9 April 2025; ChatGarment was first submitted on 23 December 2024 and revised on 3 April 2025. [S08] [S09] [S10] The first date locates the research origin. The revision date identifies the abstract actually read. Assigning every feature in the revised abstract to the original submission would create a false capability history.

HAIGEN supplies a longer revision example: its first submission was on 1 August 2024, while the version read for this chapter was revised on 27 July 2026. [S07] **Exhibit 7**, p. 36 turns that distinction into a record-keeping rule. Publication, revision, local reproduction and local acceptance are separate events with separate evidence.

EXHIBIT 07 / PROPOSED EVIDENCE SEQUENCE

Version dates change the meaning of a chronology



1. do not backdate features
2. code, data or none stated
3. only after execution
4. only after defined review

ILLUSTRATIVE [S07] [S20] Illustrative GFT evidence workflow. S07 supplies the first-submission and revision-history example; S20 motivates contextual empirical evaluation. No arrow asserts that local reproduction or acceptance has occurred for HAIGEN or any other method in this chapter.

The flow prevents two common distortions. First, an old first-submission date must not be used to imply that the original version contained every later refinement. Second, a recent revision must not be mistaken for a recent factory deployment. The local reproduction and acceptance nodes remain empty until the maker has run and reviewed the method under its own conditions. [S20]

2.4

Materials and fit became explicit research targets

Later research also exposes variables that an attractive garment image leaves unstated. Image2Garment addresses the absence of material properties in geometry-only prediction by inferring material composition and fabric attributes, then mapping those attributes to physical fabric parameters through a predictor trained on a small dataset of material-physics measurements. [S11] FIT addresses a different omission: appearance-driven virtual try-on can depict a garment as well fitted without adequately representing the person's and garment's sizes. [S13] EASE addresses yet

another: local material allowance is often a byproduct of geometry or simulation rather than an independent design variable. [S27] These contributions should not be collapsed into a general claim that models now understand garments.

An input-output card keeps the first boundary clear. Image2Garment receives a single image and produces simulation-ready garment estimates that include inferred material information and predicted physical parameters. [S11] Its research contribution is the connection between visual attributes and simulation inputs. Its limitation for a maker is equally specific: an inferred composition or parameter is not a measurement certificate for the fabric actually purchased. The small measurement dataset grounds the research method, but does not turn every subsequent image-based estimate into a local material observation. [S11]

FIT receives the virtual try-on problem as its research boundary rather than treating it as solved by image realism. The authors introduce synthetic try-on triplets accompanied by body and garment measurements, including ill-fit cases created through programmatic garments and physics simulation. [S13] The output is a dataset and a fit-aware baseline model, not a physical fit approval for a wearer. Its importance lies in evaluation design: deliberately mismatched sizes make a previously hidden failure visible. It does not follow that every commercial try-on system has the same failure or that the new baseline resolves comfort, movement or buyer acceptance. [S13]

EASE begins with a parametric body surface and represents ease as spatially varying, anisotropic per-triangle scales. It then optimises sewing patterns intended to preserve the specified ease while satisfying geometric and seam constraints. [S27] This makes ease an editable input rather than merely an outcome observed after simulation. The capability does not establish which ease distribution a Cameroonian wearer population prefers, nor whether a chosen fabric and activity will feel comfortable. Those remain product and physical-validation questions rather than transferable accuracy claims. [S27]

The three cards therefore answer different questions. Image2Garment asks how an image might inform material and simulation parameters. FIT asks whether size information and ill-fit examples enter try-on evaluation. EASE asks how a designer can specify and transfer local allowance. [S11] [S13] [S27] A combined accuracy league table would erase these distinctions. It would also hide the difference between inferred material properties, synthetic fit data and an explicit design representation.

For GFT, the chronology supports a more modest inference: research has made missing variables easier to name and test. A local evaluation should carry each variable into its own evidence field, with the source of the value, the evaluated case and the unresolved physical check recorded separately. [S20] Material estimation can propose a swatch test, fit-aware data can inspire matched and mismatched cases, and explicit ease can inform a specification review. None of these research results replaces the corresponding local evidence.

2.5

Structure and agentic refinement address different failure modes

PatternGSL and GarmentWeaver make structure the immediate object of generation, while DressAgent makes iterative refinement central. These are complementary research directions, not successive labels for the same capability. PatternGSL defines a template-free specification language containing panel boundaries, parameterised seams and explicit stitch topology, then decodes predicted specifications with lightweight deterministic validity handling. [S14] GarmentWeaver responds to flat long sequences that entangle garment structure with detailed parameters by constructing hierarchical targets and applying feasibility-aware regularisation. [S16] DressAgent addresses inverse design through planning and a closed loop of pattern editing, physical simulation, and visual or measured judgement. [S18]

The failure modes differ. PatternGSL targets a representation gap between surface reconstruction and structured garment construction. [S14] GarmentWeaver targets redundant components, inaccurate local details and poor simulation compatibility associated with its authors' account of flat sequence models. [S16] DressAgent targets the difficulty of recovering layered outfits from an image, including layer order, coarse pattern estimation and fine-grained refinement. [S18] A schema can make an invalid relationship easier to detect; an agent can propose another edit after inspecting a simulated result. Neither mechanism alone establishes that the final garment satisfies a production brief.

There is also a control distinction. Deterministic validity handling applies defined rules during decoding, whereas an agentic loop chooses further actions on the basis of tools, critiques and memory. [S14] [S18] The first can reject or repair a specified class of structural conditions without proving design intent. The second can improve alignment with a reference under its evaluated protocol without gaining authority to approve a pattern. A local pilot should inspect both the structural error log and the agent's action history rather than compressing them into a single success label. [S20]

Public-availability statements require the same precision as technical claims. The PatternGSL abstract says that code and data-processing scripts will be released, which is a promise recorded at the accessed version, not evidence that this lane installed them. [S14] The GarmentWeaver and DressAgent abstract pages read for this chapter contain no code-release statement. [S16] [S18] That absence does not prove that no artefact exists elsewhere; it means availability remains unverified in the source boundary used here. By contrast, the Sewformer and ChatGarment abstracts explicitly state that code and other artefacts are available. [S03] [S10] Even those statements do not prove that a particular revision runs, that its licence permits the intended use or that its output passes a local CAD readback.

A claim ledger should consequently use separate labels: code stated as available, code promised, no release statement in the source read, package retrieved, package executed and output accepted. The first three describe publication evidence. Retrieval and execution describe subsequent work. Acceptance describes a decision under local requirements. Treating a project link or a future-release sentence as a verified installable package would skip most of that evidence chain. [S20]

The recent agent wave therefore changes the research menu without invalidating earlier structure. A generated program can be the object an agent edits; a schema can constrain what that edit may contain; simulation can supply evidence for another iteration. [S09] [S16] [S18] The adoption question is whether the combined chain works on the maker's cases with acceptable corrections and approvals. This chapter did not execute these systems, so it reports their stated mechanisms and availability language only.

2.6

Commercial documentation establishes a narrower fact than factory success

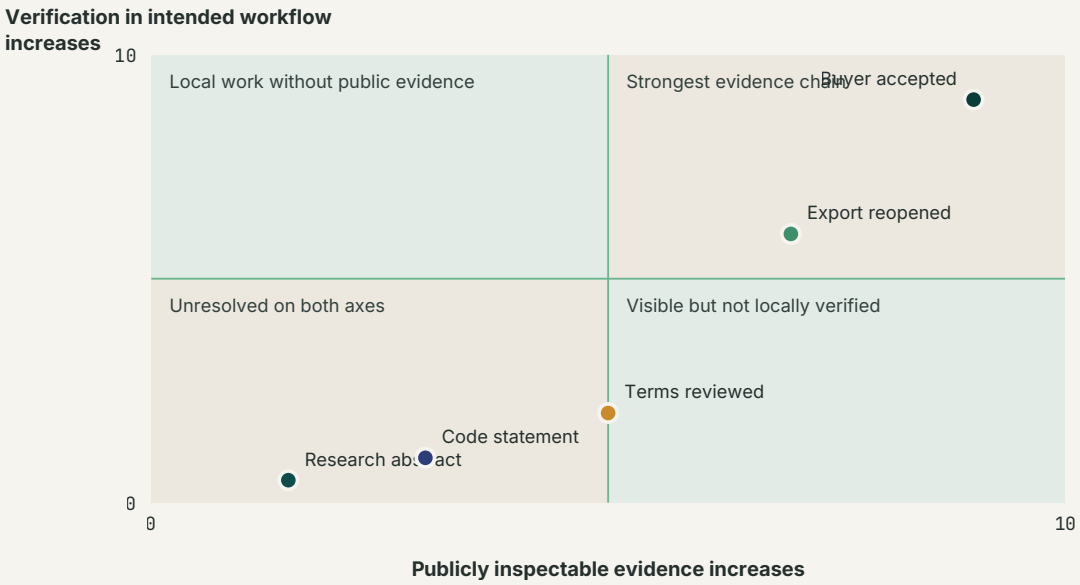
Vendor documents answer different questions from research abstracts, and both differ from local acceptance. CLO's technical documentation lists Standard DXF, AAMA and ASTM options, together with controls for units, grading, notches, seam allowances and the relationship between cutting and sewing lines. [S22] This establishes documented exchange functionality at the page's stated update. It does not establish that a file generated by a research model will reopen correctly in the maker's receiving CAD release, nor that grading and annotations will survive a particular round trip. Those are testable interoperability questions, not implications of a format list. [S22]

Adobe's Firefly FAQ states that outputs from features without the beta label may be used in commercial projects, with beta features subject to their stated conditions. [S31] That is a vendor-documented use condition for the identified service context. It does not prove that a generated garment design is non-infringing, that every partner model has identical terms or that the design is manufacturable. A procurement record must identify the actual feature, model, plan and terms reviewed rather than transferring one FAQ sentence to an entire toolchain. [S20] [S31] This is evidence discipline, not a legal clearance or legal advice.

Exhibit 8, p. 40 separates public inspectability from verification in the intended workflow. The points are generic evidence states, not scores assigned to any named research system or vendor. A method whose licence, runnable package or buyer acceptance is unknown must remain unknown rather than receiving a favourable coordinate by assumption.

EXHIBIT 08 / ILLUSTRATIVE EVIDENCE-STATE COORDINATES

Public availability and local suitability require separate checks



ILLUSTRATIVE [S20] [S22] [S31] Illustrative GFT teaching model informed by S20's evaluation logic, S22's documented CAD controls and S31's vendor-specific use conditions. Coordinates are didactic positions for generic evidence states, not measured scores for a method or product. Unknown licence, runnable-code, export and acceptance fields remain unknown and are not plotted.

The matrix deliberately leaves a large distance between a public statement and an accepted outcome. Reading an abstract can establish a claimed mechanism. Reviewing terms can establish what a particular document permits. Reopening an export can establish that specified geometry and annotations survived the tested exchange. Buyer acceptance records a different decision again. [S20] [S22] [S31] The exhibit cannot show how frequently any organisation has completed these steps because the source base contains no such census.

GFT NOTE

GFT recommends a dated procurement evidence pack for every candidate method: first-submission and revision identifiers, exact model or service release, licence and use terms, code or data statement, system requirements, support scope, data-retention and training settings, representative input and output files, and the receiving-CAD readback result. Store the page or document snapshot used for the decision and mark every unanswered field as pending. Do not convert a demonstration, download button or commercial-use sentence into factory approval. [S20] [S22] [S31]

The pack should also preserve refusal conditions. If the vendor cannot supply an export specimen, the team cannot claim exchange readiness. If data-handling terms remain unresolved, proprietary blocks and buyer material should remain outside the service. If the support commitment does not identify the release tested, a later update requires a new check. These are proposed GFT controls derived from the evidence gaps, not assertions about any vendor's current performance.

2.7

A claim ledger prevents research results becoming sales promises

The main risk in a fast-moving research review is not only technical error. It is a change of meaning as a claim travels. An author's demonstrated benchmark becomes a vendor's general capability statement, then a sponsor's assumed factory result. NIST's generative AI profile instead calls for contextual evaluation, empirical capability checks, source verification and lifecycle monitoring. [S20] Applied here, that discipline requires every statement to retain its evidence state and intended boundary. A research result can justify a test. A vendor document can justify a narrower procurement question. Neither can silently become a locally accepted result.

A useful ledger begins with the publication event. Record the method title, first-publication date, exact version read, revision date and stable source location. The first date supports the chronology; the version identifies what was actually reviewed. [S07] [S08] Add the publication venue as stated by the source, but do not treat a conference reference or journal reference as local reproduction. PatternGSLData and GarmentWeaver, for example, can occupy the same research timeline while having different revision histories and availability statements. [S14] [S16]

The technical entry should then record the input, output, intermediate representation and deterministic downstream component. DressCode's texture branch should not be merged with its sewing-pattern branch, and ChatGarment's JSON should not be confused with the pattern constructed from it. [S04] [S10] The evaluation entry needs the dataset or setting, unit of evaluation, comparator and limitation stated at the level actually read. Where only an abstract was reviewed, the ledger should say so. It should not invent a sample size, protocol detail or confidence interval that the opened text did not supply.

Availability forms another group of fields. Record whether the source states that code is available, promises a future release or makes no release statement, and preserve the sentence and access date. [S03] [S14] [S16] Separately record whether the team retrieved the artefact, checked its licence, installed it, executed it and retained the output. A link is not an execution log. An execution log is not a garment approval. This separation makes reproducibility work visible instead of allowing the word available to carry unsupported assurance.

The local entry belongs to the maker rather than the paper's authors. It should identify the illustrative or real product brief, approved block, material record, model configuration, operator role, corrections, receiving-CAD result, physical checks and final reviewer. [S20] [S22] If those fields have not been populated, the honest status is not ready or not adopted. It is untested in the intended workflow. That status can coexist with a strong published result because the two claims answer different questions.

WORKED EXAMPLE: REGISTER COMPLETENESS IS NOT ADOPTION

Suppose an illustrative review contains 10 candidate methods and 8 required ledger fields per method: first submission, revision read, input, output, evaluation setting, availability statement, local reproduction and local acceptance. The full register has 10 × 8 = 80 cells. Assume the reviewer completes the first six bibliographic and technical fields for every method, producing 10 × 6 = 60 populated cells. No method has been reproduced or accepted locally, so all 20 cells in those last two fields remain empty. The hypothetical register is therefore 60 / 80 = 75 percent complete, while local reproduction is 0 / 10 = 0 percent and local acceptance is 0 / 10 = 0 percent.

The illustrative arithmetic exposes why register completeness cannot be relabelled as readiness. The 75 percent figure describes documentation coverage under the assumed field list. It says nothing about how well any method works, how many firms use it or whether any garment should be released. Changing the field list would also change the completeness percentage without changing the underlying evidence. The decision report should consequently show the individual local-reproduction and local-acceptance rows rather than presenting one reassuring aggregate.

This worked example also clarifies what an adoption estimate would require. The research unit would need a defined population of organisations, a sampling method, an observation date and a distinction between awareness, trial, routine use and production release. The numerator would need verified instances of the chosen adoption state, and the denominator would need eligible organisations rather than papers or products. The sources opened for this chapter supply no such market sample. They support a chronology of methods and a set of local test questions, not an industry adoption rate. [S02] [S20] [S31]

2.8 EACH EVIDENCE STATE SHOULD PERMIT A DIFFERENT ACTION

A claim ledger is useful only if it changes decisions. GFT can allow a documented research mechanism to enter a test backlog when its input and output match a defined need. It can allow vendor-documented functionality to enter a procurement comparison when the exact release and terms are captured. It should require local execution, correction records and receiving-workflow evidence before calling a method operationally suitable. [S20] [S22] Local acceptance should remain a separate decision by the accountable reviewer, with unresolved material, fit, rights or geometry issues returned to their owners rather than averaged away.

The same rule applies to chronology. An early first submission proves precedence for the documented research idea, not present superiority. A recent revision proves what the current record says, not that the capability existed at the earlier date. A code promise indicates intent at the accessed version, not reproducibility. An agent trace can show which edits and simulations occurred, but it cannot authorise a cutting order. [S07] [S14] [S18] Keeping these meanings separate is more useful than declaring one wave of methods victorious over another.

For GFT's greenfield apparel context, the defensible recommendation is to maintain a versioned research register and use it to select bounded tests. Include structured sewing methods from 2023 alongside later multimodal, material, fit, schema and agentic approaches. [S02] [S03] [S11] [S13] [S16] [S18] Require common local tasks before

comparing outcomes, and require local acceptance before describing adoption. The present evidence justifies curiosity with controls. It does not justify a market-share figure, a universal factory-readiness claim or an investment, legal or tax conclusion.

A publication date can locate an idea; only local evidence can authorise its use.

CHAPTER SOURCE NOTES

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S18 Authors via arXiv. 2026-10-11.

S03 Authors via arXiv; ACM Transactions on Graphics. 2023-11-07.

S04 Authors via arXiv. 2024-01-29.

S08 Authors via arXiv. 2024-12-05.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S11 Authors via arXiv. 2026-01-14.

S16 Authors via arXiv. 2026-08-31.

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

S20 National Institute of Standards and Technology. 2024-07-26.

S07 Authors via arXiv. 2024-08-01.

S14 Authors via arXiv. 2026-06-23.

S27 Authors via arXiv; Computers & Graphics. 2026-06-28.

S05 Authors via arXiv; 3DV 2025. 2024-03-27.

S22 CLO Virtual Fashion. 2026-07-15.

S13 Authors via arXiv; SIGGRAPH 2026. 2026-04-09.

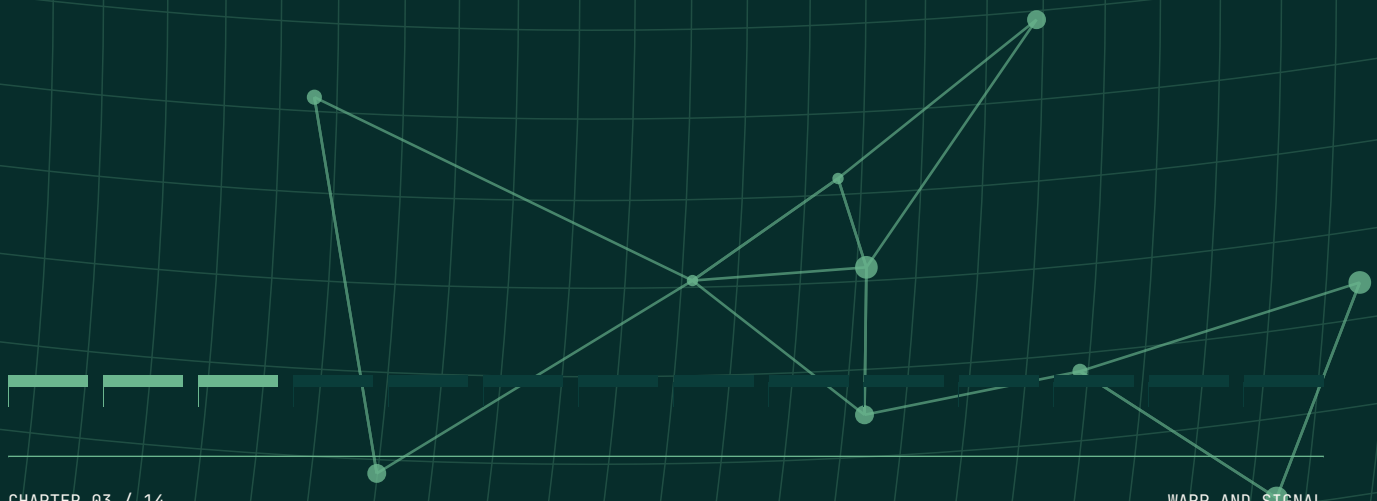
S31 Adobe. 2026-05-13.

03

PART I: PRODUCTION READINESS IS THE TEST / INTENT BEFORE IMAGES

Controlled ideation starts with a manufacturable brief

Image generation can broaden exploration when designers retain the constraints and decide which concepts deserve technical work.



CHAPTER 03
DECISION TEST

Image generation can broaden exploration when designers retain the constraints and decide which concepts deserve technical work.

A greenfield apparel maker should permit generative image and sketch exploration only after a designer records an explicit, manufacturable brief that fixes what the factory can cut, sew and source. [S09] Unconstrained prompting produces a gallery of plausible pictures whose construction contradictions, unavailable fabrics and uncleared references are discovered too late by pattern technicians and buyers. [S01] In GFT's Cameroon textile programme (prepared 31 January 2024, with sources refreshed in 2026), the industrial objective is to help shift a sector where approximately 95 percent of cotton fibre is exported raw and about 5 percent is processed locally toward the national target of 50 percent local processing by 2030. [R06] [R07] Moving fibre into domestic cut, make, trim and finishing lines requires styles that can be engineered, costed and repeated on real machines, rather than unmeasured concept renders. [R06] Under the Warp and Signal discipline, production constraints form the warp and recorded design decisions form the weft: creative assistance earns its keep only when every candidate is filtered against declared physical, material and governance boundaries before pattern development begins. [S20] Nothing in this chapter constitutes investment, legal or tax advice.

3.1

The brief must specify what can change

Creative exploration becomes productive for a factory only when the starting brief separates fixed production constraints from the design variables open to ideation. [S09] In baseline practice during 2023 to 2024, generative artificial intelligence in apparel was applied primarily to open-ended visual ideation across style, silhouette, colour and surface pattern, as noted in the programme's baseline literature record. [R30] While an unconstrained text prompt can synthesise novel silhouettes, an industrial workflow cannot price, grade or sew an adjective. [S09] Research on translating design concepts into parametric sewing patterns, including GarmentCode and Design2GarmentCode, shows that producible garments depend on component-level rules, explicit panel boundaries and valid seam topology rather than pixel coherence alone. [S02] [S09] Before a designer opens an image generator or sketch assistant, the brief must establish the operating envelope within which variation is permitted. [S09]

GFT proposes six mandatory fields for every manufacturable brief in a greenfield cut, make, trim and finishing operation, grounded in the programme's feasibility study base case and multimodal specification research. [S09] First, the brief defines the wearer task and environment, stating whether the garment is a uniform shirt, a workwear layer or a retail piece, together with movements that govern ease and seam selection. [S09] Second, it states the intended size range and base size, flagging where local wearer measurements remain uncollected so the team does not assume an untested body standard. [S13] Third, it binds the concept to fabric availability, naming the approved woven or knitted fabric family, weight band and supply route, or marking the fabric as a pending procurement enquiry. [S01] Fourth, it lists construction constraints, identifying the base block retained, the stitch classes available on the sewing floor and prohibited operations. [S02] Fifth, it bounds permitted embellishments and surface treatments, distinguishing yarn-dyed solids, piece-dyed fabrics and placed or repeat prints alongside their approval gates. [S26] Sixth, it states the target cost

basis and order-volume assumption in CFA francs (XAF) or United States dollars, keeping commercial thresholds visible before a designer adds pockets, yokes, plackets or trims that multiply sewing minutes. [S19]

Unknown buyer requirements must remain flagged as missing inputs rather than masked by fluent prose or renders. [S20] In a greenfield cycle, a design team often starts ideation before a buyer confirms label placement, packaging rules, shrinkage tolerances or trim nominations. [S19] If a brief leaves those fields blank without a pending marker, an image model or vision-language assistant fills the vacuum with arbitrary button counts, welt pockets or contrast topstitching that nobody requested. [S10] Keeping unresolved buyer specifications visible prevents technicians from mistaking a visual default for a commercial instruction. [S20]

To make these controls concrete without inventing a client or factory operating result, this chapter carries forward the paper's illustrative teaching example of a short-sleeved woven cotton shirt. [S09] In this illustrative brief, the approved torso block, armhole curve and front placket construction are held fixed so the sewing line does not re-qualify core fit and folder attachments, while the designer explores variations in collar geometry, chest pocket shape and yoke panelling. [S02] [S09] Demand context from `wp04` and prioritisation in the `use-case-portfolio` tool inform which garment families enter this bounded queue, but demand forecasts in `wp04` never act as permission to release every generated variation to the cutting room. [S19] As shown in **Exhibit 9**, p. 47, structuring the brief into fixed parameters, variable degrees of freedom and unresolved buyer inputs gives both the designer and any multimodal tool an inspectable boundary. [S09]

EXHIBIT 09 / PROPOSED BRIEF FIELDS

A constrained brief is more useful than a list of style adjectives

Brief field	Fixed constraint	Permitted variation	Unresolved buyer input
Wearer task and product role	Warm-climate daywear and institutional uniform shirt	Formal or utility visual styling within woven shirting	Specific institutional dress-code compliance rules
Base block and size range	Approved adult torso and set-in sleeve block retained	Hem profile (straight or curved) within block length	Final size-grading span and local fit-panel sign-off
Fabric availability	Plain-weave or twill cotton shirting from qualified lot	Solid shade or approved single-colour surface motif	Confirmed supplier lot ID, shrinkage and colourfastness
Collar and neckline	Two-piece stand or convertible collar compatible with neckband	Point length, spread angle and collar-tip geometry	Interfacing weight and collar-stay requirement
Pocket and front detail	Single needle or twin-needle patch attachment on chest	One or two chest pockets; mitred, rounded or flapped	Button specification, bar-tack points and logo embroidery
Target cost and assembly basis	Cut, make, trim line using standard lockstitch and overlock	Component count capped at illustrative ten cut panels	Target FOB or ex-works unit price ceiling in XAF

ILLUSTRATIVE [S09] Illustrative GFT teaching brief for a woven cotton shirt, informed by the programme's feasibility study base case and Design2GarmentCode's separation of design intent from parametric pattern structure (S09). Values and rules are teaching assumptions, not a customer order or GFT operating record.

The structure in **Exhibit 9**, p. 47 demonstrates why a constrained brief outperforms a mood board of style adjectives when bridging creative intent to pattern engineering. [S09] Each row pairs a creative degree of freedom with the production rule that keeps the garment sewable on a standard cut, make and trim line, while isolating what still awaits buyer confirmation. [S09] What **Exhibit 9**, p. 47 cannot establish is whether a given fabric lot will meet physical shrinkage tolerances, whether the base block fits a specific Cameroonian wearer cohort, or what unit cost a live sewing line will achieve; those questions require physical measurement in Chapter 6, size-aware fit verification in Chapter 7 and empirical costing in Chapter 12. [S13] [S19]

GFT NOTE

GFT recommends that no prompt or reference image be entered into a generative ideation tool until the product developer completes the six-field brief and obtains sign-off from the responsible pattern technician on which block components are frozen and which are open to variation. [S09] [S20]

3.2

References need permission and design rationale

Every visual reference attached to a brief requires a recorded origin, an explicit permission status and a documented design rationale before it is uploaded to a model or shown to a buyer. [S20] [S21] Design teams assemble mood boards from archival garments, swatch books, photography and regional motifs to guide image conditioning or sketch refinement. [S07] In a generative workflow, feeding an image into an image-to-image pipeline, a style-conditioning adapter or a multimodal prompt is an active processing step that can propagate third-party expressive elements or sensitive cultural symbols into candidate visuals. [S21] [S23] Treating search results as a free input bin exposes a greenfield maker to copyright disputes, contractual breaches and reputational harm with local communities. [S21] [S23]

A workable governance rule distinguishes three reference classes inside the brief. [S21] [S23] First, owned designs comprise pattern blocks, technical sketches, photographs and motifs created internally by the maker's staff or assigned under a written contract with clear title. [S21] Second, licensed commercial assets comprise stock photography, commissioned artwork or partner brand marks whose terms must be checked for permission to upload into external artificial intelligence services, create derivative designs and commercialise apparel. [S20] [S21] As the United States Copyright Office explained in Part 2 of its Artificial Intelligence Report in January 2025, copyright protection for generative outputs in that jurisdiction depends on whether a human author determined sufficient expressive elements, while input rights and third-party similarities remain separate legal exposures. [S21] Third, traditional cultural expressions and heritage designs require dedicated governance because commercial licensing categories do not capture community stewardship or collective identity. [S23] [S24]

References to traditional cultural expressions require particular care in a Cameroonian textile programme. [S23] [S24] The World Intellectual Property Organization notes that traditional cultural expressions are integral to the cultural and social identities of indigenous and local communities, embodying know-how and skills and transmitting core values and beliefs. [S24] Its guidance on fashion and indigenous peoples emphasises dialogue between custodian communities and fashion enterprises to shape fair and equitable interaction rather than unconsented appropriation. [S23] Within Cameroon, textile traditions such as Grassfields royal ndop resist-dyed indigo cloths or northern geometric weaving motifs carry specific social meanings and lineage associations. [S23] [S24] While GFT's visual identity uses an indigo colour token as an abstract nod to Grassfields dye heritage and ochre for northern laterite soils, a product brief for commercial garments must never treat traditional motifs as rights-free style tokens to be remixed by a model. [S23] [S24] Where a brief proposes drawing on a community textile tradition, the maker must halt automated generation and initiate documented engagement with the relevant custodians before use, as examined in Chapter 8. [S23]

To enforce this separation at the start of ideation, GFT proposes attaching a mandatory reference manifest to every brief. [S20] Each entry records six required attributes: a unique reference identifier, the verified owner or custodian community,

the source classification (owned, licensed or cultural reference requiring engagement), the permitted scope of use (internal study, local conditioning, cloud upload or commercial adaptation), the licence or consent expiry date, and restrictions such as prohibitions on model training or publication. [S20] [S21] A link on a public website is a locator string, not a permission record. [S20] If a reference lacks a verified owner and documented permission for the target model environment, the brief gate blocks its upload. [S20]

3.3

Silhouette and local detail require separate review

Reviewing a generated garment image requires separating global silhouette and colour composition from local construction details and placement accuracy. [S26] Controllable generation frameworks show why visual synthesis has become convincing while leaving manufacturing questions open. [S26] In IMAGGarment, first submitted in April 2025 and revised in September 2025, the authors introduce a fine-grained garment generation framework using a two-stage training strategy to model global appearance (silhouette sketch and colour palette) separately from local details such as logo placement. [S26] Research on garment-centric outpainting submitted in March 2025 takes a segmented garment image from a mannequin or person and synthesises showcase poses and model appearances around it via text prompts and a multi-scale appearance module. [S0302] IMAGGarment conditions a new garment image on silhouette, colour and mark constraints, whereas garment-centric outpainting preserves a photographed garment while fabricating the presentation scene. [S26] [S0302]

Neither multi-condition synthesis nor showcase outpainting establishes that a generated garment can be cut, sewn or printed to specification. [S26] [S0302] When a diffusion model places a brand mark or pocket on a rendered shirt front, it optimises perceptual plausibility in two-dimensional pixel space rather than computing physical print registration across curved panels, stretch distortion, embroidery backing stability or needle-penetration limits. [S26] Likewise, a showcase image can conceal structural impossibilities: a collar stand merging into a lapel without a seam line, buttonholes drifting off the centre front grain, or a back yoke contradicting the front shoulder seam. [S09] [S26] Mistaking fine-grained conditioning in an image generator for engineering control over the specification passes hidden rework downstream. [S09]

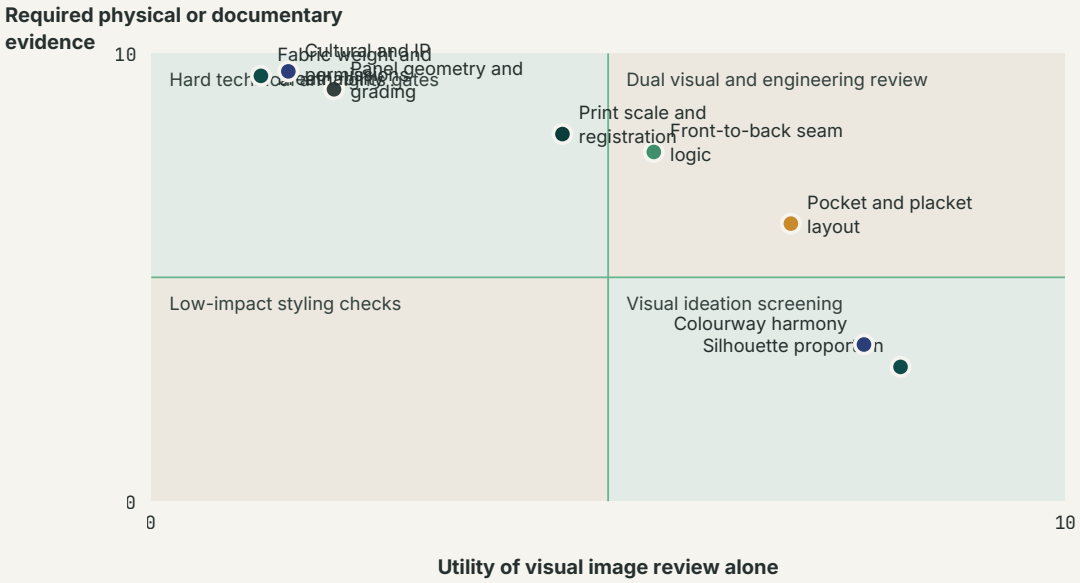
To catch visual contradictions before pattern synthesis, GFT proposes a structured review sheet completed by the designer and checked by a pattern technician. [S09] [S20] First, the sheet checks front-to-back consistency: shoulder slope, side-seam drop, hem curve and yoke height must align across views rather than shifting between generated images. [S09] Second, it verifies component counts and fastening logic: pocket count, symmetry, placket direction, button count and cuff closure style must match the brief's cost envelope and machine list. [S09] Third, it evaluates seam plausibility: every visible shape change on the shirt must correspond to a dart, panel seam, gather or fold rather than airbrushed shading on flat cloth. [S02] [S13] Fourth, where

surface graphics or logos appear, the sheet checks print scale, repeat continuity across seams and clearance status, routing any unverified motif to Chapter 8. [S23] [S26]

As shown in **Exhibit 10**, p. 50, each review dimension occupies a different position when mapped by its reliance on visual inspection versus physical and documentary verification. [S01] [S26]

EXHIBIT 10 / ILLUSTRATIVE VERIFICATION COORDINATES (0 TO 10 SCALE)

A visual review cannot clear a physical property



ILLUSTRATIVE [S01] [S09] [S23] [S26] Illustrative GFT classification of review gates across visual appearance review utility (horizontal axis) and required physical, geometric or legal evidence depth (vertical axis), informed by S01, S09, S23 and S26. Coordinates are illustrative conceptual placements, not measured benchmark scores.

The distribution of tasks in **Exhibit 10**, p. 50 clarifies why an attractive concept image clears only the lower-right quadrant of design screening. [S26] Global silhouette proportion and colourway harmony can be judged visually at the concept stage, and pocket or placket layout can be screened through dual visual and engineering review when paired with the brief's block constraints. [S09] [S26] By contrast, panel geometry, fabric weight and breathability, and cultural or intellectual-property permissions sit in the upper-left quadrant: visual inspection of a generated image provides near-zero evidence for those properties, which require explicit geometric representation (S09), physical swatch testing (S01) and documented rights records (S23). [S01] [S09] [S23] What **Exhibit 10**, p. 50 cannot provide is a failure rate for any model; its coordinates are an illustrative governance map keeping visual sign-off from substituting for engineering or legal clearance. [S20]

3.4

Prompt iteration needs a record of decisions

Iterative prompting and sketch refinement must leave a structured decision trail so the studio knows which constraints were applied, which model version produced each image, and why a designer selected or modified a candidate. [S07] [S20] As demonstrated by HAIGEN (Human-AI Collaboration for GENeration), first submitted in August 2024 and read in its July 2026 fourth revision, an assisted fashion design workflow spans four modules: cloud-based inspiration generation, a local sketch library, a style-recommendation module and an interactive colouring module. [S07] The GarmentSketch benchmark submitted in June 2026 pairs 26,249 sketches across 21 garment categories with textual descriptions refined through a multimodal large language model and human-in-the-loop pipeline, showing both the value of structured text-sketch pairing and model limits on fine-grained sketch fidelity. [S0301] HAIGEN's four-module split is evaluated here from its current revised abstract rather than back-dated to its 2024 submission, and its local-storage design is treated as a research architecture rather than an audited security control. [S07] [S20]

When designers iterate without logging their steps, three failures follow. [S20] First, the studio loses reproducibility: when a buyer or technician asks why a collar spread changed or which reference drove a pocket detail, nobody can reconstruct the inputs or model state. [S20] Second, confidential brief details or block sketches can be pasted into external cloud endpoints if the team does not record which tool and deployment mode were used. [S07] [S20] Third, unlogged exploration blurs the boundary between machine defaults and human design choices, complicating internal handoffs and intellectual-property documentation. [S21]

To prevent untraced drift, GFT proposes a mandatory design log entry for every concept batch and promoted candidate, aligned with the content-provenance principles in the National Institute of Standards and Technology Generative AI Profile. [S20] The design log captures seven fields:

- **Brief and constraint reference:** the identifier and revision of the approved brief governing the session, confirming that fixed block and fabric boundaries were active. [S09]
- **Prompt and negative constraints:** the positive text prompt and negative exclusion terms entered by the designer, preserving how construction limits (such as no pleats, no piping or single chest pocket only) were expressed. [S20] [S0301]
- **Reference manifest identifiers:** the IDs from the brief's reference manifest used for sketch conditioning, style reference or colour transfer, verifying that every input passed permission checks. [S20] [S21]
- **Model, version and execution boundary:** the model name, checkpoint or API version date, adapter weights and whether execution occurred on a local workstation or an external cloud service. [S07] [S20]
- **Sampling parameters and seed:** the random seed, guidance scale and step settings where exposed by the tool, enabling technical staff to inspect sensitivity to stochastic variation. [S20]

- **Output asset identifier and hash:** a stable file identifier and cryptographic digest for the generated sketch or image, preventing silent file overwrites during editing. [S20]
- **Human selection and modification rationale:** a short engineering note stating why the designer selected this candidate over rejected outputs and listing the manual sketch edits, proportion corrections or construction changes applied after generation. [S07] [S21]

> **WARNING**
>
> A completed design log is an internal engineering and provenance record, not a certificate of copyright ownership or freedom to operate. [S20] [S21] As the United States Copyright Office stated in January 2025, merely supplying text prompts does not establish human authorship of generative outputs under United States law, and copyrightability in Cameroon, Hong Kong SAR or export markets depends on local law and the specific human contribution. [S21]

3.5

Material appearance is not material identity

A generated fabric texture or a language model's material description must never be accepted as evidence of fibre composition, weave construction, weight, breathability or supplier availability. [S01] During ideation, multimodal models and image generators render convincing cotton twills, slub linens and cellular knits alongside fluent text descriptions of comfort and durability. [S01] [S04] In the programme's data annex, entry R31 pointed to arXiv:2509.04705 to suggest a transition toward multi-model pipelines coupling neural generation with simulation validation. [R31] Inspecting the underlying paper (S01, submitted September 2025) reveals a narrower empirical finding that product teams should heed. [S01] [R31]

In S01, the authors examined three large language model ecosystems (OpenAI with integrated DALL-E, GeminiAI and Deepseek) across textile manufacturing stages and built a custom dataset for fabric image generation and classification. [S01] On their image generation test for fabric construction, the OpenAI setup achieved a Learned Perceptual Image Patch Similarity (LPIPS) score of approximately 0.2. [S01] When classifying fabric properties from images across factors such as breathability, moisture-wicking and tactile comfort, the top tested model achieved approximately 80 percent accuracy on base construction but dropped to approximately 55 percent accuracy on detailed construction. [S01] Deepseek faced significant challenges in generating and recognising fabric images, and all tested models struggled to recognise complex fabric constructions and intricate designs from images. [S01]

This 25 percentage-point gap between base-construction recognition (approximately 80 percent) and detailed-construction recognition (approximately 55 percent) within the same study shows why pixel appearance cannot substitute for physical material specification. [S01] A model that distinguishes a basic woven from a basic knit in a curated dataset fails nearly half the time on detailed construction attributes, and an image classifier cannot measure yarn count, twist, finishes, weight or shrinkage. [S01] Other routes confirm this boundary: when DressCode (S04) generates rendering

textures alongside patterns from text prompts, its objective is computer-graphics lighting fidelity rather than mill-floor weave specification. [S04] Similarly, when appearance-based virtual try-on models render a garment on a body, the FIT benchmark (S13) notes that current methods default to generating well-fitted results regardless of garment or person size, masking how cloth would strain or drape. [S13]

Exhibit 11, p. 53 summarises the empirical findings in S01 and their implications for a maker's brief and ideation gate. [S01]

EXHIBIT 11 / REPORTED BENCHMARK METRICS IN S01

The model study reports a task-specific recognition gap

Evaluated task in S01	Reported metric	Study scope and finding	GFT specification implication
Base fabric construction classification	~80%	Top tested LLM setup on custom fabric-image dataset	Coarse visual sorting only; cannot certify mill specification
Detailed fabric construction classification	~55%	Same top tested setup across detailed construction factors	Near coin-flip reliability on intricate construction from images
Fabric construction image synthesis (LPIPS)	~0.2	Perceptual patch similarity for OpenAI DALL-E setup	Measures visual similarity, not yarn count, gsm or shrinkage
Cross-model complex fabric recognition	Struggled	All three tested LLMs struggled on complex weaves and designs	Route any novel fabric visual back to procurement and physical test

SOURCED [S01] Sourced from S01 (arXiv:2509.04705, submitted 2025-09-04). Accuracy and LPIPS values are approximate figures reported in the study abstract on the authors' custom fabric dataset. They describe that study's specific tasks and tested model versions, not universal model capability or factory readiness.

As **Exhibit 11**, p. 53 makes explicit, the figures in S01 are task-specific research results on a custom image dataset, not an industry benchmark or factory qualification test. [S01] Moving from base to detailed construction reduces classification accuracy from approximately 80 percent to approximately 55 percent, and perceptual similarity (LPIPS of approximately 0.2) measures visual patch resemblance rather than mechanical or thermal behaviour. [S01] What **Exhibit 11**, p. 53 cannot establish is how these or later models perform on Cameroonian cotton shirting lots, nor can any image score establish compliance with buyer wash and wear standards. [S01]

Whenever a selected concept image displays a surface texture, slub effect or drape that differs from the pre-qualified cotton shirting lots named in the brief, the designer cannot type that texture adjective into the tech pack. [S01] [S20] Instead, the concept must be constrained back to a qualified swatch from the mill library or routed as a material enquiry to procurement and laboratory measurement under Chapter 6. [S01] In our illustrative cotton-shirt example, if a render suggests a textured dobby weave whereas the brief specified plain-weave shirting, the pattern technician holds the concept until the designer confirms the plain weave or procurement supplies a measured dobby swatch with verified weight, thickness and shrinkage. [S01] [S09]

3.6

A designer must translate selected concepts into technical intent

Before any concept leaves the studio for pattern generation or technician drafting, the designer must translate the selected image into an explicit technical-flat package and a structured textual feature list. [S09] [S10] Multimodal pattern systems such as Design2GarmentCode (S09) and ChatGarment (S10) accept images, sketches or text descriptions and map them to parametric pattern programs or JSON parameter sets. [S09] [S10] Passing a raw concept render directly into a pattern generator delegates engineering interpretation to a statistical model when visual ambiguity is highest. [S09] [S20] A single front-facing concept image leaves the back view, facings, collar-stand construction, placket fold order and underarm geometry under-specified. [S09] Similarly, as the GarmentSketch benchmark (S0301) illustrates across 26,249 sketches and 21 garment categories, aligning visual sketches with accurate textual descriptions requires multi-stage refinement and human-in-the-loop verification rather than unreviewed captioning. [S0301]

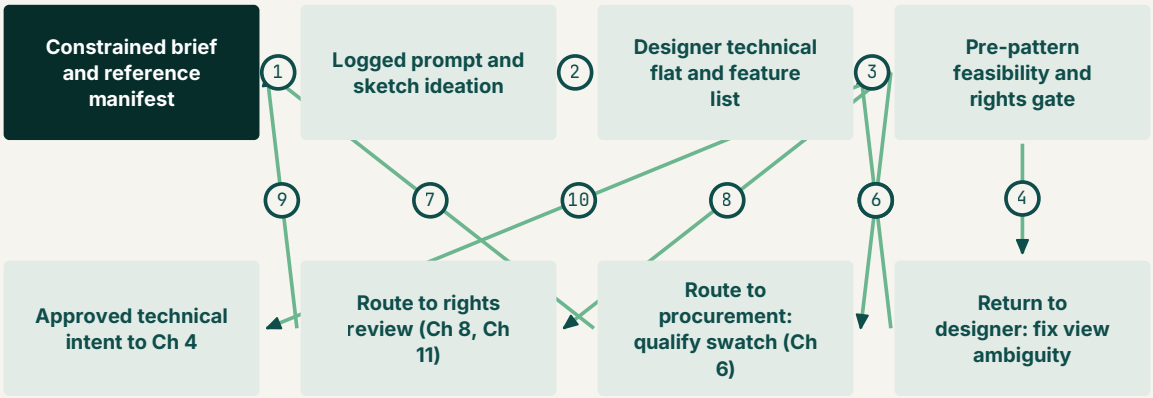
To bridge visual ideation and Chapter 4's structured pattern representations, GFT proposes a mandatory translation handoff owned by the designer. [S09] [S10] Rather than forwarding a raw diffusion image to the pattern room, the designer prepares two inspectable artefacts for every candidate concept that survives initial screening: 1. **A paired front-and-back technical flat checklist:** orthographic line drawings (human-drawn or human-corrected from an assisted sketch tool such as S07 or S0301) showing explicit seam lines, topstitching lines, dart terminations, closure orientations, collar-stand joins, cuff splits and pocket placements on both front and rear views. [S07] [S09] [S0301] 2. **A structured textual feature specification:** a field-by-field declaration mapping each visible feature of the illustrative cotton shirt to the base block's parameter vocabulary, including collar type, point length, stand height, placket width, button count, chest pocket count and dimensions, yoke presence, sleeve length, cuff style and hem shape. [S09] [S10]

During this translation step, the designer and pattern technician apply a pre-pattern rejection gate to eliminate unsewable or contradictory concepts before CAD drafting or program synthesis begins. [S09] [S20] The gate rejects any concept with ambiguous rear views (where the back yoke, pleat or hem drop contradicts the front side seam), contradictory sleeve geometry (such as a set-in sleeve crown on the front view paired with a raglan seam on the shoulder), or unsewable openings (such as pockets overlapping a dart leg, plackets terminating without a seam release, or collars lacking a constructible neckband). [S02] [S09] No empirical rejection-frequency percentage is asserted here, because rejection rates depend on prompt discipline, tool choice and product complexity, which must be measured locally during a pilot. [S19] [S20]

When a concept fails one of these checks, it must not be patched by a pattern technician guessing what the designer meant, nor discarded without recording why it failed. [S20] As shown in **Exhibit 12**, p. 55, every rejected candidate follows a typed return path to the accountable decision owner. [S20]

EXHIBIT 12 / PROPOSED GFT WORKFLOW HANDOFFS

Rejected concepts should return to the relevant decision owner



1. Approved constraints
2. Selected candidates
3. Complete dossier
4. Ambiguous geometry
5. Revised flat
6. Unqualified fabric
7. Update fabric constraint
8. Uncleared motif
9. Update permission scope
10. Cleared intent

ILLUSTRATIVE [S20] Illustrative GFT workflow for routing rejected ideation candidates back to the accountable owner before pattern synthesis begins, informed by NIST Generative AI Profile governance controls (S20). Box-and-arrow sequence represents decision ownership, not a measured flow volume.

The routing structure in **Exhibit 12**, p. 55 ensures that every failure mode discovered at the end of Chapter 3 returns to the role capable of resolving it rather than polluting Chapter 4's pattern queue. [S20] If a concept fails because the back view contradicts the front view or a pocket overlaps a seam, it returns to the designer (`design_fix`) to revise the technical flat. [S09] If it relies on an unverified weave or weight, it routes to procurement and Chapter 6 (`proc_fix`), which supplies a measured swatch or tightens the brief's fabric constraint. [S01] If it incorporates an uncleared third-party reference or a traditional cultural expression requiring community dialogue, it routes to rights review (`rights_fix`) under Chapters 8 and 11 and updates the brief's permission scope. [S21] [S23] Only candidates passing all four checks advance as approved technical intent (`ch4`). [S09] [S20] What **Exhibit 12**, p. 55 cannot show is the calendar duration of each return loop, which must be tracked in the management system. [S19]

METHOD

To audit the translation handoff from Chapter 3 to Chapter 4, verify that every promoted concept file contains: (1) the active brief ID and revision; (2) a reference manifest with zero uncleared inputs; (3) the seven-field design log for the selected render; (4) matched front and back orthographic flats; (5) a textual feature table with no blank required parameters; and (6) the pattern technician's pre-pattern gate signature confirming fixed block boundaries. [S09] [S20]

3.7

Exploration should be bounded by the cost of review

Because generating candidate images is near-instantaneous whereas engineering review by a designer and pattern technician consumes scarce skilled hours, a greenfield maker must cap the concept queue and measure review effort across both accepted and rejected candidates. [S19] [S20] A frequent error in generative design evaluations is comparing the seconds required to render an image against the hours required to draw a manual sketch, while ignoring the hours spent sifting, verifying, correcting and rejecting unbuildable variations. [S19] [S20] Where experienced pattern technicians and product developers are the primary bottleneck, flooding the review gate with unconstrained batches degrades throughput and delays approved tech packs. [S19]

GFT recommends governing ideation through a capped concept-review queue per style brief, aligned with the empirical evaluation principles of the National Institute of Standards and Technology AI Risk Management Framework (S19) and Generative AI Profile (S20). [S19] [S20] Rather than asserting a universal queue size for every factory and garment category, the product-development lead sets a bounded candidate cap at the start of a local pilot based on available technician hours, and adjusts that cap after inspecting measured pass rates and review times. [S19] Time tracking in the performance-management system must record touch time spent on all generated candidates, including those rejected at visual screening, returned at the technical-flat gate or halted for fabric and rights queries, dividing total review and translation hours by the number of styles that achieve full specification release. [S19] [S20]

WORKED EXAMPLE

Consider an illustrative workload model for one woven cotton-shirt brief to see why uncapped exploration erodes technical capacity before pattern drafting starts. Assume the following explicitly illustrative inputs for a single style brief:

- **Unbounded Case A:** the designer generates 80 raw concept images, forwards 25 images to technical screening without a cap, and translates 8 candidates into technical flats before 1 style concept is approved for Chapter 4 pattern engineering (an illustrative ratio of 80 generated to 25 screened to 8 translated to 1 approved).
- **Constrained Case B:** the brief enforces a capped queue: the designer logs at most 15 generated images within the brief's block and fabric rules, promotes at most 5 candidates to screening, and translates 2 candidates into flats and feature lists to yield 1 approved concept (an illustrative ratio of 15 generated to 5 screened to 2 translated to 1 approved).
- **Assumed touch times:** 3 illustrative minutes per raw image for logging and culling, 12 illustrative minutes per screened candidate for the review sheet (silhouette, seam plausibility, fabric check and reference manifest verification), and 45 illustrative minutes per translated candidate for preparing paired front-and-back flats, the feature list and technician pre-pattern sign-off.

The illustrative arithmetic for human review and translation effort per approved concept is:

- **Unbounded Case A:** (80 × 3 minutes) + (25 × 12 minutes) + (8 × 45 minutes) = 240 + 300 + 360 = 900 illustrative minutes, or **15.0 illustrative hours** per approved concept handoff.
- **Constrained Case B:** (15 × 3 minutes) + (5 × 12 minutes) + (2 × 45 minutes) = 45 + 60 + 90 = 195 illustrative minutes, or **3.25 illustrative hours** per approved concept handoff.
- **Illustrative difference:** Capping the queue and filtering against a manufacturable brief saves 900 - 195 = 705 illustrative minutes (**11.75 illustrative hours**) of skilled review and translation time per approved style before Chapter 4 pattern generation begins.

These figures are hypothetical scenario inputs constructed solely to demonstrate the queue-cost accounting method; they are not observed GFT factory times or a forecast of savings. [S19] If Constrained Case B requires a second iteration of the capped queue, effort doubles to 195 × 2 = 390 illustrative minutes (**6.50 illustrative hours**), remaining below Unbounded Case A's 15.0 illustrative hours while keeping every failure reason logged. [S19] [S20]

This queue accounting connects Chapter 3 to the rest of the paper and the wider GFT library. [S19] [S20] Only technical intent passing the capped pre-pattern gate transfers to Chapter 4, where panel geometry, parametric programs (S02, S09), JSON parameters (S10) and deterministic validators are examined. [S02] [S09] [S10] Any candidate whose appeal rests on a surface repeat motif or cultural reference routes to Chapter 8 for repeat-tile geometry, colour proof and community-permission checks before release. [S23] [S26] Finally, all recorded designer and technician hours spent on selected and rejected candidates in Chapter 3 flow into Chapter 12's economic model and

the `sampling-cost` tool as incremental review cost, ensuring that financial evaluation counts total effort per accepted style rather than celebrating cheap image generation in isolation. [S19] [S20]

A prompt widens what the studio can imagine, but only a constrained brief defines what the factory may build.

CHAPTER SOURCE NOTES

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S01 Authors via arXiv. 2025-09-04.

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

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

S20 National Institute of Standards and Technology. 2024-07-26.

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

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S13 Authors via arXiv; SIGGRAPH 2026. 2026-04-09.

S26 Authors via arXiv. 2025-04-17.

S19 National Institute of Standards and Technology. 2023-01-26.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S21 U.S. Copyright Office. 2025-01-29.

S07 Authors via arXiv. 2024-08-01.

S23 World Intellectual Property Organization. 2026.

S24 World Intellectual Property Organization. 2026.

S0302 Authors via arXiv. 2025-03-03.

S0301 Authors via arXiv. 2026-06-12.

S04 Authors via arXiv. 2024-01-29.

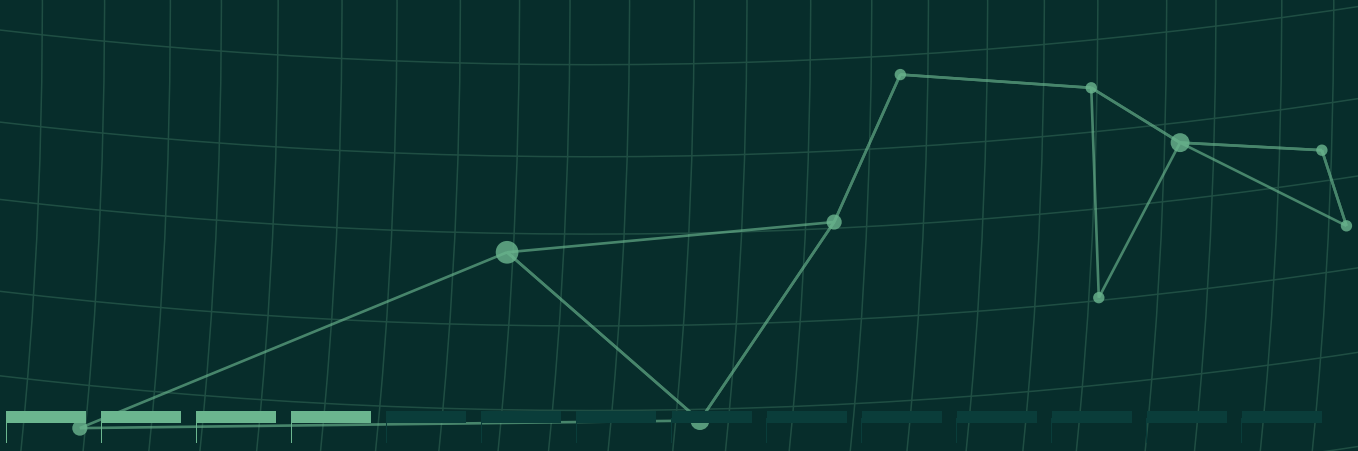
R31 arXiv 2509.04705. Publication date unrecorded; inherited source-text check.

04

PART II: SPECIFICATIONS REQUIRE GEOMETRY, MATERIALS AND FIT / FROM INTENT TO PANELS

Editable patterns require geometry and seam relationships

A structured representation makes errors inspectable, but validity of a representation does not guarantee a good garment.



CHAPTER 04
DECISION TEST

A structured representation makes errors inspectable, but validity of a representation does not guarantee a good garment.

A greenfield apparel manufacturer should accept generative pattern assistance only when the system emits an explicit, editable two-dimensional geometry and seam-relationship record rather than a flat picture or an unindexed three-dimensional mesh. [S02] [S09] In Cameroon, where approximately 95 percent of raw cotton fibre is exported unprocessed and roughly 5 percent undergoes local transformation against a 50 percent target by 2030, capturing value in domestic cut-make-trim plants depends on converting creative briefs into producible cutting instructions without wasting cloth or sample-room time. [R06] [R07] An approved concept from Chapter 3 defines visual intent and permitted feature changes for an illustrative woven cotton shirt, yet no spreading table or cutter acts on pixels alone. [S09] [S22] Between the selected concept and the pattern technician's computer-aided design (CAD) workstation sits a representation choice: how panels, boundary curves, seam pairings, grain directions and dimensional parameters are encoded, validated and edited before cutting. [S02] [S14] Published research from 2021 to 2026 represents sewing patterns as domain-specific programs, JSON parameter files, discrete token sequences, explicit specification languages or implicit signed-distance fields. [S02] [S08] [S10] [S14] [S0401] [S0403] Because each representation exposes different control points and fails differently when parameters conflict, GFT recommends placing a deterministic validation and technician-editing boundary between model inference and CAD release. [S16] [S20] [S22]

4.1

Panel geometry is a different object from a garment image

A garment image records projected light, shading and folds from one viewpoint, whereas a sewing pattern defines the intrinsic two-dimensional rest shape of flat fabric pieces and the topological rules joining their boundaries into a three-dimensional shell. [S02] [S03] [S0402] When an image generator renders a cotton shirt on a figure, it resolves occluded rear panels, yokes, collar stands and underarm seams by visual plausibility rather than planar area conservation. [S03] [S09] Receiving only that rendered image leaves the technician to draft every panel boundary by hand, place darts or seams, assign seam allowances and verify that opposing stitching lines match within tolerances. [S02] [S22] Treating image synthesis and pattern generation as interchangeable confuses visual ideation with geometric engineering. [S09] [S17]

4.2 PARAMETRIC COMPONENTS SEPARATE DESIGN INTENT FROM VERTEX COORDINATES

Computational pattern research models garments as structured two-dimensional panels rather than pixel arrays or deformed meshes. [S02] [S0401] [S0402] Early synthetic dataset generation introduced parametric templates across 19 base garment categories to produce over 20,000 design variations paired with three-dimensional draped models, establishing flat panel geometry as an essential prior. [S0401] NeuralTailor then framed pattern recovery from three-dimensional point clouds as set regression with variable cardinality, using point-level attention to disentangle flat panel shapes and stitch links from deformations caused by gravity, body contact and fabric properties. [S0402] In 2023, Sewformer extended structured recovery to single

two-dimensional photos via a two-level Transformer trained on the synthetic Sew-Factory dataset of roughly 1 million image-and-pattern pairs, predicting parametric panel shapes alongside stitch relationships. [S03]

While reconstruction models infer patterns from visual observations, programmatic systems construct patterns forward from explicit rules. [S02] GarmentCode introduced an object-oriented domain-specific language (DSL) organising sewing patterns into hierarchical, interchangeable components. [S02] Instead of dragging vertices on a flat canvas, a GarmentCode program defines modular sub-assemblies such as a bodice, sleeve, collar or skirt, and computes two-dimensional contours algorithmically from design parameters and body measurements. [S02] That separation matters in product development because changing a high-level dimension, such as neck circumference or armhole depth, must propagate consistently across connected boundaries without breaking loop closure or creating orphan stitch links. [S02] [S30]

4.3 A FACTORY PATTERN RECORD REQUIRES OPERATIONAL METADATA BEYOND GEOMETRY

Academic representations encode panel outlines as ordered straight segments, Bezier curves or signed distance fields in a two-dimensional plane, paired with stitch correspondences for simulation. [S02] [S03] [S0403] A simulation-ready research file remains narrower than a factory pattern record. [S14] [S22] In GFT's Cameroon textile programme base case, centred on cut, make, trim and finishing operations with optional partnered upstream supply, every pattern piece entering the sample room must state cut quantity, warp and weft alignment, fabric or interlining assignment, and notch and seam-allowance placement relative to the net stitching line. [R06] [S22]

To bridge research geometry and sample-room execution, GFT proposes an illustrative panel identity schema for the running cotton-shirt example ({{xr:wp10-ch04-pattern-record-schema}}). [S02] [S14] [S22] Each piece carries a stable panel identifier, component role, cut quantity, mirror rule, explicit warp grainline vector, material and interlining link, and indexed boundary edges separating net sewing lines from outer cutting lines. [S02] [S22] These cutting-room fields wrap the geometric output so technicians inspect a complete specification candidate rather than bare coordinate arrays. [S19] [S22]

EXHIBIT 13 / PROPOSED SCHEMA FIELDS

A usable pattern record carries more than coordinates

Field group	Proposed attribute	Illustrative shirt value	Downstream check
Panel identity	panel_id / role	SHIRT-01-FRONT-L (Left front bodice)	Unique ID per piece; no orphan panels
Cutting rule	cut_qty / mirror_pair	1 self, mirror with SHIRT-01-FRONT-R	Left-right symmetry or declared placket asymmetry
Coordinate basis	unit / closed_loop	millimetres; tolerance 0.1 mm	Unit readback in CAD; zero open contours
Fabric alignment	grainline_vector	0 deg to warp selvedge	Prevent off-grain twist during spreading
Boundary indexing	edge_id / curve_type	E01..E07 (line, cubic Bezier)	Every stitchable edge addressable by ID
Seam mapping	seam_pair_id / allowance	SP-SHOULDER-L; 10 mm net-to-cut	Pass to Chapter 5 topology and Chapter 9 DXF
Material link	material_id / interlining	MAT-COT-POPLIN-01; collar fusible	Pass to Chapter 6 physical swatch record

ILLUSTRATIVE [S02] [S14] [S22] Illustrative GFT teaching schema for the running cotton-shirt example. Geometric and topology concepts draw on S02 and S14; CAD exchange fields draw on S22. Field codes are a GFT design proposal, not a standard adopted by those authors.

The schema in {{xr:wp10-ch04-pattern-record-schema}} clarifies what a geometric representation can and cannot establish. [S02] [S14] [S22] It organises the metadata required before CAD export, ensuring coordinates travel with unit declarations, grain orientation and edge identifiers. [S22] What the table cannot establish is whether the dimensions fit a target wearer group or whether the cotton poplin drapes without puckering along the armhole; those questions require seam-assembly verification in Chapter 5, physical material testing in Chapter 6 and size-aware fit trials in Chapter 7. [S12] [S27]

4.4

Programs and parameter files expose different edit controls

Technical teams evaluating structured pattern generation must choose among intermediate representations published between 2023 and 2026. [S02] [S08] [S09] [S10] [S14] [S16] Multimodal models do not natively output cutting files; they emit an intermediate representation that a decoder or compiler translates into flat two-dimensional panels. [S09] [S10] [S14] That intermediate structure determines whether a technician can inspect, edit and repair the output when a model misinterprets a brief. [S09] [S15]

4.5 FOUR REPRESENTATION FAMILIES TRADE EXPRESSIVE FREEDOM AGAINST DETERMINISTIC CONTROL

Across the opened literature, four representation strategies connect multimodal inputs to two-dimensional sewing patterns, alongside implicit fields used for differentiable meshing. [S08] [S09] [S10] [S14] [S16] [S0403] [S0404]

1. **Program synthesis in a domain-specific language.** Design2GarmentCode uses large multimodal models to translate images, text or sketches into parametric programs in the GarmentCode DSL. [S02] [S09] Because the model emits modular program calls rather than raw vertices, the deterministic GarmentCode compiler executes the script to construct size-precise panels and stitch rules. [S02] [S09] DressAgent similarly builds its inverse design loop on an enhanced GarmentCode representation covering over 15,000 outfits, using an agentic Tailor, Critic and Orchestrator trio to refine parameters against simulation. [S18]
2. **Structured JSON descriptions and continuous parameter vectors.** ChatGarment fine-tunes a vision-language model to emit a JSON file containing style descriptions and numerical attributes. [S10] That JSON record drives a refined GarmentCode model to produce two-dimensional sewing patterns draped on three-dimensional bodies, restricting the network output to a parameter dictionary consumed by a deterministic generator. [S10]
3. **Discrete sewing-pattern token sequences.** Alpparel fine-tunes large multimodal models on over 120,000 garments annotated with text, images and patterns, encoding panel contours and stitch instructions into concise token sequences for autoregressive prediction and editing. [S08] DressCode earlier introduced SewingGPT, using cross-attention and text-conditioned embeddings to generate sewing patterns autoregressively alongside textures. [S04]
4. **Explicit specification languages, hierarchical schemas and implicit fields.** Because flat token sequences entangle topology with curve parameters and fixed templates limit variety, two 2026 frameworks introduced stricter abstractions. [S14] [S16] PatternGSL defines a template-free specification language encoding panel boundaries, parameterised seams and stitch topology, trained on the 300,000-sample PatternGSLData dataset with deterministic validity handling during decoding. [S14] GarmentWeaver constructs compact hierarchical targets by activating garment-relevant structural branches and applies feasibility-aware regularisation to encourage executable patterns. [S16] In parallel, implicit-field methods such as ISP and latent flow matching represent panel boundaries via signed distance fields and seam endpoints via unsigned distance fields in continuous latent space, extracting edges for stitch prediction. [S0403] [S0404]

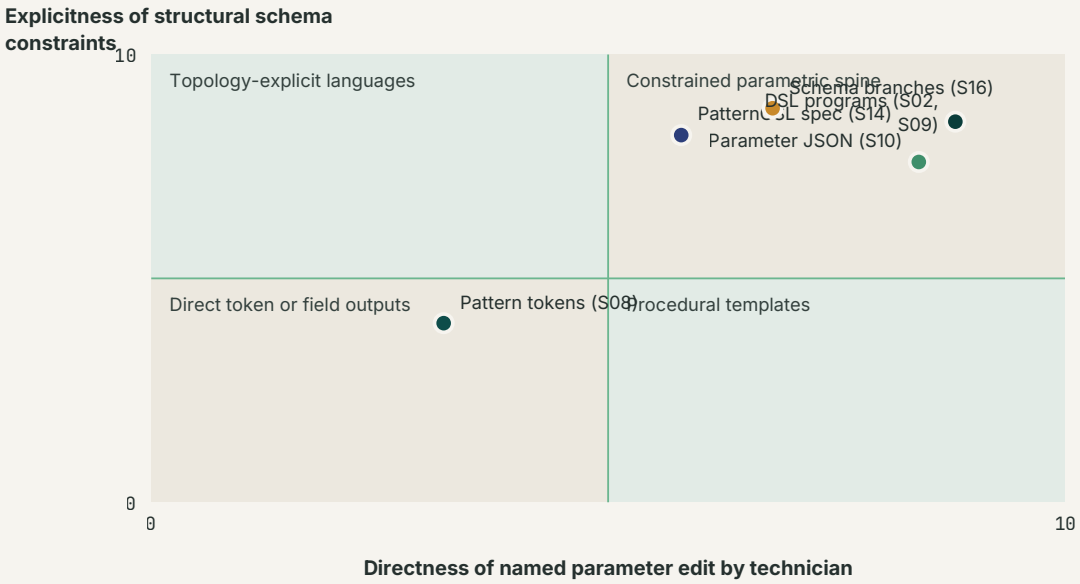
4.6 REPRESENTATION SELECTION SHOULD PRIORITISE TECHNICIAN EDITABILITY AND DETERMINISTIC VALIDATION

A greenfield manufacturer should evaluate these families by how well they support human engineering control rather than by academic leaderboard scores ({{xr:wp10-ch04-representation-matrix}}). [S19] [S20] GFT proposes three selection criteria before any representation enters a pilot:

- **Technician editability:** Can a technician inspect the intermediate artefact and adjust one dimension, such as lengthening our illustrative shirt's cuff by 15 millimetres or widening the yoke by 10 millimetres, without regenerating the whole garment and risking collar or sleeve drift? [S09] [S10] [S15]
- **Deterministic pre-simulation validation:** Does the representation pass through a rule-based compiler or schema checker catching syntax errors, missing parameters and unclosed contours before simulation or CAD import? [S14] [S16] [S20]
- **CAD export fidelity:** Do decoded panels, seam lines, grainlines and notches map cleanly into industrial exchange formats such as AAMA DXF or ASTM DXF for receiving-CAD readback? [S22]

EXHIBIT 14 / ILLUSTRATIVE CONTROL SCORE (0 TO 10)

Representation choice should follow the required edit



ILLUSTRATIVE [S02] [S08] [S09] [S10] [S14] [S16] Illustrative GFT placement of published representation families on two qualitative engineering axes: directness of named parameter editing by a technician (horizontal) and explicitness of built-in structural schema constraints before CAD export (vertical). Mechanism definitions are sourced from S02, S08, S09, S10, S14 and S16; coordinate positions are illustrative GFT analytical judgements for workflow design, not empirical benchmark measurements.

As {{xr:wp10-ch04-representation-matrix}} illustrates, representations exposing named parameters or hierarchical schema branches ([^S02], [^S09], [^S10], [^S16]) give technicians stronger control than flat token sequences that must be decoded before their geometry is visible ([^S08]). [S14] [S16] Still, {{xr:wp10-ch04-representation-matrix}} classifies edit surfaces rather than model accuracy; it cannot prove that a DSL or JSON pipeline will succeed on a Cameroonian buyer's specification until tested locally under Chapter 13. [S19] [S20]

GFT NOTE

For a greenfield maker establishing its first digital pattern workflow, GFT recommends prioritising representations emitting inspectable parameter tables or schema-checked component programs ([^S02], [^S09], [^S10], [^S16]) over opaque coordinate regressions. When a collar stand or placket fails a dimensional check, the technician must edit an explicit parameter value and recompile deterministically rather than re-prompting a neural network. [S15] [S20]

4.7

Parameter constraints should be explicit and testable

Valid syntax in an intermediate representation does not guarantee a geometrically valid or sewable pattern. [S14] [S16] [S30] A JSON file can parse cleanly while specifying a cuff wider than its sleeve opening, and curve tokens can decode into a self-crossing polygon. [S16] [S30] GarmentWeaver's authors note that flat long sequences entangle structure with detailed parameters, causing redundant components, inaccurate local details and poor simulation compatibility. [S16] Similarly, GarmentCodeData's creators report that incompatible parameter combinations remain a common pitfall of complex parametric systems, including their own. [S30]

4.8 STRUCTURAL LANGUAGES AND FEASIBILITY REGULARISATION TARGET INVALID GEOMETRY

Recent research addresses these failures during training and inference via hierarchical schemas and regularisation ([^S16]), and during decoding via deterministic validity handling ([^S14]). GarmentWeaver activates only structural branches relevant to the garment category and applies feasibility-aware regularisation to penalise invalid predictions. [S16] PatternGSL encodes panel boundaries, parameterised seams and stitch topology in a standardised language and applies lightweight deterministic validity handling during decoding, preserving simulation readiness without expensive optimisation cleanup. [S14]

Even when a decoder enforces closure, industrial pattern engineering requires auditable invariants before a file enters CAD. [S20] [S22] Allowing an invalid panel into a 3D simulator or 2D marker queue wastes compute and technician hours on defects that computational geometry rules catch in milliseconds. [S20] [S30]

4.9 DETERMINISTIC PRE-CAD RULES CATCH FOUR RECURRING GEOMETRY FAILURES

Drawing on GarmentCodeData's exclusions ([^S30]) and DXF exchange parameters ([^S22]), GFT proposes a mandatory deterministic validation gate between pattern generation and technician review ({{xr:wp10-ch04-geometry-failures}}). Before any generated cotton-shirt pattern is opened for refinement, an automated validator should check six rules:

1. **Positive enclosed panel area:** Every panel must enclose a positive planar area above an illustrative minimum (1,500 square millimetres for a collar tab), excluding empty or belt-only outputs. [S30]

2. **Closed boundary loops:** Ordered boundary segments must close within an illustrative 0.05 millimetre tolerance so open gaps never block cutting-line offset generation in CAD. [S14] [S22]
3. **No edge self-intersections:** Non-adjacent boundary segments must not cross in the two-dimensional plane, eliminating bow-tie inversions from incompatible curve tangents or dart depths. [S30]
4. **Unit and bounding-box sanity:** Coordinates must declare a linear unit (`mm`, `cm` or `inch`), and bounding boxes must lie within category bounds so an illustrative 74-centimetre shirt front is never imported as 74 millimetres or 74 inches. [S22] [S30]
5. **Component cardinality and pairing:** A symmetric two-sleeve shirt brief must not emit duplicate right sleeves, omit the left front panel or leave a collar fall without its stand. [S16] [S30]
6. **Cross-component ratio bounds with expert exceptions:** Dependent dimensions must satisfy construction inequalities, such as cuff length below sleeve length ([^S30]) and collar-stand neckline length matching the neck opening within an illustrative plus or minus 2.0 millimetres, unless a technician flags an intentional gathering, pleating or ease exception. [S12] [S27] [S30]

EXHIBIT 15 / ILLUSTRATIVE SEVERITY INDEX (0 TO 100)

The same concept can conceal different geometry failures



ILLUSTRATIVE [S22] [S30] Illustrative comparison of four structural failure modes across five pre-CAD check dimensions, grounded in the quality-control exclusions of S30 (self-intersecting panels, empty/trivial components, out-of-bound lengths) and CAD unit/contour requirements of S22. Severity values (0, 50, 100) are illustrative rule-trigger indicators showing which check catches each defect, not observed defect frequencies.

The profile in {{xr:wp10-ch04-geometry-failures}} shows why one validity flag cannot diagnose failures. [S20] [S22] [S30] An open contour fails loop closure while passing unit-scale and part-count checks; centimetre-versus-millimetre confusion produces closed polygons ten times too small in CAD; and a duplicate sleeve passes single-panel tests while corrupting the bill of materials. [S16] [S22] [S30] What {{xr:wp10-ch04-geometry-failures}} does not provide is an empirical defect rate; failure shares must be measured in a local pilot. [S19] [S20]

DETERMINISTIC PRE-CAD GEOMETRY CHECK PROTOCOL

GFT proposes a five-stage check on every generated pattern revision before CAD opening: 1. **Schema and unit parse:** Confirm `unit`, `panel_id`, `grainline_vector` and `cut_qty` exist and convert coordinates into millimetres ([^S22]). 2. **Planar topology test:** Verify vertex closure within an illustrative 0.05 millimetre threshold and sweep contours for intersections ([^S30]). 3. **Bounding-box and area envelope:** Check that panel area and dimensions lie inside the product family's envelope ([^S30]). 4. **Cardinality reconciliation:** Compare component roles and cut quantities against the Chapter 3 brief to block missing or duplicated panels ([^S16]). 5. **Exception routing:** Flag cross-edge length mismatches above an illustrative 2.0 millimetre threshold; block release unless a technician attaches a signed exception code (such as `EX-EASE-SLEEVE-CROWN`) for Chapter 5 ([^S12], [^S20]).

4.10

Template reuse and template-free research solve different needs

A recurring question in the 2023 to 2026 literature is whether generative garment systems should rely on parametric templates or operate in a template-free regime. [S02] [S14] In computer vision, *template dependence* limits open-world reconstruction: when recovering arbitrary clothing from street photographs, a fixed library of programmatic blocks cannot represent novel cut-outs, drapery or seams absent from the library. [S03] [S14] PatternGSL targets this gap, preserving pattern-based rigour while removing template dependence across 300,000 training examples. [S14]

4.11 INDUSTRIAL MANUFACTURING RELIES ON APPROVED BLOCKS FOR REPEATABILITY

What solves open-world reconstruction in computer vision does not automatically solve quality control in a greenfield garment factory. [S19] [S20] In industrial production, a base block (or sloper) is an accumulated record of verified fit, tested grading increments, sewing-line balance and buyer-approved tolerances. [S02] [S15] [S22] When a plant manufactures woven cotton shirts, uniforms or school garments, starting every style from an unconstrained, template-free coordinate prediction introduces unnecessary variance in armhole scye depth, shoulder slope and neck-drop balance. [S02] [S27]

Even within the literature, systems oriented toward controllable authoring and multi-layer refinement build on structured component libraries. [S02] [S09] [S10] [S18] GarmentCode's component abstraction lets users swap, combine and parameterise modular elements while maintaining validity. [S02] ChatGarment refines GarmentCode to translate vision-language JSON parameters into valid patterns ([^S10]), and DressAgent extends GarmentCode to recover multi-layer outfits with closed-loop simulation refinement ([^S18]). TailorCoPilot's version-control backend, TailorTrace, captures expert practice by recording discrete state transformations and operation sequences over panels, edges, vertices and stitches inside a conventional interface. [S15]

4.12 A GREENFIELD MAKER SHOULD START WITH CONTROLLED VARIANTS OF KNOWN BLOCKS

For GFT's staged Cameroon programme, distinguishing template modification from template-free synthesis yields a clear rule. [R06] [S02] [S14] During plant launch and qualification, the factory should anchor its workflow in an owned, technician-verified library of base blocks for its core product families, including the illustrative woven cotton shirt. [S02] [S15] Generative tools should first map a constrained brief from Chapter 3 onto controlled parametric variations of those blocks, such as altering collar point length, switching between single-piece and two-piece yokes, adjusting sleeve length or resizing a chest pocket. [S09] [S10]

Template-free generation ([^S14]) and schema-aware synthesis ([^S16]) remain valuable to test in bounded pilots where a designer requests a silhouette outside existing blocks, but they should never silently overwrite an approved block. [S14] [S16] [S20] Whenever a technician accepts a departure beyond an approved block's parameters, the pattern record should log an exception reason, model version and physical fit-sample gate before cutting is authorised. [S15] [S19] [S20]

4.13

Synthetic training data carries design assumptions

Learning-based pattern generators and stitch predictors depend on synthetic datasets because paired industrial archives of garment photographs, graded DXF patterns, measured fabric physics and multi-size fit approvals are rarely public. [S03] [S08] [S14] [S30] [S0401] Understanding how synthetic corpora are constructed and filtered is essential when judging whether a published model will generalise to local garments and wearers. [S19] [S20] [S30]

4.14 SYNTHETIC PIPELINES IMPOSE SAMPLING PRIORS AND EXCLUDE FAILED DESIGNS

Public datasets illustrate both the scale and built-in assumptions of synthetic garment data: the 2021 NeurIPS dataset generated over 20,000 variations from 19 base types with 7 held-out topologies ([^S0401]); SewFactory synthesised roughly 1 million image-and-pattern pairs ([^S03]); GarmentCodeData constructed 115,000 samples across three textile presets fitted to a neutral body and 5,000 sampled body shapes ([^S30]); AIpparel curated over 120,000 multimodal garments ([^S08]); PatternGSLData generated 300,000 image-to-GSL pairs ([^S14]); AutoSew annotated over 18,000 multi-edge patterns ([^S12]); and DressAgent built over 15,000 multi-layer outfits ([^S18]).

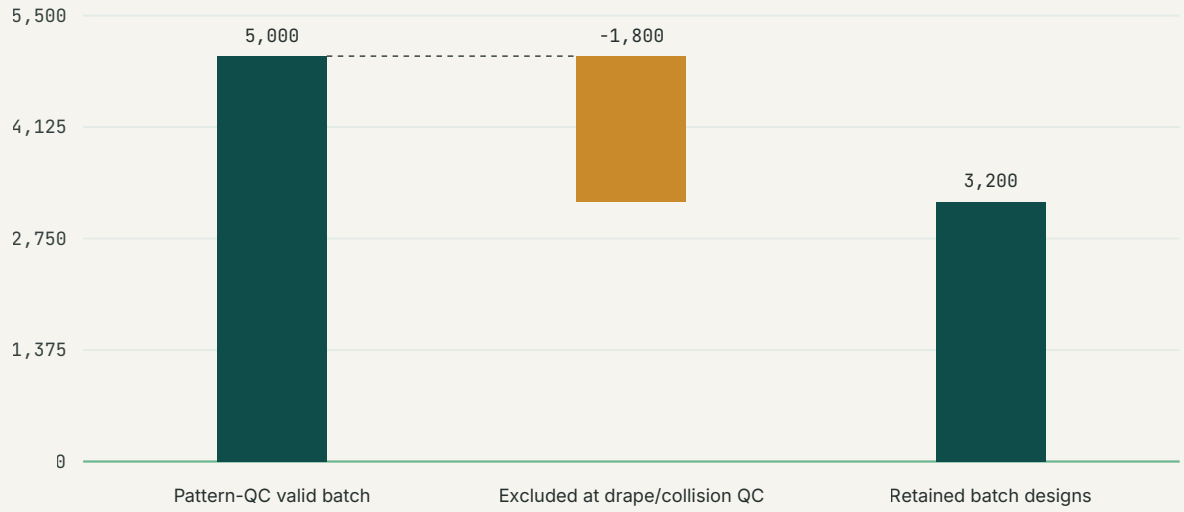
Because the full text of GarmentCodeData ([^S30]) discloses its generation and quality-control mechanics, it reveals what synthetic training data contains. First, naive uniform parameter sampling produced non-representative garments, such as skirts or sleeves longer than body dimensions. [S30] To prevent this, the authors aligned parameter bounds with body measurements and introduced manual coin-flip probabilities p_i to favour daily-wear conventions, dividing the corpus into roughly one third upper-body, one third lower-body and one third full-body garments, and setting $p_i = 0.8$ on neckline and sleeve asymmetry to favour left-right symmetry. [S30]

Second, the body shapes were generated by registering a template mesh to the European subset of the CAESAR database (1,700 three-dimensional scans), applying principal component analysis (PCA) to sample 5,000 body shapes, and computing 25 body measurements automatically along balance planes. [S30] Compared against manual tailor measurements in CAESAR, automatic convex-hull algorithms overestimated bust circumference by 2.81 centimetres, hip circumference by 1.99 centimetres and leg circumference by 3.49 centimetres, while underestimating shoulder width by 4.57 centimetres by measuring across shoulder ball joints rather than deltoid muscles. [S30]

Third, the pipeline relied on two quality-control filtering stages ({{xr:wp10-ch04-synthetic-filtering-waterfall}}). [S30] At the two-dimensional pattern stage, the authors excluded samples with self-intersecting panels, full-body lengths exceeding floor length, empty patterns, belt-only outputs, and lower garments prone to sliding off. [S30] Passing patterns were generated in batches of 5,000 designs (10,000 patterns per batch across random and neutral bodies) and draped in an Extended Position-Based Dynamics (XPBD) simulator with temporary half-space constraints, capped at 5 minutes and 2,400 frames per drape. [S30] Drapes failing static equilibrium or showing high cloth-body or self-collision intersections were discarded: out of each batch of 5,000 simulated drapes, approximately 3,200 final designs passed simulation quality control, a batch retention of 64 percent (and an overall reported simulation success rate of about 72 percent). [S30]

EXHIBIT 16 / DESIGNS PER 5,000-SAMPLE BATCH

Quality filtering in GarmentCodeData removes about 36 percent of simulated drapes per batch



DERIVED [S30] Derived by GFT arithmetic from Section 3.4 of S30: each generation batch comprises 5,000 design samples that have already passed 2D pattern quality checks; after XPBD draping and collision/equilibrium quality control, approximately 3,200 final designs pass per batch. Derived exclusion = 5,000 - 3,200 = 1,800 designs (36 percent of the 5,000 batch). Pattern-stage rejections prior to the 5,000 batch are discarded continuously until 5,000 valid patterns are collected and are not quantified as a fixed count in S30.

The reconciliation in {{xr:wp10-ch04-synthetic-filtering-waterfall}} carries an operational warning. [S30] Training only on the surviving 3,200 designs per batch that passed two-dimensional geometric filtering and three-dimensional simulation filtering yields a curated survivor distribution rather than unconstrained parametric sampling. [S20] [S30] Moreover, {{xr:wp10-ch04-synthetic-filtering-waterfall}} reflects only simulation-stage attrition inside the authors' pipeline; it cannot predict how often a generative model emits invalid geometry when prompted with local Cameroonian styles, non-European body proportions or fabrics outside the three presets in GarmentCodeData. [S19] [S30]

WORKED EXAMPLE: ILLUSTRATIVE BATCH-REVIEW AND GEOMETRY-FILTERING ARITHMETIC

Consider an illustrative local evaluation testing an assisted pattern generator across 50 woven shirt candidates before manual CAD refinement:

- Candidate queue: an illustrative batch of 50 shirt candidates submitted to generation.
- Deterministic pre-CAD gate: an illustrative 16 percent (8 patterns: 3 unclosed collar loops, 3 self-intersecting armholes and 2 neck-ratio mismatches) fail automatic checks in under 1 minute.
- Candidates entering review: an illustrative $50 - 8 = 42$ valid patterns.
- Technician effort with screening: at an illustrative 25 minutes per valid pattern, review takes $42 \times 25 = 1,050$ minutes (17.5 hours).
- Counterfactual without screening: if those 8 defective files entered CAD and took an illustrative 45 minutes each for diagnosis, they would add $8 \times 45 = 360$ wasted minutes (6.0 hours), raising review time to 23.5 hours (an illustrative 34.3 percent increase).

This illustrative case shows how deterministic checks protect technician hours ([^S14], [^S20], [^S30]). All figures are illustrative scenario inputs, not observed GFT results.

4.15 HELD-OUT LOCAL TESTING MUST SEPARATE INTERPOLATION FROM EXTRAPOLATION

Because synthetic corpora encode specific body-shape priors (the European CAESAR subset in [^S30]), strong symmetry preferences ($p_i = 0.8$ in [^S30]) and three textile presets, a greenfield maker must test distribution shift explicitly. [S19] [S20] [S30] GFT recommends splitting any local test suite into two buckets:

- **In-distribution interpolation cases:** Symmetric woven tops and shirts mirroring standard bodice, sleeve and collar templates in public synthetic datasets. [S02] [S30]
- **Held-out structural and anthropometric variants:** Styles with deliberate asymmetry (offset pocket, concealed placket or asymmetric neckline), regional proportions, and consented local wearer measurements from Chapter 7 rather than European CAESAR PCA components. [S19] [S20] [S30]

Reporting accuracy separately on these buckets prevents strong performance on symmetric shirts from masking failures on asymmetric or locally proportioned garments. [S19] [S20]

4.16

Generated code needs a constrained execution path

When a system produces sewing patterns via program synthesis ([^S09], [^S18]), JSON parameter files ([^S10]) or specification languages ([^S14]), software must parse and execute that output to build two-dimensional geometry and export a CAD file. [S02] [S22] The NIST Generative Artificial Intelligence Profile ([^S20]) notes that generative models can emit syntactically plausible but hazardous or unverified outputs, recommending capability checks, input/output validation and lifecycle monitoring ([^S19], [^S20]).

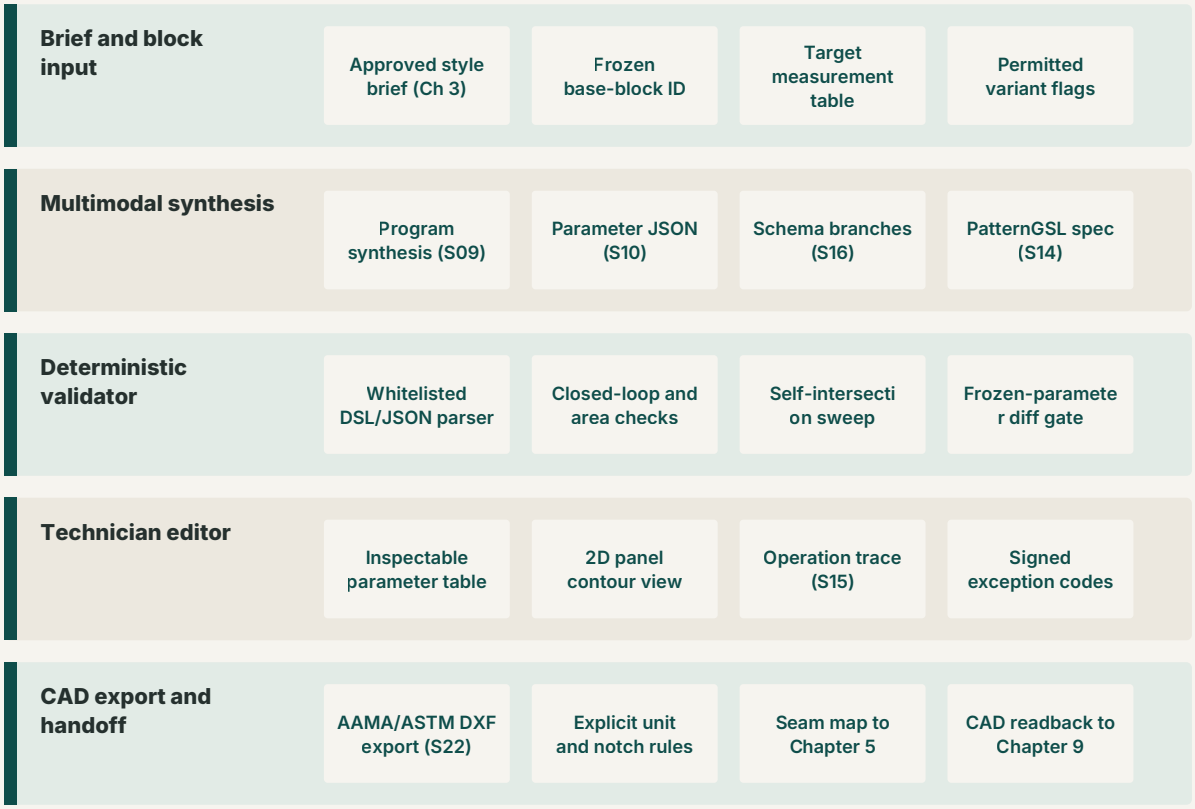
4.17 PROGRAM SYNTHESIS AND PARAMETER DECODING REQUIRE A SANDBOXED GEOMETRY WORKER

In a program-synthesis pipeline such as Design2GarmentCode ([^S09]), the multimodal model writes executable Python code targeting the GarmentCode library ([^S02]). Executing model-generated scripts directly in the technician's primary environment with file-system and network privileges creates two risks. [S20] First, an unconstrained interpreter could run arbitrary system calls, overwrite approved blocks or exfiltrate tech-pack records. [S20] Second, generated code could import unapproved helpers, alter unit settings, enter an infinite loop during curve subdivision or silently modify parameters on components marked as frozen. [S02] [S20]

To contain both failure modes, GFT proposes a deterministic architectural boundary between the generative model and the factory's CAD environment ({{xr:wp10-ch04-deterministic-boundary}}). [S09] [S10] [S14] [S16] [S20] [S22] The model never interacts directly with the pattern database or CAD application; outputs enter an isolated, network-disabled geometry worker restricted to a whitelist of declarative primitives and bounded schemas. [S10] [S14] [S20]

EXHIBIT 17 / SYSTEM RESPONSIBILITY LAYERS

A deterministic boundary belongs between generation and CAD



ILLUSTRATIVE [S09] [S10] [S14] [S16] [S22] Illustrative GFT five-layer architecture separating probabilistic model generation from deterministic validation, technician editing and CAD exchange. Informed by program/JSON/schema generation in S09, S10, S14 and S16, risk controls in S20 and DXF exchange options in S22. This is a proposed engineering control architecture, not an installed GFT software deployment or a security certification.

The five-layer separation in {{xr:wp10-ch04-deterministic-boundary}} ensures probabilistic inference stops at the second layer. [S09] [S10] [S14] [S16] The third layer acts as an airlock: it parses only schema-compliant JSON dictionaries ([^S10]) or syntax trees containing whitelisted GarmentCode or PatternGSL constructors ([^S02], [^S09], [^S14]), enforces timeouts, runs the geometric validators from Section 4.3, and blocks unrequested parameter changes against the frozen base block. [S20] Only patterns passing this validator reach the technician's editor and DXF export utility. [S15] [S22] While {{xr:wp10-ch04-deterministic-boundary}} provides an integration blueprint, it is an illustrative GFT proposal rather than a security certification or a substitute for receiving-CAD readback tests in Chapter 9. [S20] [S22]

4.18 EXECUTION LOGS MUST RECORD EVERY PARAMETER DIFF AGAINST THE BASE BLOCK

The validator layer should also write an immutable execution log for every candidate. [S15] [S20] When the illustrative cotton shirt starts from an approved base block and the Chapter 3 brief permits only collar-point and chest-pocket modifications, the geometry worker computes a parameter diff between the base block and the generated specification. [S02] [S15] If the model also altered chest width by an illustrative 12 millimetres or sleeve cap height by 8 millimetres without authorisation, the diff

gate rejects the candidate before technician inspection. [S20] Retaining the model identifier, weight version, brief hash, emitted code or JSON, validator flags and parameter diff log aligns the workflow with companion paper wp14. [S20]

> **UNVERIFIED CODE REPOSITORIES AND SECURITY BOUNDARIES**
>
> Several research abstracts link to project pages or state that code and data are or will be released ([^S02], [^S03], [^S08], [^S10], [^S14]). A public link in a preprint is not evidence of production packaging, dependency stability, licence compatibility or safe execution against untrusted inputs ([^S20]). Before running external code in a pilot, the team must inspect the licence, pin dependency hashes, disable outbound network calls inside the geometry worker and verify that no proprietary blocks reach external interfaces ([^S19], [^S20]).

4.19

The output must remain editable by the technician

Even when a generated pattern passes every syntax, closure, self-intersection and parameter-diff check, it remains a candidate specification rather than an authorised cutting instruction. [S19] [S22] Computational validity proves only that polygons close and parameters lie within mathematical bounds; it does not prove that the collar roll sits cleanly, that armhole pitch balances movement, or that seam allowances match factory folders. [S02] [S22] [S27] Therefore, pattern output must arrive in two linked forms that a technician can inspect and modify: a structured intermediate parameter or operation view, and an industrial 2D CAD exchange specimen. [S09] [S15] [S22]

4.20 INTERMEDIATE OPERATION LOGS AND INDUSTRIAL DXF OPTIONS SERVE COMPLEMENTARY ROLES

At the intermediate level, TailorCoPilot ([^S15]) shows the value of representing pattern making as discrete, version-controlled state transformations over panels, edges, vertices and stitches. When a technician adjusts a generated shirt pattern, recording the operation sequence and edit reason turns a one-off repair into a reusable training and governance record, developed in Chapter 10. [S15]

At the CAD exchange level, the pattern must cross into industrial pattern-making and marker-making software. [S22] Vendor documentation for CLO's two-dimensional pattern DXF import and export ([^S22]) illustrates the settings governing exchange survival:

- **Exchange format conventions:** Tools distinguish between Standard DXF, AAMA DXF and ASTM DXF formats, each with specific layer and annotation rules. [S22]
- **Measurement units and scale:** Import and export require explicit unit selection (millimetres, centimetres or inches) to preserve physical dimensions. [S22]
- **Seam allowances and cutting lines:** Export options control whether seam allowances are included and whether the exporter swaps cutting and sewing lines, a critical choice because cutting on a sewing line trims away the seam allowance. [S22]

- **Grading, notches and annotations:** Exchange settings govern whether grading rules, notches, grainlines, internal lines and piece names survive translation. [S22]

Because documented support for AAMA or ASTM DXF in one manual does not guarantee a lossless round trip into a receiving CAD release, GFT requires every pattern specimen to carry an explicit handoff manifest into downstream chapters. [S19] [S22]

4.21 THE PATTERN HANDOFF DOSSIER CONNECTS GEOMETRY TO SEAMS, MATERIALS, FIT AND CAD READBACK

Before a pattern candidate leaves Chapter 4, the technician must sign off a pattern-geometry record with seven mandatory elements: 1. **Verified panel inventory:** Complete panel count with `panel_id`, role, `cut_qty`, mirror rules and warp `grainline_vector` matching the Chapter 3 brief ({{xr:wp10-ch04-pattern-record-schema}}). [S02] [S22] 2. **Canonical unit and coordinate basis:** Confirmed unit (millimetres), bounding-box dimensions and net enclosed area, verified by the geometry worker. [S22] [S30] 3. **Intermediate snapshot and diff:** Archived DSL program, JSON file or PatternGSL specification, parameter diff against the approved base block and recorded technician edits. [S09] [S10] [S14] [S15] 4. **Indexed boundary and seam-pair map:** Edge identifiers and proposed stitch partners handed to Chapter 5 ([^S12]) so edge correspondence and whole-garment topology are verified separately from single-panel geometry. [S12] [S14] 5. **Material and interlining links:** Material identifiers for self fabric and fusibles handed to Chapter 6 for physical swatch calibration. [S30] 6. **Base-size points of measure and ease intent:** Key garment dimensions handed to Chapter 7 ([^S27]) for size-aware fit and grading evaluation. [S22] [S27] 7. **Exported DXF specimen and settings:** Saved AAMA or ASTM DXF file with declared unit, seam-allowance, cutting-line swap and notch settings, handed to Chapter 9 for receiving-CAD readback and to `wp20` for digital sampling. [S22]

Enforcing this explicit structure between creative intent and specification ensures generative pattern tools act as inspectable assistants under technician authority rather than black-box generators of unverified cutting files. [S15] [S19] [S20] Nothing in this chapter constitutes investment, legal or tax advice, and no pattern representation is production-ready until its seam assembly, material behaviour, physical fit and receiving-CAD readback pass their local gates. [S12] [S19] [S22]

A pattern is a contract in flat geometry: until every boundary closes and every seam has an owner, the machine has only drawn a picture.

CHAPTER SOURCE NOTES

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

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

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

S22 CLO Virtual Fashion. 2026-07-15.

S14 Authors via arXiv. 2026-06-23.

S08 Authors via arXiv. 2024-12-05.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S0401 Authors via arXiv; NeurIPS 2021 Datasets and Benchmarks. 2021-09-12.

S0403 Authors via arXiv. 2023-05-23.

S16 Authors via arXiv. 2026-08-31.

S20 National Institute of Standards and Technology. 2024-07-26.

S03 Authors via arXiv; ACM Transactions on Graphics. 2023-11-07.

S0402 Authors via arXiv; ACM Transactions on Graphics. 2022-01-31.

S17 Authors via arXiv. 2026-09-16.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S19 National Institute of Standards and Technology. 2023-01-26.

S12 Authors via arXiv; WACV 2026. 2026-02-25.

S27 Authors via arXiv; Computers & Graphics. 2026-06-28.

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

S0404 Authors via arXiv. 2026-01-25.

S18 Authors via arXiv. 2026-10-11.

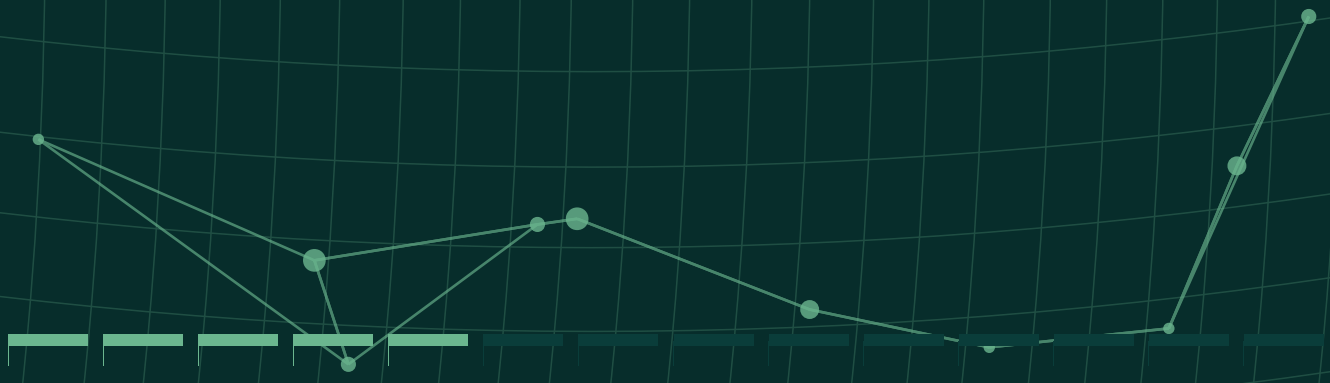
S04 Authors via arXiv. 2024-01-29.

05

PART II: SPECIFICATIONS REQUIRE GEOMETRY, MATERIALS AND FIT / TOPOLOGY AND CONSTRUCTION

A high seam score can coexist with failed garment assembly

Whole-garment assembly must be checked separately from edge-level prediction metrics.



CHAPTER 05
DECISION TEST

Whole-garment assembly must be checked separately from edge-level prediction metrics.

A greenfield apparel maker should not release a generated pattern because a model reports a high seam-prediction score. The score may describe individual relationships, while the production question is whether every required relationship in one garment is correct, physically constructible and approved. AutoSew makes that distinction visible: its test results include strong total relation metrics alongside a lower share of completely resolved digital garments. [S12] GFT therefore recommends a gate with separate outcomes for relation prediction, complete topology, simulated assembly, technician correction and physical seam approval. Each outcome needs its own numerator, denominator, evidence record and authorised decision. A digital pass can narrow the search for errors. It cannot convert a research benchmark into a factory acceptance result.

5.1

Seam correspondence determines how panels become a garment

A flat pattern becomes an assembly instruction only when its boundaries carry relationships. Panel geometry describes where fabric ends. Seam correspondence states which boundary meets another, in which orientation and, where relevant, in which order. AutoSew addresses this problem by treating pattern contours as a graph-matching task, using geometric context to predict stitch relationships without requiring panel names. Its formulation also permits edges to remain unmatched and supports multi-edge connections. [S12] PatternGSL takes a different representation route but makes the same underlying information explicit through panel boundaries, parameterised seams and stitch topology. [S14] Neither representation should be reduced to a drawing of touching lines.

This distinction matters when a technician edits a pattern. Moving a vertex changes geometry. Changing an edge reference changes the assembly. A revision can preserve the visible silhouette while connecting a sleeve to the wrong armhole boundary, reversing one edge or leaving an opening unintended by the designer. Conversely, a large geometric edit can remain correct if the relationship record follows the affected edges. The review unit must therefore include both the contour and the topology, with stable identifiers linking them.

GFT proposes a seam record for the illustrative cotton shirt that contains a seam-pair identifier, the source and destination panel and edge identifiers, sewing direction, gathering or easing intent, notch references and an exception comment. These fields are a release-dossier design, not fields claimed to exist in AutoSew or PatternGSL. Direction should describe the oriented relationship between boundaries. It should not silently be treated as the complete machine sequence, which depends on the construction method and the other operations around the seam.

Stable identifiers need revision discipline. An array position is not a durable identity because splitting a curve or inserting a notch can reorder the data. The corrected record should carry forward the prior seam-pair identifier where intent is unchanged, create a new identifier where construction intent changes and record supersession rather than overwriting the approved version. This allows a reviewer to distinguish a geometric refinement from a new assembly decision.

The record should also distinguish required, optional and intentionally unstitched boundaries. AutoSew's partial-assignment formulation is useful precisely because not every edge must be forced into a pair. [S12] A factory validator needs the same conceptual freedom, but with an important addition: an unmatched edge requires a reason when the product specification expects a closed seam. Free hems, openings and construction-specific exceptions belong in the technician's intent record rather than being inferred from absence alone.

> **GFT NOTE**
>
> Release the seam relationship as an identified engineering object, not as a visual overlay. For every proposed join, require panel and edge identifiers, orientation, intended fullness, notch references and a reviewer exception field. Where an edge remains unmatched, record whether that state is intended, unresolved or rejected. This is GFT's proposed control, informed by the explicit topology in PatternGSL and the partial-assignment problem in AutoSew. [S12] [S14]

5.2

Edge-level F1 is not a whole-garment acceptance rate

AutoSew reports total precision of 96.1 percent, total recall of 96.3 percent and total F1 of 96.2 percent on its multi-edge test setting. Its garment sewn percentage, meaning the share of test patterns for which all ground-truth stitch correspondences are predicted correctly, is 73.3 percent. The abstract rounds the total F1 result to 96 percent. [S12] These are compatible results because the denominators and acceptance rules differ. Relation metrics count correspondences. Garment sewn percentage counts complete patterns and gives no partial credit to a garment with mostly correct relationships.

Precision asks how many predicted relationships agree with the reference among all predicted relationships. Recall asks how many reference relationships are recovered among all reference relationships. F1 is their harmonic mean. None of these measures, by itself, states how errors are distributed across garments. [S12] A small number of relation errors can be concentrated in a few difficult garments or dispersed across many. Those distributions have different consequences for complete digital assembly even when the aggregate relation counts are identical.

The distinction becomes sharper for difficult subsets. AutoSew reports multi-edge precision of 80.4 percent, multi-edge recall of 89.5 percent and multi-edge F1 of 84.7 percent, below its total relation F1 in the same evaluation. [S12] A headline total can therefore coexist with weaker performance on the construction cases that most need technician attention. This is not a universal ranking of models or garment types. It is a reason to read the test protocol, subset definitions and complete-garment outcome together.

{{xr:wp10-ch05-autosew-metric-boundaries}} places the reported measures side by side without putting them on a false common denominator.

EXHIBIT 18 / PERCENT, WITH DIFFERENT DENOMINATORS

AutoSew reports distinct relation and assembly metrics

96.2%

Total relation F1

Precision 96.1%; recall 96.3%; all evaluated relations

84.7%

Multi-edge relation F1

Precision 80.4%; recall 89.5%; multi-edge subset

73.3%

Error-free digital assembly

Fully resolved test patterns divided by all test patterns

SOURCED [S12] Sourced from AutoSew full-text sections 8.1 and 8.2. Total and multi-edge metrics concern stitch correspondences on the authors' test set. Garment sewn percentage is fully resolved test patterns divided by all test patterns. The abstract rounds total F1 to 96 percent. These are digital benchmark results, not physical sewing yield or GFT performance.

The exhibit does not show that 26.7 percent of physical garments would fail. The arithmetic complement, 100 percent minus 73.3 percent, is 26.7 percent of the authors' digital test garments not fully resolved under their stated benchmark rule. [S12] It does not measure cutting defects, seam strength, production rework, operator time or a Cameroonian factory outcome. It also does not establish that a human technician could not correct the digital outputs. Those are separate questions requiring separate observations.

WORKED EXAMPLE: ONE RELATION SCORE CAN HIDE DIFFERENT ASSEMBLY OUTCOMES

Consider an illustrative set of 100 garments, each with 10 required reference relationships, giving 1,000 reference relationships. Suppose the system emits 1,000 predictions: 960 true positives, 40 false positives and, consequently, 40 false negatives against the reference set. Precision is 960 / (960 + 40) = 96 percent. Recall is 960 / (960 + 40) = 96 percent. F1 is therefore 96 percent. These counts are constructed solely to demonstrate the arithmetic and are not AutoSew observations. [S12]

Now distribute the 80 relation errors across the illustrative garments. If all errors affect 20 garments, the other 80 garments are error-free. If at least one error affects each of 73 garments, only 27 garments are error-free. Relation precision, recall and F1 remain 96 percent in both cases. The complete-assembly outcome changes because its unit is the garment and its pass rule is all required relationships correct. No independence assumption is needed, and none is justified by an F1-score.

{{xr:wp10-ch05-f1-assembly-scenarios}} preserves that worked example in a compact comparison. It also shows why raising a relation score to a power is not a valid way to infer garment success: F1 is not an independent per-garment survival probability, and errors need not be independent within or between garments. [S12]

EXHIBIT 19 / ILLUSTRATIVE COUNTS

Identical relation F1 can produce different assembly outcomes

Error allocation	TP / FP / FN	Relation F1	Garments with errors	Error-free garments
Concentrated	960 / 40 / 40	96%	20	80 of 100
Dispersed	960 / 40 / 40	96%	73	27 of 100

ILLUSTRATIVE [S12] Illustrative GFT arithmetic, not observed AutoSew error allocation. Both rows use 1,000 predictions with 960 true positives, 40 false positives and 40 false negatives, giving 96 percent precision, recall and F1. The only change is how the 80 errors are distributed across 100 garments. S12 motivates the separate whole-garment metric.

For a local pilot, the useful report is consequently a small metric family rather than one headline. Keep relation precision and recall for diagnosis, complete digital assembly for system-level screening, technician-corrected assembly for the assisted workflow and physical seam approval for release. Record all four with their denominators. NIST's framework calls for accuracy measures to use clearly defined, realistic test sets and documented methods, while leaving risk tolerance to the organisation. [S19] GFT's proposed acceptance thresholds should likewise be agreed for the intended shirt, fabric and buyer before results are reviewed, not imported from the AutoSew benchmark.

5.3

Many-to-one relationships require construction intent

Real construction does not always present one edge on one panel and one identical edge on another. AutoSew's authors created their multi-edge dataset after identifying cases where a sleeve attaches to both front and back torso panels. Their representation also merges split sleeves, cuffs and torso pieces to reduce artificial one-to-one assumptions in the source data. [S12] The technical problem is therefore not just finding two similar curves. It is deciding how a boundary segmented in one representation corresponds to several boundaries in another while preserving the intended assembly.

Segmentation can change for innocent reasons. A curve may be divided at a notch, split by a design edit or represented as one continuous edge on the opposing panel. A validator that expects equal edge counts would reject a workable correspondence. A matcher that forces every edge into exactly one pair could invent crossings or leave part of a required seam uncovered. The release record should instead describe coverage: which portions of each boundary participate, their directions and the references that align them.

Length equality is also an incomplete test. The illustrative shirt can include intentional fullness where a sleeve meets a cuff, or controlled easing around an armhole. The technician should record that intent, the relevant notches and the accepted construction method. A raw length difference then becomes a prompt to inspect the easing

specification, not automatic proof of an error. Equally, calling every mismatch ease would conceal a genuine drafting fault. The exception is valid only when a responsible technician records it against identified edges.

A useful many-to-one review asks four questions. Does every required portion of the longer boundary map to a destination? Do the mapped portions avoid unintended overlap or gaps? Do orientations and notches place fullness where the design intends? Does the physical construction instruction explain how the operator should manage that fullness? The first three questions concern digital topology. The fourth belongs to the sample room. Keeping them separate prevents a plausible simulation from filling in a missing instruction.

AutoSew's lower multi-edge F1 than total F1 reinforces the need to inspect this subgroup separately in its benchmark. [S12] It does not follow that every sleeve, cuff or gathered seam in another dataset will behave similarly. The local test should select actual pattern variants, including held-out segmentation changes, and report failures by construction family. A greenfield maker can then learn whether its own representations preserve many-to-one intent before it considers automation beyond a narrow block library.

> **PROPOSED MANY-TO-ONE SEAM CHECK**
>
> For each identified many-to-one relationship, freeze the pattern revision and display the participating edges in contrasting annotation rather than altering their geometry. Record source coverage, destination coverage, direction, notch correspondence and intended fullness. Run deterministic checks for missing coverage, duplicate coverage and reversed segments. The technician then accepts an intentional ease exception or returns the relationship for correction. Finally, the sample technician records whether the physical construction method realises the approved intent. No universal length-difference tolerance is assumed.

This protocol creates evidence at the point where a metric loses detail. If the digital check fails, the reviewer receives edge identifiers and a coverage map. If the physical check fails after a digital pass, the reviewer knows that topology alone was insufficient and investigates the construction specification or material interaction. Both outcomes are useful because they locate the next decision instead of merely labelling the garment defective.

5.4

Geometric validity does not settle the sewing method

A correct relationship says that two boundaries should meet. It does not specify the seam type, stitch formation, thread, reinforcement, allowance, finishing operation or machine arrangement needed to join them. PatternGSL's explicit seams and stitch topology make structure available for simulation, but its abstract does not establish those factory instructions or production tolerances. [S14] AutoSew's geometric input and digital assembly metric answer a similarly bounded question. [S12] A production release needs an additional construction layer.

GFT proposes that this layer begin with the approved material and trim identifiers, then name the seam and stitch requirement, seam allowance, edge finish, reinforcement and machine or attachment constraint. The record should identify where a construction choice is mandatory and where the sample technician may select an equivalent method. Thread and needle details should be completed from actual product and equipment evidence rather than generated as plausible text. If a choice remains open, the field should read awaiting technician decision, not inherit a generic shirt recipe.

The distinction between stitching line and cutting line is central. Seam correspondence belongs to the intended join. Seam allowance determines the material retained beyond that join for assembly and finishing. An automated topology check can pass while an allowance is absent, inconsistent with the selected construction or incorrectly transferred during export. The reviewer must therefore inspect allowance geometry and annotations alongside the relationship, without mistaking a closed outer contour for a complete sewing instruction.

Construction order also needs an owner. Attaching a collar, closing a cuff or finishing a placket can change access for a later operation. The digital relationship graph may remain correct while the proposed sequence is impractical. For the illustrative shirt, the sample technician should review closure order, visible finishing and access to layered areas before approving the sample plan. This is a proposed GFT review, not an output guarantee from either cited garment model.

{{xr:wp10-ch05-construction-specification}} separates what the digital relationship contributes from what the release dossier must add.

EXHIBIT 20 / PROPOSED RELEASE FIELDS

Digital correspondence omits physical construction choices

Specification element	What topology can carry	What the technician must add or confirm
Seam identity	Source and destination edges	Stable seam-pair ID and intended operation
Direction and alignment	Oriented relationship and endpoints	Notches, starting reference and easing distribution
Allowance	Associated boundary geometry where encoded	Net-to-cut allowance and suitability for the selected seam
Seam and stitch	No complete machine method	Seam type, stitch requirement and acceptable alternative
Material interaction	Panel identity where linked	Approved fabric, trim, thread and interlining records
Reinforcement	No stress or durability evidence	Reinforcement location and agreed physical test
Machine arrangement	No equipment instruction	Machine, folder or attachment constraint and sample-room review
Assembly sequence	Relationships between panels	Operation order, access and closure review
Acceptance	Digital assembly result	Physical sample evidence and authorised approval

ILLUSTRATIVE [S12] [S14] Illustrative GFT teaching checklist for the running cotton-shirt example. S12 supports the bounded digital stitch-correspondence task and S14 supports explicit topology. The physical construction fields and review questions are GFT proposals, not outputs or guarantees reported by those papers.

The table is deliberately a handoff checklist rather than a claim about a universal seam library. The appropriate construction depends on the chosen product, material, equipment and buyer requirement. A changed fabric or machine can reopen a previously acceptable choice even if topology stays unchanged. Material review belongs to Chapter 6, and receiving-CAD verification belongs to Chapter 9, but the seam gate must preserve the identifiers those later checks need.

Physical approval should also state what was actually inspected. A sample can demonstrate the selected construction under recorded conditions. It does not prove every size, fabric batch or production operator will reproduce it. Where seam performance matters, the responsible team should define the test method and acceptance condition before sampling. Without that record, a neat photograph is only appearance evidence, not evidence of strength, durability or production capability.

5.5

Assembly review needs an explainable failure taxonomy

A single pass or fail field is too coarse to manage correction. Parsing can fail before relationships are predicted. Prediction can produce an incorrect pair. Deterministic validation can detect an unintended opening. Simulation can expose an assembly problem. A technician can correct the topology without changing the model. A physical sample can fail after every digital stage passes. These are different events, and combining them into one defect rate would hide the intervention that resolved the problem.

GFT proposes six starting failure categories: missing seam, reversed orientation, incorrect pairing, unexpected opening, duplicated panel and unsupported construction. Layer collision is added when multi-layer review applies. The taxonomy is a teaching and pilot design, not a measured distribution. Each code should attach to a pattern revision, model and configuration record, panel and edge identifiers, an annotated view, the stage that detected it and the responsible correction role. A free-text explanation can add context, but it should not replace the structured code.

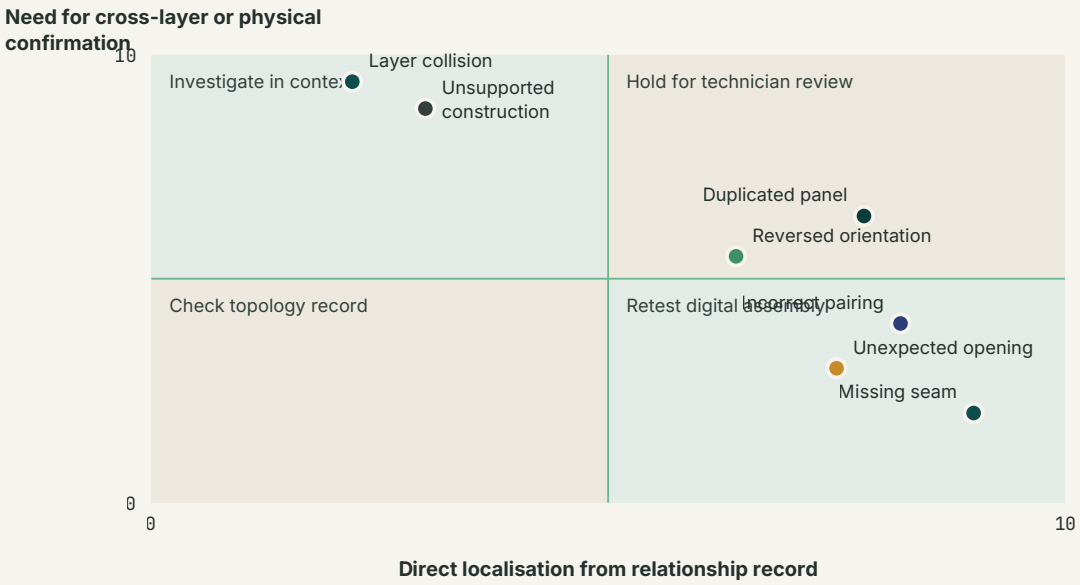
The correction record needs an outcome as well as a cause. Useful states include returned to geometry, returned to construction intent, accepted as an equivalent valid topology, corrected and awaiting retest, physically rejected and released after approval. The equivalent-topology exception matters because AutoSew's authors report that identical or interchangeable panel shapes can cause valid solutions to be treated as failures against a single ground-truth arrangement. [S12] A technician must be able to record why an alternative is construction-equivalent rather than allowing the benchmark label to dictate the factory decision.

Retest scope should follow the failure. A reversed edge requires a topology and orientation retest. A changed allowance requires geometric readback and construction review. A duplicated panel may require component-count, material-use and assembly checks. An unsupported construction requires sample-room evidence even if the corrected file simulates. The record should name the retest and prevent an unrelated successful check from clearing the original failure automatically.

{{xr:wp10-ch05-failure-taxonomy}} maps the proposed categories by the kind of review they need, not by invented frequency or severity scores.

EXHIBIT 21 / ILLUSTRATIVE REVIEW PLACEMENT (0 TO 10)

Assembly errors need categories before rates



ILLUSTRATIVE [S12] [S20] Illustrative GFT analytical placement, not observed defect frequency, probability or severity. Horizontal position expresses how directly the proposed failure can be localised from a two-dimensional relationship record. Vertical position expresses the proposed need for cross-layer or physical confirmation. S12 motivates topology failure review and S20 motivates contextual empirical evaluation. Positions are workshop judgements to be tested locally.

The positions are not measurements and should not be averaged into a risk score. Their purpose is to prompt the right evidence attachment. A missing seam should be visible in the relationship record, while unsupported construction may remain invisible until a technician examines the method or a sample is made. Local review can replace the illustrative placements with observed detection stages and correction effort after enough cases have been recorded.

Severity should remain separate from detection ease. An error that is obvious in a topology view can still have serious construction consequences, while a subtle visual anomaly may have no effect on the approved assembly. The responsible team should define criticality by the affected product function and release decision, not by the matrix coordinates above. Any criticality rule should be agreed before the pilot and applied consistently to manual and assisted work. NIST leaves such tolerance decisions to the organisation rather than prescribing a universal threshold. [S19]

Aggregation needs equal care. Several failed relationships inside one garment can explain one complete-assembly failure, but they must not be counted as several failed garments. Conversely, one garment-level pass must not erase the number of corrections needed to reach it. The pilot register should therefore link relation events to a garment identifier and preserve both views. Missing simulation results, abandoned runs and cases withdrawn for unrelated material reasons should remain visible with an exclusion reason rather than disappearing from the denominator.

GFT NOTE

Report detection stage and correction stage separately. A model may propose the error, a deterministic validator may detect it and a technician may resolve it. Counting only the final released garment would erase the assistance burden. Counting every initial model error as a production defect would overstate factory failure. Preserve both the initial event and the final outcome so the pilot measures the workflow rather than promoting or condemning the model by definition. [S19] [S20]

5.6

Layering increases the need for inspection

Layered garments add relationships that a single visible shell can conceal. Interfacing, linings, cuffs, collars and overlapping panels introduce questions about which layer sits where, which edges remain accessible and when a closure should occur. A clean rendered exterior cannot answer all of them. ContourCraft addresses collisions and intersections in learned multi-garment simulations, describing that problem as largely unsolved and proposing an intersection-contour loss with collision-avoiding repulsion. [S06] That research is relevant to simulation quality. It is not evidence that a rendered layered garment can be sewn to specification.

DressAgent studies a different but related problem: reconstructing layered outfits through layer planning, pattern estimation and closed-loop refinement grounded in physical simulation. Its agents edit, simulate and judge patterns against a reference using visual and measured evidence. [S18] This makes layer order and iterative correction explicit research concerns. It does not make an agent's final render equivalent to a technician-approved construction sequence. The abstract reports a research method and benchmark claims, not a factory release protocol or a locally reproduced result.

The seam gate should therefore request a layer-stack view in addition to the flat topology. For each relevant region, the technician should identify outer fabric, lining or interfacing, the order in which layers are assembled and the operation that closes the region. Hidden edges need the same stable identifiers as visible ones. If the simulator resolves an initial intersection, the record should preserve that event and the correction rather than presenting only the cleaned final image. A repaired visual state is not proof that the underlying pattern or construction intent was repaired.

The illustrative cotton shirt offers a bounded test without needing a complex outerwear example. Collar stand, collar fall and interfacing create a layered assembly. Cuffs and plackets create closure and access questions. The technician should review these regions with annotated cross-sections or an equivalent physical construction reference, then confirm the sequence on a sample. This is a GFT proposed inspection method. No claim is made that ContourCraft or DressAgent supplies the required cross-section, seam allowance or factory operation order.

Simulation evidence remains useful within this boundary. It can draw attention to intersections, implausible layer order or geometry that cannot initialise cleanly. The reviewer should log the simulator and settings, avatar or body representation, material revision, initial failures and corrections. Those records allow a later reviewer to distin-

guish a topology correction from a solver workaround. Material parameters still require the physical evidence process in Chapter 6, and fit requires the separate size-aware checks in Chapter 7.

A layered test set should include deliberate hidden-layer errors as well as ordinary approved variants. Otherwise, a pilot may demonstrate only that the system handles cases already close to its training or tuning examples. NIST guidance supports contextual testing and expert involvement rather than assuming that a laboratory result transfers unchanged to operational conditions. [S19] [S20] The maker's acceptance question is not whether a render looks resolved. It is whether the specified layers can be assembled in the recorded order with the actual materials and equipment.

5.7

A seam gate should return evidence with the rejection

The final gate should be a decision record, not a score threshold copied from a paper. GFT proposes releasing a seam specification only when every required relationship is identified, orientation and coverage checks pass, intentional fullness has an approved exception, no unexplained opening or duplicate component remains, and the construction record is complete. Simulated assembly must have no unresolved issue relevant to release. Where physical sampling is required, the responsible technician must record the sample outcome and any conditions before bulk authorisation.

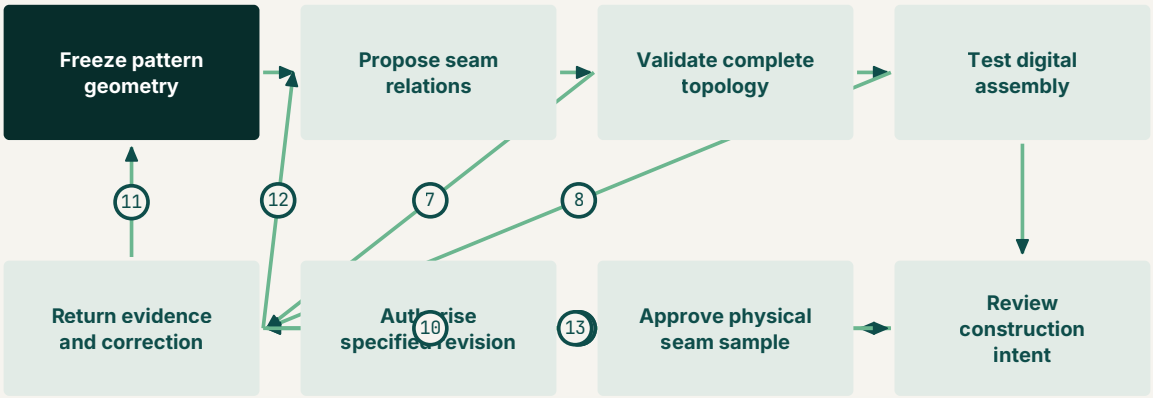
A rejection should contain enough evidence for another person to act without rerunning the entire investigation. The minimum packet is an annotated panel view, failed relationship or region identifiers, the detecting stage, the failure code, the current pattern revision and the required correction owner. If the reviewer already knows the appropriate correction, the packet can propose it, but the corrected topology must remain distinguishable from the rejected version. The last approved pattern revision remains available throughout.

Routing should follow the cause. Ambiguous panel boundaries return to the geometry process in Chapter 4. Fabric-dependent seam behaviour returns to the material review in Chapter 6. Missing or distorted annotations after export go to the receiving-CAD readback in Chapter 9. Unsupported construction returns to the sample technician. This prevents every problem becoming a prompt revision and prevents a model update being used as a substitute for engineering correction.

{{xr:wp10-ch05-seam-evidence-loop}} shows the proposed loop, including the return path that a simple linear generation pipeline omits.

EXHIBIT 22 / PROPOSED DECISION SEQUENCE

A failed seam must return to an identified correction point



6. Pass with approval
7. Fail
8. Fail
9. Hold
10. Reject
11. Correct geometry
12. Correct topology
13. Correct method

ILLUSTRATIVE [S12] [S19] [S20] Illustrative GFT workflow informed by the digital stitch task in S12 and the contextual evaluation and monitoring guidance in S19 and S20. The physical sample and release roles are GFT proposals. The flow is not a reported implementation or a measured factory process.

The flow separates correction of geometry, topology and method because they do not always require the same retest. It also makes physical approval a distinct state rather than an automatic consequence of simulation. What the diagram cannot show is workload. A pilot should record technician correction time, repeated failures and waiting time so that a high final approval share does not conceal an unacceptable review burden. Economic analysis belongs to Chapter 12; the seam gate supplies the events and denominators.

Four proposed measures can make the gate inspectable. First-pass complete digital assembly divides garments passing the defined digital gate without technician topology correction by all evaluated garments. Corrected digital assembly divides garments passing after recorded corrections by all evaluated garments. Physical seam approval divides samples meeting the agreed construction conditions by all evaluated samples. Escaped seam failures count relevant failures detected after the declared release point, with the release population and observation window stated. None has a local baseline yet, so no target value is asserted here.

Relation precision and recall can remain diagnostic measures alongside these outcomes, particularly for multi-edge subsets. They should not replace the complete-garment measures. Nor should corrected digital assembly be compared directly with first-pass assembly without disclosing technician effort. The measurement design follows NIST's emphasis on defined test contexts, realistic test sets and human judgement about thresholds, while the organisation retains responsibility for its own risk tolerance. [S19] [S20]

For GFT's Cameroon textile programme, the practical recommendation is to qualify this gate on approved local blocks before expanding generative specification work. Begin with a narrow shirt family, preserve manual technician authority and test ordinary variants alongside many-to-one, layered and deliberately incorrect cases. Record the evidence needed for correction and stop release whenever a critical relationship or construction decision remains unresolved. AutoSew, PatternGSL, ContourCraft and DressAgent provide research mechanisms and benchmark questions, not local acceptance evidence. [S06] [S12] [S14] [S18]

The resulting decision is deliberately conditional. A maker can adopt seam-prediction assistance as a diagnostic aid when its errors are visible, correctable and evaluated at the whole-garment boundary. It should not adopt the headline relation score as a proxy for manufacturing readiness. The evidence required to move from digital assistance to production release remains a complete specification, a technically reviewed construction method and the prescribed physical approval.

A seam is accepted only when the whole garment, not merely the average relationship, survives review.

CHAPTER SOURCE NOTES

S12 Authors via arXiv; WACV 2026. 2026-02-25.

S14 Authors via arXiv. 2026-06-23.

S19 National Institute of Standards and Technology. 2023-01-26.

S20 National Institute of Standards and Technology. 2024-07-26.

S06 Authors via arXiv; SIGGRAPH 2024. 2024-05-15.

S18 Authors via arXiv. 2026-10-11.

GREEN FUTURE TECH

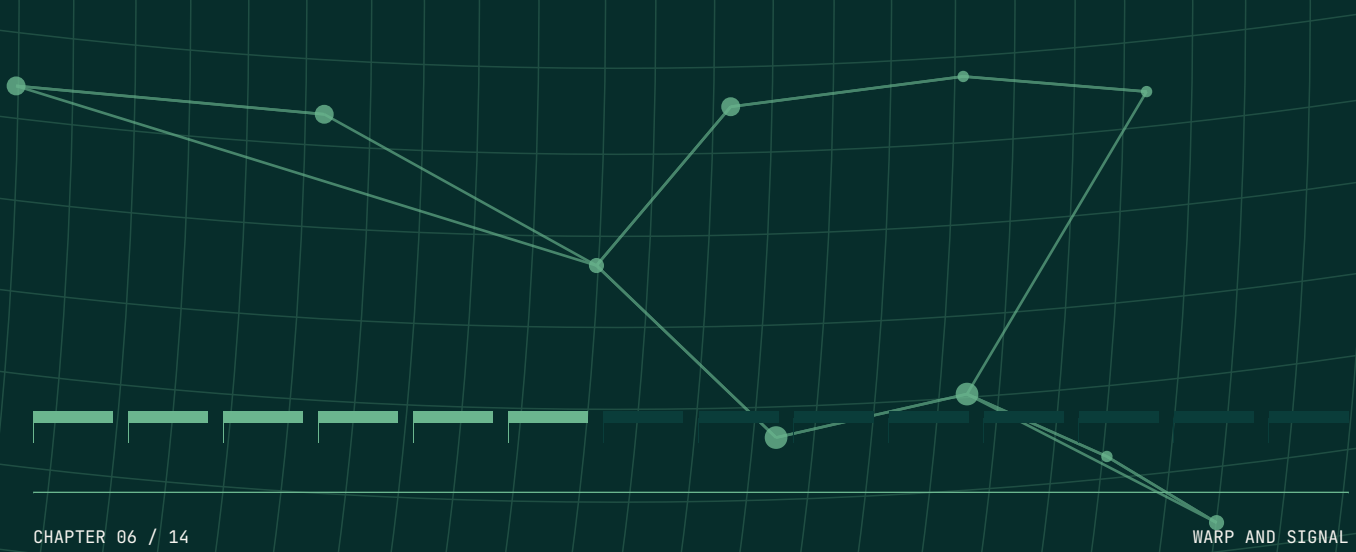
LIMITED / HONG KONG SAR

06

PART II: SPECIFICATIONS REQUIRE GEOMETRY, MATERIALS AND FIT / MATERIAL REALITY

Fabric physics must be checked against the fabric being bought

Predicted material properties can guide a test, but a simulation must remain traceable to the actual material and its measurement basis.



CHAPTER 06
DECISION TEST

Predicted material properties can guide a test, but a simulation must remain traceable to the actual material and its measurement basis.

A generative model can synthesise a convincing cloth surface or infer plausible mechanical parameters from a single photograph, yet neither output certifies how a real roll of cotton behaves under a cutting knife, sewing needle or steam press. [S11] Green Future Tech Limited (GFT), Hong Kong SAR, recommends a strict rule for a greenfield apparel maker: predicted material properties may guide early digital exploration, but a garment specification must not advance to pattern adjustment, fit sign-off or bulk cutting until its simulation inputs are linked to a measured physical swatch from the actual fabric lot being bought. [S32] In Cameroon, where the inherited sector baseline records approximately 95 percent of cotton fibre exported raw against roughly 5 percent processed locally, and where national policy targets 50 percent local processing by 2030, capturing value locally depends on converting domestic lint into specification-compliant garments rather than digital renderings. [R06] [R07] Confusing a texture map with verified cloth mechanics hides errors in mass, stiffness, elongation and shrinkage until assembly. [S28]

6.1

A texture map does not describe mechanical behaviour

Visual appearance and mechanical deformation solve different problems and consume different inputs. [S28] A surface texture pipeline estimates how light reflects off a yarn structure across two-dimensional ultraviolet (UV) coordinates, whereas a cloth solver integrates internal elastic, shear and bending forces against gravity and contact across a panel mesh. [S11] A workflow breaks down whenever software or staff treat a plausible poplin render as evidence that digital cloth will drape, stretch or recover like physical fabric from a mill. [S01]

6.2 RENDERING CHANNELS AND CONSTITUTIVE PARAMETERS OCCUPY SEPARATE SPACES

Recent research separates surface-appearance synthesis from constitutive mechanical estimation, and factory governance must preserve that boundary. [S28] On the rendering side, OmniFabric synthesises textures for reconstructed three-dimensional garments in coherent sewing-pattern UV space while removing lighting contamination. [S28] That helps a team inspect motif continuity, sheen and weave appearance without shadows distorting the preview. [S28] However, an albedo, normal or roughness map carries zero values for areal density, yarn crimp, tensile modulus or friction. [S32] Even in visual classification, multimodal models struggle to identify textile structures reliably: a comparative study of large language models in apparel design reports that Deepseek faced significant challenges in generating and recognising fabric images, with evaluated models achieving approximately 80 percent accuracy on base fabric construction and roughly 55 percent on detailed construction within the authors' test set. [S01] An image model that misclassifies detailed weaves nearly half the time in a benchmark cannot certify materials in a factory. [S01]

On the mechanical side, Image2Garment augments geometry prediction with a material estimation branch, using a vision-language model to infer fabric attributes and a predictor grounded in a small measurement dataset to output simulation-ready physical parameters from a single image. [S11] Depth-capture research likewise shows that estimating mechanical parameters from casual depth images requires sim-to-real

transfer protocols because naive parameter regression fails to match human perception of folds. [S0601] Furthermore, the Fabric-101 study across over 101 common fabrics shows that realistic cloth behaviour involves three distinct deformation components: self-recoverable elastic deformation, friction-driven recoverable deformation and unrecoverable plastic deformation measured through cyclic loading and unloading. [S0602] A flat RGB texture map records none of those three regimes. [S0602]

6.3 FOUR DISTINCT MATERIAL IDENTIFIERS PREVENT SILENT SUBSTITUTION

Because appearance assets, inferred presets and lab measurements evolve on different schedules, GFT proposes assigning four versioned identifiers within the specification spine rather than one generic fabric label. [S32] In this illustrative GFT architecture (**Exhibit 23**, p. 94), every material reference resolves into four linked records:

1. **Supplier material lot identifier (**`supplier_lot_id`**)**: records the mill, yarn count, construction, finishing route, dye lot, roll width and delivery batch for the physical cloth under review. [S32]
2. **Physical swatch identifier (**`swatch_id`**)**: identifies the retained physical specimen cut from a roll or sample yardage, stored in the sample room with its custody date and conditioning log. [S20]
3. **Measured property record identifier (**`measured_property_id`**)**: stores quantitative mechanical and dimensional observations from the swatch, alongside test method, instrument, orientation and uncertainty, or flags the parameters as an unverified model estimate. [S11] [S32]
4. **Visual asset identifier (**`visual_asset_id`**)**: points to the colour-calibrated diffuse, normal and roughness texture maps used for rendering in three-dimensional design tools. [S28] [S32]

EXHIBIT 23 / ARCHITECTURE LAYERS

Physical and visual material records need a common identity



ILLUSTRATIVE [S11] [S28] [S32] Illustrative GFT data-architecture proposal separating physical lot custody, laboratory measurement, simulation physics and rendering maps across linked revision identifiers. Informed by S11, S28 and S32; not an installed factory system.

Separating these four layers in **Exhibit 23**, p. 94 prevents a common failure mode: a designer swaps a poplin texture map for a twill texture map in the viewport without updating bending stiffness, or a buyer substitutes a supplier lot that looks identical on a table but exhibits different weft elongation and shrinkage. [S28] [S32] While **Exhibit 23**, p. 94 defines an inspectable data relationship for a specification system, it is an illustrative GFT governance proposal and cannot prove that vendor software enforces foreign-key integrity across those records. [S20]

6.4

Material inference is a hypothesis that needs validation

When a generative workflow infers physical parameters from an image or text description, the output is a prior hypothesis about what the garment in the photograph might be made of, not a measurement of cloth in the warehouse. [S11] Treating an inferred preset as an initial search hypothesis allows a team to screen early silhouettes while requiring measured material evidence before physical release. [S19]

6.5 SINGLE-IMAGE PHYSICAL PREDICTION RELIES ON BOUNDED TRAINING PRIORS

Image2Garment introduces two datasets, FTAG and T2P, to connect visual garment attributes to simulation parameters without iterative optimisation at inference time. [S11] A fine-tuned vision-language model predicts semantic fabric attributes from an input photograph, and a neural network trained on a small measurement dataset maps

those attributes to simulator parameters so predicted sewing panels can be draped directly. [S11] A related route, Garment3DGen, deforms template meshes and synthesises textures so three-dimensional stylized garments can be draped and simulated on bodies. [S05]

Both approaches advance computer-graphics asset creation, yet both operate under ambiguities that a factory cannot ignore. [S05] [S11] First, monocular photographs are under-constrained with respect to fabric mass and construction: a crisp fold can result from a heavy cotton twill, a light poplin with a resin finish, or a synthetic blend captured mid-motion. [S11] [S0601] Second, an attribute-to-physics predictor only interpolates within its measurement dataset; if a Cameroonian mill produces a carded cotton weave with a yarn twist or finish outside FTAG and T2P, inferred parameters reflect the training prior rather than local cloth. [S11] [S30] Third, neither tool observes chemical finishing, pilling, seam slippage or wet-processing relaxation. [S05] [S11]

6.6 THE RUNNING COTTON SHIRT SEPARATES PRESET EXPLORATION FROM SWATCH EVIDENCE

Consider the illustrative short-sleeve cotton shirt carried through this paper as a teaching case. [S09] Suppose a designer uses a multimodal generator to explore collar and pocket variations on an approved base block, and the system attaches an inferred "medium cotton shirting" preset derived from a reference image. [S09] [S11] At concept screening, that preset lets the technician run a preliminary simulation to check whether pocket placement pulls the front panel or whether the collar stand collapses under a plausible weight. [S11]

Before that shirt moves to sample cutting or grading, the hypothesis must be tested against the fabric sourced by procurement. [S32] In GFT's proposed workflow, the technician assigns a provisional status (`status: estimated_preset`) to the simulation and requests the physical swatch (`swatch_id`) from the candidate roll. [S20] [S32] Once the swatch is weighed, measured for thickness and checked for directional bending and stretch, the technician reruns the shirt pattern on the same avatar using the measured record (`measured_property_id`). [S27] [S32] If collar roll, shoulder line or hem flare shifts between the inferred preset and the measured swatch simulation, or between the simulation and a sewn toile, the difference is logged as a calibration deviation. [S19] No calibration trial outcomes are asserted here; the protocol is an illustrative design for local execution. [S19]

6.7

Material assets need measurement provenance

A number entered into a simulation material dialog has no engineering value unless a reviewer can inspect how, when, in which direction and on what specimen that number was obtained. [S20] [S32] In many three-dimensional workflows, digital fabrics are exchanged as opaque files whose numeric sliders were adjusted visually until screen folds looked pleasing. [S28] For manufacturing release, GFT recommends treating every simulation material preset as a metrological record requiring explicit provenance metadata. [S20]

6.8 VENDOR ASSET FIELDS START THE RECORD BUT DO NOT COMPLETE THE AUDIT TRAIL

Commercial platform documentation recognises the distinction between visual maps and physical descriptors, even if its mandatory fields are minimal. [S32] In CLO Virtual Fashion's CONNECT documentation for creating a three-dimensional fabric asset, the author is instructed that weight and thickness of the fabric must be entered in the asset description, alongside tags for fabric type, fibre content and visual texture maps, while measurement information is requested separately. [S32] That vendor requirement establishes a baseline: a digital fabric asset without stated areal weight and thickness is incomplete even by commercial asset-store standards. [S32]

Yet asset-upload instructions for a digital marketplace are narrower than an industrial quality record. [S20] [S32] Entering a nominal mass per unit area from a supplier's flyer does not reveal whether the delivered roll is five percent heavier, whether thickness was measured under a standardised presser-foot pressure, or whether bending and stretch parameters were measured on a test rig, inferred from a depth camera or copied from a library template. [S32] [S0601] Following the source-verification logic in the National Institute of Standards and Technology (NIST) Generative Artificial Intelligence Profile, a maker should distinguish between vendor defaults, model-inferred estimates and empirically verified measurements. [S20]

6.9 CONDITIONING, INSTRUMENT CONTEXT AND UNCERTAINTY GOVERN COMPARABILITY

Cotton is a hygroscopic natural fibre whose stiffness, mass and dimensions vary with ambient relative humidity and temperature. [R06] [S0602] A cotton shirting swatch weighed and tested in an unconditioned room during the dry season in northern Cameroon will not yield the same mass or bending response as the same cloth tested in a humid coastal facility without conditioning. [R03] [S20] To make material records auditable across time and partners, GFT proposes expanding the baseline asset record documented in S32 into the disclosure schema shown in **Exhibit 24**, p. 97. [S32]

EXHIBIT 24 / SPECIFICATION FIELDS

Every simulation input should disclose its basis

Property field	Unit	Basis flag	Method and instrument	Provenance and uncertainty
Areal mass density	g/m2	Measured required	Cut specimen scale; state standard	Swatch ID, conditioning, roll dispersion
Fabric thickness	mm	Measured required	Thickness gauge; state foot pressure	Swatch ID, mean and range across width
Warp and weft tensile	N/m or %	Estimated or measured	Inferred (S11) or tensile rig (S0602)	Disclose training prior or cyclic curve
In-plane shear stiffness	N/m or deg	Estimated or measured	Bias extension or simulator device	Record bias locking angle and hysteresis
Bending rigidity	uN.m	Estimated or measured	Cantilever/drape test or depth capture	Separate warp and weft face/back values
Laundering shrinkage	% warp/weft	Measured required	Wash-dry template test on lot swatch	Mandatory before marker or grading release

ILLUSTRATIVE [S11] [S20] [S32] [S0602] Illustrative GFT specification table for the teaching cotton shirt. Shows how estimated presets and measured swatch records are distinguished without inventing unsourced fabric test values. Informed by S11, S20, S32 and S0602.

As **Exhibit 24**, p. 97 sets out, every property row carries five fields: parameter name, engineering unit, basis flag (`estimated_preset`, `supplier_nominal` or `measured_swatch`), method or inference model identifier, and provenance and uncertainty (direction, conditioning, date, provider and range). [S11] [S20] [S32] Where a greenfield maker lacks a full suite of mechanical testing instruments, **Exhibit 24**, p. 97 lets the team record areal mass, thickness, cantilever overhang and wash shrinkage locally while marking complex hysteresis parameters as model estimates (`estimated_preset`) awaiting external verification. [S32] [S0602] The table is an illustrative GFT specification template and omits invented numeric test thresholds; actual tolerances must come from buyer specifications and measured lots. [S19]

6.10

Direction and fabric variation belong in the model assumptions

Woven and knitted textiles are anisotropic materials whose mechanical response depends on orientation relative to yarn structure, and whose physical properties vary across roll width and between finishing batches. [S02] [S0602] A workflow that treats cloth as an isotropic sheet or relies on a generic "cotton" preset mispredicts how panels hang on the bias, how collars resist curling, and how finished garments measure after washing. [S11] [S27]

6.11 ANISOTROPY ACROSS WARP, WEFT AND BIAS ALTERS PANEL DEFORMATION

In woven cotton shirting, warp yarns held under loom tension exhibit different crimp and elongation from weft yarns inserted across the width, while the 45-degree bias direction is governed by trellis shear until yarns jam. [S02] [S0602] Knitted wale and course directions show even wider differences in extensibility and recovery. [S0602] Chapter 4 established that every panel in a structured representation must carry an explicit grainline vector linking the two-dimensional polygon to the fabric axis. [S02] [S09] That grainline has physical consequences only if the solver distinguishes warp stiffness, weft stiffness and shear resistance in its material model. [S11]

When Image2Garment infers parameters from a single photograph, or when a depth camera estimates mechanics from a draped specimen, grain orientation relative to folds determines which directional stiffness components are excited. [S11] [S0601] Furthermore, as Fabric-101 demonstrates across over 101 fabrics, real textiles exhibit non-linear hysteresis under cyclic loading and unloading, leaving friction-driven or plastic wrinkles. [S0602] For the illustrative cotton shirt, a back yoke aligned to the weft or a placket cut on the bias deforms differently from front panels aligned with the warp. [S02] [S09] If a simulator uses an isotropic approximation or swaps warp and weft during import, digital review misses armhole distortion or yoke pulling. [S22] [S27]

6.12 SHRINKAGE AND BATCH DISPERSION REQUIRE PHYSICAL TEST TOLERANCES

Elastic drape simulation models a garment in its dry, pre-laundered state unless post-wash dimensional change is built into pattern allowances. [S09] [S32] In cotton processing, tensions from spinning, weaving, bleaching, dyeing and drying relax when the garment is steamed or washed. [R06] [S32] Two lots delivered under the same style name can exhibit different warp and weft shrinkage if finishing overfeed or sanforisation settings drift. [S20]

GFT therefore requires verifying dimensional stability and batch dispersion through physical testing on each lot before cutting markers are authorised. [S19] No mechanical test numbers or standard-method thresholds are invented here; a factory must obtain supplier certificates and run bench wash-and-dry shrinkage trials under the buyer's standard. [S19] [S32] The worked example below steps through an arithmetic check on the running cotton shirt. [S09]

WORKED EXAMPLE: ILLUSTRATIVE SHRINKAGE AND DIRECTIONAL ALLOWANCE ON THE TEACHING COTTON SHIRT

Consider an illustrative teaching calculation for one base size of the running cotton shirt. All numerical inputs in this box are hypothetical scenario assumptions constructed to demonstrate the reconciliation method; they are not measurements from a GFT facility or a Cameroonian mill. [S09]

Assume the post-wash specification for the illustrative shirt requires a finished half-chest width of 540 millimetres (weft) and a centre-back length of 750 millimetres (warp), with an illustrative tolerance of plus or minus 10 millimetres on both. [S27] Suppose a generative pattern tool emits a base pattern whose pre-wash dimensions match those targets (540 millimetres half-chest and 750 millimetres back length), and an inferred preset assumes zero shrinkage. [S10] [S11]

Now assume the technician tests a swatch from the candidate lot (`swatch_id: SW-ILL-01`) across three wash-and-dry cycles and records an illustrative mean warp shrinkage of 4.0 percent and weft shrinkage of 2.5 percent. [S32]

1. **Unadjusted outcome if cut directly from the generated pattern**:
2. Illustrative post-wash centre-back length (warp): 750 mm multiplied by (1 - 0.040) equals 720.0 mm, 30.0 mm below target and breaching the illustrative -10 mm lower bound by 20.0 mm.
3. Illustrative post-wash half-chest width (weft): 540 mm multiplied by (1 - 0.025) equals 526.5 mm, 13.5 mm below target and breaching the illustrative -10 mm lower bound by 3.5 mm.

1. **Required pre-wash pattern cutting dimensions to hit the post-wash target**:
2. Illustrative required pre-wash centre-back length: 750 mm divided by 0.960 equals 781.25 mm (an illustrative warp allowance addition of 31.25 mm, or 31 mm rounded).
3. Illustrative required pre-wash half-chest width: 540 mm divided by 0.975 equals 553.85 mm (an illustrative weft allowance addition of 13.85 mm, or 14 mm rounded).

This illustrative arithmetic shows why an uncalibrated simulation is misleading: a pattern that looks proportioned on a 750 mm by 540 mm avatar before laundering fails buyer inspection in both warp and weft after one wash unless measured lot shrinkage is fed back into the pattern parameters. [S09] [S27]

6.13

Simulation failures must remain visible

A simulation engine is an engineering diagnostic tool, not a marketing renderer. [S19] When a pattern panel has an incompatible seam length, excessive curvature or an unrealistic stiffness parameter, the cloth solver may produce interpenetrating layers, mesh jitter, extreme membrane strain or unstable collar folding. [S06] In a production workflow, those numerical artefacts are warning signals that must be logged and investigated rather than hidden. [S20]

6.14 COLLISION RESOLUTION IN LEARNED SIMULATORS CAN MASK GEOMETRIC INFEASIBILITY

As sewing-pattern generation and multi-garment styling move towards neural and agentic pipelines, the distinction between physical equilibrium and learned visual plausibility becomes critical. [S06] [S18] In physics-based simulation, a discrete shell and contact solver enforces non-penetration constraints step by step; if a cuff opening is narrower than the wrist and lacks tensile compliance, the solver stalls, reports strain or distorts the mesh. [S02] By contrast, neural surrogate simulators and template-deformation methods predict draped three-dimensional meshes directly for speed, where handling collisions and intersections in neural multi-garment simulations remains a largely unsolved problem. [S06]

ContourCraft addresses this limitation in learned multi-garment simulation through collision and intersection losses alongside contour-aware resolution mechanisms to untangle interpenetrating layers. [S06] At the multi-garment reconstruction level, agentic systems such as DressAgent edit sewing patterns inside a closed loop with physical simulation and visual judgement to converge on layered outfit geometry. [S18] While automatic intersection recovery and closed-loop editing advance computer graphics, they introduce a specification hazard: if a post-processing network or mesh-relaxation step resolves a collision by locally stretching vertices or sliding a seam line without updating the two-dimensional cutting pattern and strain report, the three-dimensional preview looks clean while the physical garment remains unsewable. [S06] [S14]

6.15 SOLVER LOGS AND MESH DIAGNOSTICS MUST ACCOMPANY RENDERED PREVIEWS

To prevent visual cleanup from concealing failures, GFT recommends that every simulation run attached to a technical review package retain solver configuration and diagnostic telemetry alongside viewport images: [S20]

- **Solver and discretisation settings**: simulator name, version, time-step size, iteration count, gravity vector, damping coefficients and mesh edge length in millimetres. [S02] [S20]
- **Avatar and pose state**: avatar identifier, body measurement table revision, surface friction coefficient and static pose or motion sequence identifier. [S27] [S30]
- **Material record binding**: the exact `measured_property_id` and revision number consumed by each panel, confirming whether any piece fell back to a default material. [S32]
- **Collision and strain exceptions**: count of unresolved self-intersections, body penetrations, maximum warp and weft strain percentages, and any manual pinning, tacking or vertex sculpting applied by the operator. [S06] [S20]

UNRESOLVED CLOTH-INTERSECTION ARTEFACTS

Never accept two-dimensional image inpainting, manual three-dimensional vertex sculpting or an unlogged collision-offset increase as a fix for a simulation failure. [S06] [S20] If a collar stand interpenetrates the neckband, a pocket bag pokes through a front panel or an armhole bunches in simulation, the technician must determine whether the cause is an invalid seam correspondence (Chapter 5), an incorrect panel contour (Chapter 4), an uncalibrated parameter (`measured_property_id`) or a mesh-resolution artefact. [S02] [S12] Resolving the defect only in the mesh or rendered image leaves the two-dimensional DXF cutting boundary unchanged, guaranteeing that the defect reappears when cloth is sewn. [S22]

6.16

Physical checks must calibrate the intended use

Even a deterministic finite-element solver driven by laboratory-measured fabric parameters is a mathematical idealisation of a complex fibrous assembly. [S02] [S0602] Seams add multi-layer thickness, sewing thread introduces local tension and pucker, fusible interlinings stiffen collars and plackets non-uniformly, and human bodies compress and shift under clothing in ways that rigid avatars do not replicate. [S12] [S27] Therefore, a digital material workflow must be calibrated against physical reality for the specific garment category and wearer use case. [S19]

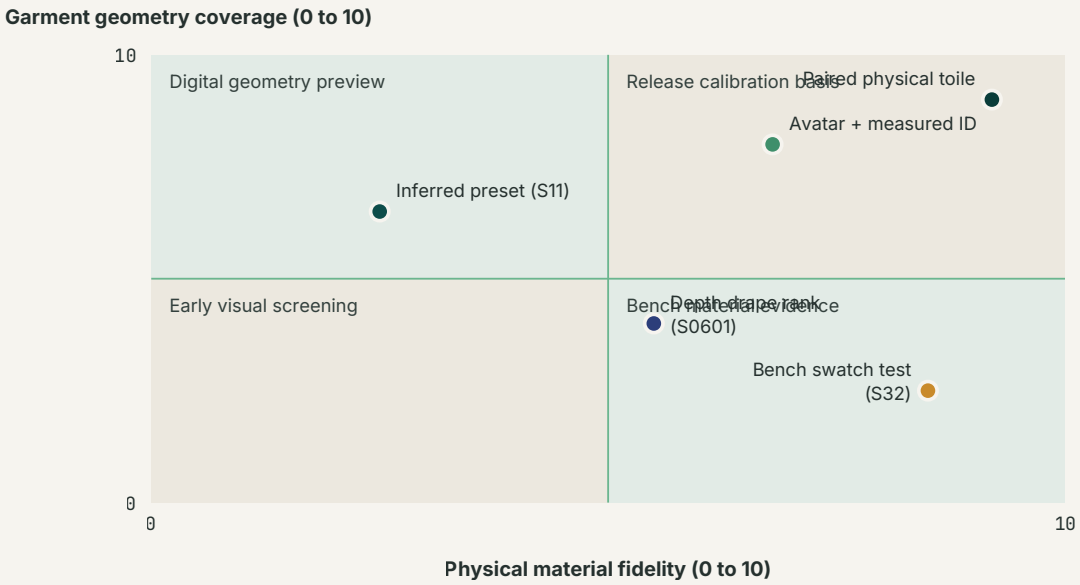
6.17 DRAPE SIMILARITY IN IMAGE SPACE DIFFERS FROM DIMENSIONAL CONFORMANCE

How should a technical team compare a simulated garment against a physical specimen? [S19] Research on capturing fabric mechanics from depth images shows that comparing estimated and ground-truth material parameters purely by Euclidean distance in numerical parameter space correlates poorly with human perception of how a fabric hangs; to address this, S0601 introduces a drape similarity metric operating in the image domain to rank perceptual similarity of cloth folds. [S0601] Conversely, the Fabric-101 study evaluates its differentiable elasto-plastic simulator against cyclic hysteresis curves and three-dimensional wrinkle geometry across over 101 fabrics to verify that persistent creases reflect physical deformation regimes. [S0602]

For an apparel factory, both perceptual drape similarity and mechanical hysteresis matter, yet neither substitutes for dimensional conformance at the specification's points of measure. [S19] [S27] A simulated shirt can exhibit convincing fold frequencies in an image-domain comparison while hanging ten millimetres longer at the hem than the physical sample because seam-allowance weight or shoulder-seam tape was omitted from the model. [S09] [S0601] GFT therefore separates calibration into four complementary checks, mapped across physical fidelity and geometric coverage in **Exhibit 25**, p. 102. [S11] [S27]

EXHIBIT 25 / ILLUSTRATIVE VERIFICATION SCOPE (0 TO 10)

A material check needs both an avatar and a physical swatch



ILLUSTRATIVE [S11] [S19] [S27] [S32] [S0601] Illustrative GFT comparison of verification methods across physical material fidelity (x-axis) and whole-garment geometric coverage (y-axis) on a 0 to 10 teaching scale. Positions are conceptual proposals informed by S11, S19, S27, S32 and S0601, not empirical trial scores.

In **Exhibit 25**, p. 102, an inferred preset (S11) sits in the upper-left quadrant: it spans full garment geometry on an avatar (y = 6.5 on the illustrative 0 to 10 scale) with low physical fidelity (x = 2.5) because the fabric is untouched. [S11] A bench swatch test (S32) sits in the lower-right (x = 8.5, y = 2.5), providing high physical fidelity on mass, thickness and shrinkage for a flat specimen without testing multi-panel construction. [S32] Depth-captured drape ranking (S0601) bridges flat tests and visual hang (x = 5.5, y = 4.0), while combining a calibrated avatar simulation consuming the measured record (x = 6.8, y = 8.0) with a paired physical toile (x = 9.2, y = 9.0) reaches the upper-right release quadrant. [S27] [S0601] The coordinates in **Exhibit 25**, p. 102 are illustrative conceptual placements on a 0 to 10 teaching scale, not empirical trial scores. [S19]

6.18 LOCAL EASE AND WEARER MOVEMENT MUST BE SEPARATED FROM FABRIC STRETCH

Calibrating a material model requires separating geometric ease from mechanical elongation. [S27] As the EASE study states, garment fit and comfort depend critically on ease, the local allowance of excess material relative to the body. [S27] When a wearer raises an arm or reaches forward in a woven shirt, the garment accommodates the pose first by redistributing geometric ease around the chest, back and armhole, and only secondarily by straining yarns and seams once slack is exhausted. [S27]

If a simulation underestimates tensile stiffness or overestimates shear compliance, a tight armhole with insufficient ease appears comfortable in a posed render because virtual cloth stretches across the back. [S11] [S27] Cut in low-stretch cotton poplin, that pattern binds and risks seam tearing. [S27] Following NIST AI RMF 1.0, which notes that a risk framework helps prioritise risk but does not prescribe risk tolerance, the

team and buyer must define dimensional and strain tolerances for each product family before relying on digital previews to reduce physical samples. [S19] No universal percentage agreement between digital and physical drape is imported here. [S19]

CALIBRATION PROTOCOL BETWEEN DIGITAL PRESET AND PHYSICAL SWATCH

To qualify a fabric lot (`supplier_lot_id`) and its simulation record (`measured_property_id`) on an approved base block, GFT proposes a five-step sequence:

1. **Condition and measure the flat specimen**: condition the swatch (`swatch_id`) under documented temperature and humidity, then record areal mass (`g/m2`), thickness (`mm`), cantilever overhang, bias extension and wash-dry shrinkage. [S32]
2. **Verify standard drape geometry**: mount a specimen on a physical drape stand, compare overhang radius and fold amplitude against a matching simulation, and log adjustments alongside raw test values. [S0601] [S0602]
3. **Simulate the approved base block under explicit ease**: assign measured parameters and grainlines to the base block on a matched avatar, keeping local ease maps (S27) separate from elastic strain maps. [S02] [S27]
4. **Construct and measure the physical reference toile**: cut and sew one reference garment from the same roll using the specified seams and interlinings from Chapter 5, then measure points of measure before and after washing. [S09] [S19]
5. **Reconcile deviations and lock the envelope**: record differences between simulated and physical pre-wash dimensions, and between pre-wash and post-wash dimensions; approve the material record for variant styling only within agreed tolerances. [S19] [S20]

6.19

A changed fabric can reopen a previously approved design

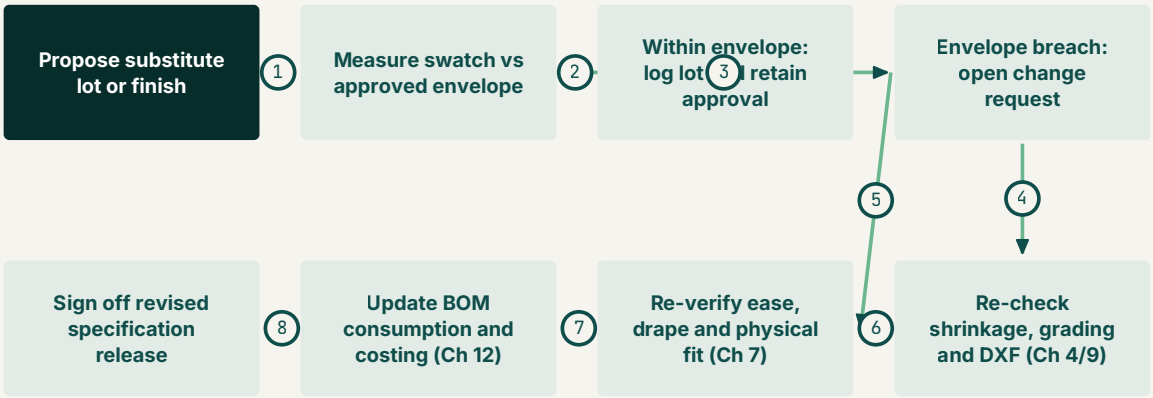
In factory operations, material selection rarely stays static between initial sign-off and bulk production. [S20] Yarn prices fluctuate, a dye lot runs out, a buyer requests a lighter weight for a repeat order, or a greenfield Cut-Make-Trim (CMT) line transitions from imported greige cloth to domestically woven Cameroonian cotton as local spinning and weaving partnerships mature. [R06] [R10] A specification system is safe only if it defines when a material change invalidates earlier pattern, fit and costing approvals. [S19] [S20]

6.20 ENVELOPE BREACHES TRIGGER MANDATORY DOWNSTREAM RE-EVALUATION

Under GFT's proposed governance model, an approved style specification is bound not to a generic fabric name ("100% cotton poplin") but to a bounded material envelope defined around the qualified `measured_property_id` and `supplier_lot_id`. [S20] [S32] Whenever procurement or a buyer proposes a substitute roll, finishing batch or mill, the quality team compares the candidate swatch against the approved envelope across five dimensions: areal mass, thickness, construction, directional stretch and recovery, and laundering shrinkage. [S32] [S0602]

EXHIBIT 26 / WORKFLOW STEPS

Material substitution is a controlled change



1. Assign swatch_id
2. Pass all tolerances
3. Exceed bound
4. Shrinkage/stretch shift
5. Drape/handfeel shift
6. Updated panels
7. Approved sample
8. Technicians sign off

ILLUSTRATIVE [S19] [S20] [S22] Illustrative GFT change-control workflow when a candidate fabric departs from the approved swatch envelope, routing back through pattern, fit, costing and release gates. Informed by S19, S20 and S22.

As **Exhibit 26**, p. 104 shows, if the candidate swatch stays within the tolerance envelope on all physical checks, the technician logs the new `supplier_lot_id` against the style revision and preserves earlier approvals. [S19] [S20] If any parameter breaches the envelope, **Exhibit 26**, p. 104 opens a change request routing the style to the affected gates:

- **Pattern and CAD exchange gate (Chapters 4 and 9)**: a shift in shrinkage or width requires recalculating allowances, updating graded sizes, re-exporting Standard DXF, AAMA or ASTM files (S22) and running a CAD readback check. [S22]
- **Seam and construction gate (Chapter 5)**: a change in thickness, slippage or weight requires checking stitch density, needle size, seam allowance and interlining. [S12]
- **Fit and ease gate (Chapter 7)**: a change in bending, shear or elastane alters hang and local ease, requiring an updated simulation and, where thresholds are crossed, a physical fit sample. [S27]
- **Surface artwork and colour gate (Chapter 8)**: a change in whiteness, mercerisation or texture requires a physical strike-off or lab dip rather than digital relighting previews. [S28] [S32]

- **Bill of materials and economic gate (Chapter 12)**: changed roll width, shrinkage allowance or rework risk alters fabric consumption per garment and updates costing. [S19]

While **Exhibit 26**, p. 104 structures fabric changes, it is an illustrative GFT governance workflow and does not replace buyer approval contracts. [S19] [S20]

> **GFT NOTE**
>
> GFT recommends never overwriting an existing material record (measured_property_id) in place when a supplier changes a finish or when a lab test replaces an initial model estimate (S11). [S11] [S20] Create an immutable revision, retain the earlier record linked to its prototype runs, and require technician sign-off before binding the new revision to a tech pack. [S20] In a phased roadmap where data capture and inspection precede generative pilots, material versioning costs little and prevents unrecorded fabric drift. [S19] [S32]

6.21 MATERIAL RECORDS CONNECT DESIGN SPECIFICATIONS TO RESOURCE AND CARBON BOUNDARIES

Linking digital specifications to physical material lots also provides the physical denominator for resource and environmental accounting. [S29] In the GFT library, wet-processing water, chemical and energy measurement belongs to wp06, cotton traceability to wp05, waste routing to wp11, and product carbon accounting to wp16. [R06] [S29] As the Greenhouse Gas Protocol Product Standard explains, product lifecycle rules require consistent boundaries before any comparative or inventory claim is supported. [S29]

A generative tool cannot substantiate an environmental saving merely by rendering a garment or labelling a texture "organic cotton". [S28] [S29] Only when the tech pack records the verified supplier_lot_id, measured areal mass (g/m2), shrinkage-adjusted marker consumption and physical sample rounds sewn (Chapter 12) can wp06 and wp16 calculate material mass, liquor ratio or embodied emissions. [S29] [S32] Anchoring simulation to a measured swatch carries specification discipline into production control and sustainability reporting. [R06] [S29] Nothing in this chapter constitutes investment, legal or tax advice, and no physical test, software certification or factory trial is claimed here. [S19] [S20]

A pixel can suggest a weave, but only a measured swatch can authorise the blade.

CHAPTER SOURCE NOTES

S11 Authors via arXiv. 2026-01-14.

S32 CLO Virtual Fashion, CONNECT Support. 2025-04-17.

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

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

S28 Authors via arXiv; SIGGRAPH Asia 2026 acceptance listed. 2026-09-24.

S01 Authors via arXiv. 2025-09-04.

S0601 Authors via arXiv. 2023-04-13.

S0602 Authors via arXiv. 2026-09-12.

S20 National Institute of Standards and Technology. 2024-07-26.

S19 National Institute of Standards and Technology. 2023-01-26.

S05 Authors via arXiv; 3DV 2025. 2024-03-27.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S27 Authors via arXiv; Computers & Graphics. 2026-06-28.

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

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S22 CLO Virtual Fashion. 2026-07-15.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S06 Authors via arXiv; SIGGRAPH 2024. 2024-05-15.

S18 Authors via arXiv. 2026-10-11.

S14 Authors via arXiv. 2026-06-23.

S12 Authors via arXiv; WACV 2026. 2026-02-25.

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

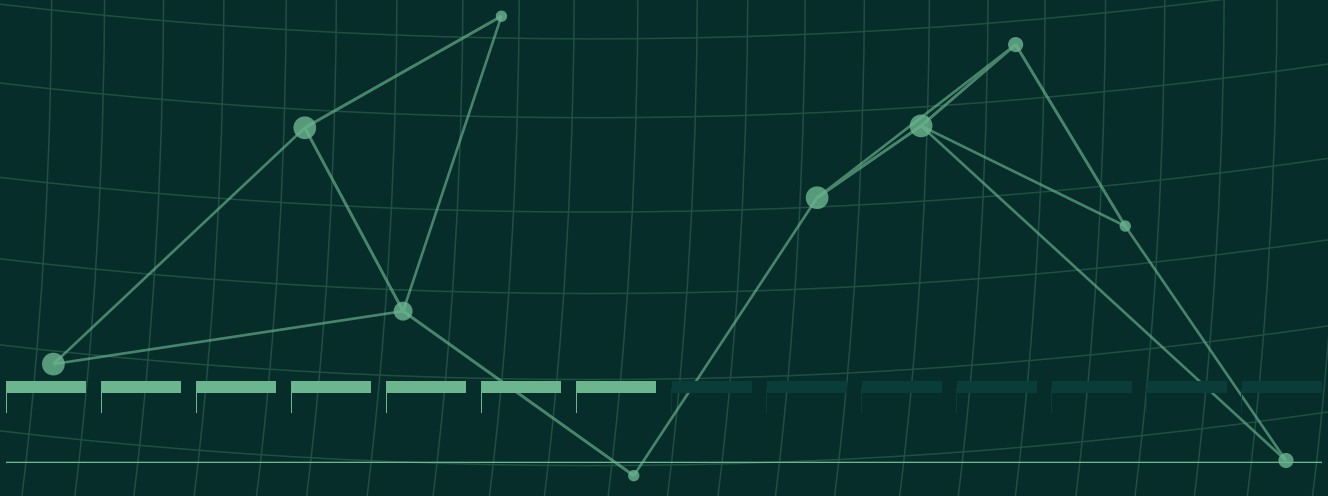
S29 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2026.

07

PART II: SPECIFICATIONS REQUIRE GEOMETRY, MATERIALS AND FIT / BODIES AND EASE

A well-fitted image is not evidence that the garment fits

Fit review needs actual garment and wearer dimensions, plus explicit ease and movement requirements.



CHAPTER 07
DECISION TEST

Fit review needs actual garment and wearer dimensions, plus explicit ease and movement requirements.

A rendered garment can look smooth, balanced and correctly proportioned while representing the wrong size for the person shown. Green Future Tech Limited (GFT), Hong Kong SAR, therefore recommends that a greenfield apparel maker treat virtual try-on as a visual review until body measurements, finished garment dimensions, local ease and intended movement are explicit inputs to the evaluation. [S13] [S27] Fit approval should remain a separate physical decision, supported by a measured sample across the proposed size range and by records that distinguish dimensional failure, fabric behaviour, wearer response and buyer acceptance. [S19] The central control is simple: no image, avatar or generated approval field can substitute for evidence about the actual body, the actual pattern and the actual cloth.

7.1

Appearance-based try-on can hide a size mismatch

Appearance-based virtual try-on learns to produce a plausible picture of a person wearing a garment. That objective does not necessarily require the system to know either the person's dimensions or the garment's size. The authors of the FIT dataset identify the resulting failure directly: current methods in their reviewed setting tend to generate well-fitted results regardless of garment or person size, partly because training data rarely contain precise measurements and deliberately ill-fitting pairs. [S13] This is not an audit of every commercial product. It is a precise warning about what visual plausibility can conceal.

7.2 AN IMAGE OBJECTIVE CAN REPAIR THE VERY EVIDENCE A FIT REVIEW NEEDS

A fit technician looks for drag lines, gaping, restricted movement, seam displacement and insufficient or excessive volume. An image generator optimised for a clean result may instead remove wrinkles, redistribute cloth or regularise the silhouette until those signals disappear. [S13] The output can then be aesthetically preferable precisely because it has suppressed the diagnostic evidence. A reviewer who rewards only visual quality teaches the workflow to hide fit problems rather than detect them.

This failure is different from an inaccurate avatar. A dimensionally correct avatar can still receive a garment whose digital pattern has no verified finished measurements. Conversely, an accurate pattern can be placed on an avatar whose chest, shoulder, arm length or posture does not represent the intended wearer. The two errors can also cancel visually: excessive garment ease on an undersized avatar may resemble an approved block on a correctly sized one. A single polished render cannot identify which combination produced it.

GFT proposes that every virtual fit record carry the body representation identifier, its measurement revision, the garment pattern revision, the stated size, the material revision and the simulation configuration. These are proposed evidence fields, not features attributed to every try-on system. If any field is absent, the output remains an appearance study. It may guide a design conversation, but it cannot support a dimensional conclusion.

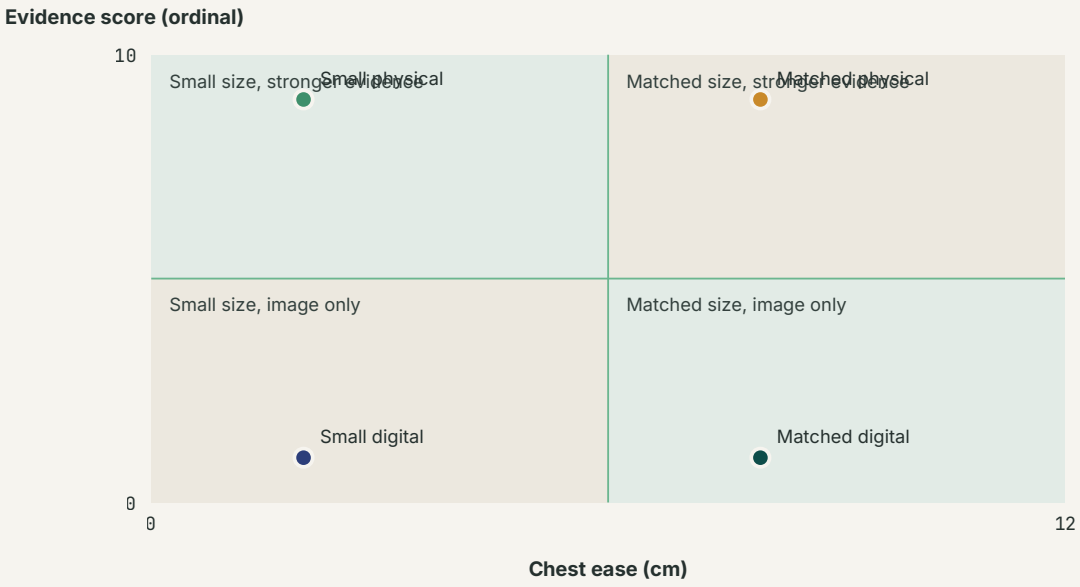
7.3 DELIBERATELY WRONG PAIRINGS REVEAL WHETHER SIZE INFORMATION MATTERS

A meaningful evaluation needs counterfactuals. The team should present the same body with the intended garment size, a smaller size and a larger size while holding pose, lighting, camera and material constant. It should then ask whether the system preserves the dimensional difference or normalises every result towards the same attractive fit. [S13] Incorrect pairings must be part of the test set, not rare accidents discovered after procurement.

The illustrative design in **Exhibit 27**, p. 109 combines those pairings with two levels of evidence. The chest measurements are teaching assumptions, not preferred ease values or observations of any wearer. A low evidence score denotes an image-only review; a high score denotes the chapter's proposed measured sample, dimensional readback and activity review. The score is an ordinal design device, not a measured model accuracy.

EXHIBIT 27 / CM CHEST EASE AND ILLUSTRATIVE EVIDENCE SCORE

Fit assessment needs matched and deliberately mismatched cases



ILLUSTRATIVE [S13] [S27] Illustrative test design. Body chest is assumed as 100 cm. Matched and small garments are assumed at 108 cm and 102 cm, giving 8 cm and 2 cm ease. Evidence scores are ordinal: 1 means image-only review and 9 means dimensional, physical and activity evidence. No wearer outcome or model accuracy is asserted.

The important comparison in **Exhibit 27**, p. 109 is vertical as well as horizontal. Moving from a small garment to the matched garment changes a dimensional input. Moving from digital preview to physical evaluation changes the evidence available to the reviewer. A procurement test should report both movements separately. If two garment sizes produce visually similar outputs, the finding concerns size sensitivity, not whether the smaller garment is comfortable. If the physical small sample restricts an agreed movement, that finding belongs to the physical fit record, not to an image-quality leaderboard.

7.4 PROCUREMENT QUESTIONS SHOULD EXPOSE THE EVALUATION BOUNDARY

Before adopting a system, the maker should request its input schema and evaluation cases. Does it consume numerical body and garment measurements, or only photographs? Can a reviewer supply an incorrect size deliberately? Does the evaluation set include such mismatches, and are outcomes reported by pairing rather than only in aggregate? [S13] Can the team export the garment measurements used for a result and reconcile them with the pattern revision?

These questions follow the contextual evaluation logic of the NIST AI Risk Management Framework: test conditions should represent intended use, and the organisation must set its own tolerance for residual risk. [S19] A vendor demonstration with one attractive, correctly matched image answers none of them. The buyer needs retained cases, declared denominators and the freedom to record a failure.

7.5

A target size range is a product decision

A size range is not a neutral technical setting inherited from a model. It is a product commitment that connects intended wearers, order volumes, grading effort, sampling work and inventory. The maker must decide which bodies and use cases the product is meant to serve before asking software to fit them. An avatar library or synthetic body model can support that work, but it cannot decide the commercial range or establish the distribution of Cameroonian wearers.

7.6 A SYNTHETIC BODY DISTRIBUTION IS NOT A LOCAL WEARER CENSUS

GarmentCodeData illustrates both the value and the boundary of synthetic evidence. The research pipeline samples body shapes from a custom statistical body model based on the European subset of CAESAR, generates sewing patterns, drapes them and applies quality controls to remove invalid samples. [S30] This creates controlled variation with known measurements and pattern labels. It does not turn the resulting distribution into an anthropometric survey of Cameroon, a buyer's size curve or evidence about a particular workforce.

The distinction matters when a generated result transfers poorly. A system may have learned the design conventions and body-shape coverage of its development data rather than the proportions, postures or product preferences of the intended wearer group. [S30] The answer is not to discard synthetic data. It is to declare its coverage, test beyond familiar cases and collect local evidence where the product decision requires it. No Cameroonian anthropometric dataset is supplied in this paper, so no local size distribution or preferred garment dimension is asserted here.

For the illustrative cotton shirt, the product team should first record the intended wearer task, whether the shirt is for uniform or general use, the planned size labels, the required fit character and the activities that must be possible. Procurement and planning should then estimate the intended size mix through their normal demand process. That estimate belongs to the product plan, not to the fit model. A rare but required size still needs a controlled grading and approval route.

7.7 MEASUREMENT DEFINITIONS MUST SURVIVE COLLECTION AND REUSE

Local wearer measurement should begin with a protocol rather than a scanner purchase. Each dimension needs a name, landmarks, tape or sensor method, posture, clothing condition, unit and reviewer. Chest circumference measured over a thick garment is not interchangeable with a measurement over close-fitting clothing. A shoulder width taken between different landmarks cannot be merged into the same field without qualification. These are measurement-definition requirements proposed by GFT, not numerical garment standards.

The record should also distinguish a direct body measurement from an inferred scan measurement and from a synthetic avatar parameter. [S30] Repeated measurements can be retained where the protocol requires a reliability check, with discrepancies resolved before pattern work. The technician needs the final approved value and its revision, while the wearer panel coordinator needs the consent and participation record. Combining those responsibilities in one unrestricted file would expose more personal information than the fit task requires.

A minimum viable size-range dossier therefore contains four linked but distinct artefacts: the product range decision, the measurement protocol, the consented wearer evidence and the approved body representations used in digital work. Missing evidence stays visible as missing. The team should not fill gaps with an unnamed "standard African body" or convert an available avatar into a population claim. Synthetic bodies remain useful test instruments when their assumptions are documented and their limits are tested. [S30]

7.8 SIZE INCLUSION SHOULD BE TESTED AT THE DECISION POINTS

The chosen range must survive more than base-size pattern generation. A style can fail because grading produces an unsuitable armhole at an extreme, because a closure no longer aligns, because fabric behaviour changes the intended drape, or because a digital avatar lacks the required posture. The pilot should therefore include the range endpoints and an intermediate size before the product is released, with additional sizes selected where the technician identifies a nonlinear grading or construction risk.

This is a product-coverage rule, not a claim that three sizes statistically represent every wearer. The eventual test set follows the actual range and risk. Its purpose is to expose where the specification stops being valid. The release record should name the authorised size interval and any exceptions rather than allowing a successful base-size sample to imply approval of every label in the catalogue.

7.9

Ease should be a deliberate specification

Ease is the local allowance of garment material relative to the body. The EASE research framework makes it an explicit, spatially varying and anisotropic design variable rather than an incidental result of geometry or simulation. [S27] That distinction is central to production control. If a team cannot state where material allowance is intended, it cannot tell whether a smooth render represents the design intent, an arbitrary avatar relationship or a solver adjustment.

7.10 A SINGLE CHEST ALLOWANCE CANNOT DESCRIBE THE WHOLE SHIRT

A finished chest circumference is useful, but fit also depends on how volume and length are distributed around the body and across connected panels. Shoulder slope, armhole depth, sleeve cap geometry, cuff opening, neck circumference and hem allowance interact. Two shirts with the same chest circumference can position that volume differently and produce different movement and appearance. The release dossier should therefore use an ease map linked to defined points of measure, not one global label such as "regular".

GFT proposes separating three intentions in the illustrative shirt record. Functional allowance supports the agreed activity. Design allowance creates the chosen silhouette. Construction allowance supports assembly and finishing, but belongs to the seam specification rather than being silently counted as wearer ease. These categories are a proposed review language. Their values and tolerances must come from the technician, buyer and product use case, not from this teaching chapter.

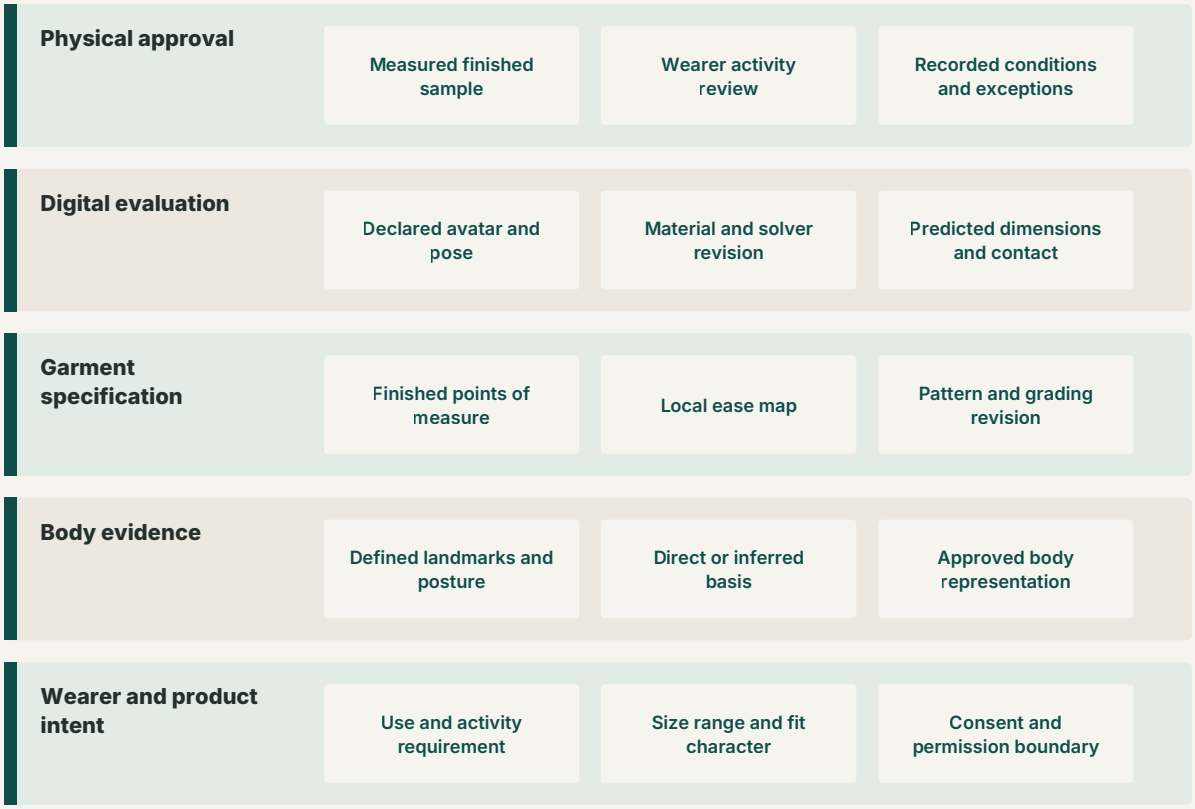
The measurement state also needs a name. A pattern dimension, a predicted digital finished dimension and a measured physical sample dimension are related but not identical. Fabric distortion, sewing and finishing can change the result. The reviewer should compare like with like: body dimension to specified finished garment dimension for intended ease, specified point of measure to physical point of measure for manufacturing conformity, and digital prediction to physical observation for model calibration. [S27]

7.11 THE EVIDENCE CHAIN SHOULD CONNECT BODY, PATTERN AND VALIDATION

Exhibit 28, p. 113 shows the proposed chain from wearer requirement to physical approval. Its purpose is to prevent a field from being copied across boundaries without reinterpretation. A body chest value is not a finished garment chest value. An ease value is not a seam allowance. A simulated displacement is not a wearer assessment. Each layer has its own evidence and owner, while identifiers preserve traceability between them.

EXHIBIT 28 / EVIDENCE LAYERS

Body measurements and ease connect intent to evaluation



ILLUSTRATIVE [S27] [S30] Illustrative GFT evidence architecture informed by explicit local ease in S27 and body-shape assumptions in S30. It is a proposed record structure, not an implemented GFT system or a claim of physical validation.

The chain in **Exhibit 28**, p. 113 also makes correction more precise. A wrong body landmark returns to the measurement protocol. A wrong finished dimension returns to the pattern or grading record. A mismatch between digital and physical dimensions returns to simulation calibration and construction review. A garment that measures correctly but fails an agreed reach test returns to product intent and ease distribution. The response follows the failed relationship rather than a generic request to improve the image.

WORKED EXAMPLE: CIRCUMFERENCE, EASE AND FLAT WIDTH ARE DIFFERENT QUANTITIES

For an illustrative teaching calculation, assume a wearer chest circumference of 100 cm and a specified finished shirt chest circumference of 108 cm. Total chest ease is 108 minus 100, which equals 8 cm. A deliberately smaller test garment at 102 cm would instead have 2 cm total ease. If the point of measure is a flat half-chest width with the shirt laid in the agreed position, the corresponding nominal widths are 54 cm and 51 cm, obtained by dividing each finished circumference by two. Every value is illustrative, no preferred allowance is implied, and the physical measuring method must define how the closed garment is laid and aligned.

The worked example exposes two common specification errors. First, comparing a 100 cm body circumference directly with a flat garment width would mix geometric bases. Second, reporting only the resulting 8 cm allowance would omit where that allowance sits and whether it supports the intended activity. A technician still needs the full point-of-measure table, the local ease intention and the physical sample result. [S27]

7.12 EASE TRANSFER REMAINS A HYPOTHESIS UNTIL RETESTED

Research that transfers an ease field to another body shape offers a useful way to preserve design intent during pattern generation. [S27] It does not prove that the transferred garment will fit a local wearer, that the selected fabric will reproduce the digital behaviour or that every size can share one rule. The maker should treat each transferred result as a candidate pattern requiring dimensional checks and physical confirmation.

A controlled change rule should reopen fit review when the body representation, size, fabric, construction or activity requirement changes outside its approved envelope. This avoids two opposite errors: freezing a sound pattern because an irrelevant visual asset changed, and preserving fit approval after a material or grading change that could alter the garment. The revision record should state which evidence remains applicable and which test must be repeated.

7.13

Grading requires checks beyond the base size

A base-size sample demonstrates that one realised garment can satisfy the agreed conditions. It does not demonstrate that the rules used to create every other size preserve fit, construction or appearance. Grading changes related dimensions according to a rule set; it is not a uniform enlargement or reduction of the whole pattern. The technician must inspect how those rules affect each point of measure and how connected features continue to work together.

7.14 SIZE EXTREMES EXPOSE DIFFERENT FAILURES FROM AN INTERMEDIATE SIZE

The smallest and largest planned sizes test the boundaries of the grading rules. An intermediate size can reveal a discontinuity or an implementation error that endpoints alone would not explain. For the illustrative shirt, the reviewer should check chest and shoulder dimensions, armhole and sleeve relationships, neck and collar closure, pocket and placket placement, and hem balance. The exact list follows the design, while the principle is consistent: measure the released artefact, not just the base-size intention.

Nested patterns are useful for visual review, but a technician also needs the numerical point-of-measure table and seam checks. A notch that remains geometrically present can move to an unhelpful position. A sleeve can retain its stated length while its cap no longer matches the armhole as intended. A pocket can remain attached to the front panel yet become poorly placed relative to the closure or wearer. These are proposed inspection categories, not reported defect frequencies.

Tolerance belongs to this review, but it must be agreed for the product and point of measure. NIST's framework explicitly leaves risk tolerance to the organisation and its context rather than prescribing a universal threshold. [S19] GFT therefore recommends that the maker and buyer approve dimensional tolerances before seeing pilot outcomes. A generated table should never fill an unknown tolerance with a plausible industry-sounding number.

7.15 EXPORT SUPPORT DOES NOT PROVE THAT GRADING SURVIVED EXCHANGE

CLO's documentation lists Standard DXF, AAMA and ASTM options, a selectable size group, units, grading export choices, notch controls and cutting-line or sewing-line handling. [S22] Those controls make exchange settings inspectable. They do not prove that a particular receiving system will reconstruct every size, annotation and relationship correctly. The maker must test the actual source and receiving releases rather than treating a format label as an acceptance result.

Exhibit 29, p. 115 sets out a proposed review grid. The size labels are deliberately generic because the teaching example has no authorised product range. Each row links the fit evidence to the export evidence, so a physical fit pass cannot conceal a broken receiving file and a successful file opening cannot conceal a failed garment.

EXHIBIT 29 / REVIEW STATES

A grading review must survive export

Size state	Fit and construction review	Export declaration	Receiving-CAD readback	Release evidence
Base	All POMs, seams, closure and activity	Size group, unit, format and rule file	Dimensions, notches, allowances and labels	Approved sample and source revision
Intermediate	Rule continuity and feature placement	Included grading and size identity	Intermediate POMs and seam matching	Exception or pass with reviewer
Lower extreme	Small-end closure, armhole and movement	Endpoint included in selected group	Endpoint geometry and annotations	Physical result and conditions
Upper extreme	Large-end balance, sleeve and movement	Endpoint included in selected group	Endpoint geometry and annotations	Physical result and conditions

ILLUSTRATIVE [S19] [S22] Illustrative GFT review grid. Generic size states avoid inventing a range. S22 supports documented export controls, not tested interoperability. S19 supports context-specific acceptance criteria. No actual readback or fit result is asserted.

The readback in **Exhibit 29**, p. 115 should compare measurements in the receiving workflow with the authorised source table. It should also verify that the selected export contains all intended sizes, that units are unchanged, that notches and seam allowances retain their meaning, and that size labels map to the correct geometry. [S22] A technician should record pass, failure or an explained exception against the pattern revision. Bulk release then uses the readback-approved file, not an untracked attachment sent alongside it.

7.16

Pose and intended activity change the acceptance question

A neutral standing pose can show balance and silhouette, but it cannot represent every demand placed on a working garment. EASE studies pose adaptation that redistributes local allowance to relieve strain in high-stretch regions, while retaining physics-based simulation for final visualisation. [S27] This makes pose an explicit design and evaluation variable. It does not establish guaranteed comfort, safety or freedom of movement for a physical wearer.

7.17 THE ACTIVITY BRIEF SHOULD PRECEDE THE SIMULATION

For an illustrative uniform-shirt review, the product team might agree reaching, sitting and fastening checks with the technician and wearer panel. These are proposed teaching activities, not a universal test suite. Each activity needs a defined starting position, movement, observation point and acceptance condition. A reviewer should know whether the concern is hem lift, closure gaping, sleeve restriction, pocket access or another product-specific outcome before a result is classified.

The digital protocol should retain the same activity identifiers as the physical protocol. If the avatar performs a reach that the physical panel cannot reproduce, the two outcomes are not directly comparable. If the physical wearer uses a different movement because the digital brief was ambiguous, the deviation belongs in the record. The purpose is not theatrical realism. It is to ask the same fit question at both boundaries.

7.18 VISUAL CONTACT AND WEARER RESPONSE REMAIN SEPARATE EVIDENCE

Simulation can help inspect predicted contact, deformation and local strain under declared material and pose assumptions. [S27] A wearer can report pressure, restriction, thermal sensation or preference. Those observations may disagree. A garment can show no visible digital contact problem yet feel restrictive, or show an unattractive fold that does not impede the task. The release record should preserve both forms of evidence rather than collapsing them into one score.

Material identity is equally important. A pose adaptation computed with one material preset should not be transferred silently to a different physical cloth. The material chapter's measured-swatch rule therefore remains active during movement review. A changed fabric or finish can require a new digital run and a repeated physical activity check even when the pattern coordinates are unchanged.

> **GFT NOTE**
>
> Write the activity acceptance card before generating the first fit image. Name the intended movement, body and garment revisions, material, observations and decision owner. Reject any automatic approval that arises only because a rendered pose looks natural. Keep digital prediction, measured garment conformity and wearer evaluation as separate fields, then require the accountable fit reviewer to reconcile them.

7.19

Wearer data needs its own permission boundary

Fit work can require photographs, body scans, dimensions, avatars and subjective wearer comments. These records concern identifiable people and may be more intrusive than the design files that a generative system normally handles. NIST's Generative AI Profile identifies privacy risks involving sensitive information, unauthorised use, disclosure and inference, and recommends lifecycle actions rather than a single privacy switch. [S20] A fit workflow therefore needs a purpose and permission record before collection begins.

7.20 COLLECTION SHOULD BE LIMITED TO THE DECLARED FIT PURPOSE

GFT proposes a data inventory that separates the fit task from optional secondary uses. The core record should contain only the measurements and activity observations needed for the defined product review. A body photograph should not become the default when a tape measurement and posture note will answer the question. Where a scan or image is necessary, its identifier, purpose, access group and retention rule should travel with it.

Consent for a fit session is not consent to train a model, publish an image or reuse an avatar for another product. Those destinations need separate, explicit permission decisions. A withdrawn or expired permission should be visible to the system and to the responsible coordinator, with affected derivatives identified for review. This is a governance design, not a conclusion about the legal effect of any consent form in Cameroon, Hong Kong SAR or another jurisdiction.

The measurement protocol should also minimise unnecessary identification. The pattern technician can work with an approved body representation and dimensions without receiving a wearer's name, contact details or full image set. The panel coordinator can hold the participation and permission mapping under a separate access rule. Role separation reduces exposure while preserving the link needed to investigate a measurement or approval dispute.

7.21 EXTERNAL MODELS AND RETENTION RULES NEED EXPLICIT DESTINATIONS

Before body data enters an external service, the maker should record the provider, model or service version, processing location disclosed by the provider, retention setting, training setting and export or deletion route. Unknown terms remain unresolved procurement items. The NIST profile's attention to third-party components and lifecycle risk supports checking the whole service chain rather than only the visible application. [S20]

Retention should follow the declared decision. A long-lived approved block may need a dimensional validation summary, but that does not automatically justify retaining every source photograph or raw scan. The team can preserve the approved measurements, protocol, reviewer and result while deleting source material under the agreed rule. Any exception should have an owner and a reason rather than becoming permanent storage by default.

Access reviews should include exported avatars and test datasets, not only the original database. A synthetic-looking avatar can remain traceable to a person if its identifiers or source files reveal the connection. Training or benchmark use requires a fresh destination check, and public presentation requires a separate publication decision. This chapter does not provide legal advice; applicable privacy, employment and buyer obligations require jurisdiction-specific review before implementation.

7.22

Physical fit approval remains a release condition

A digital fit prediction is useful when it helps a technician find a problem before cloth is cut. It becomes dangerous when its success is treated as proof that the physical garment needs no further review. GFT recommends retaining a physical fit sample, measured deviations and explicit approval conditions before bulk release. This is a proposed production gate, not a claim that a GFT factory or local wearer trial has already operated.

7.23 THE PHYSICAL REVIEW SHOULD TEST THE RELEASED SPECIFICATION

The sample must carry the same style, pattern, grading, material and construction revisions that the release dossier identifies. Before the wearer tries it, the reviewer should measure the agreed points of measure using the authorised method and compare them with the specification and tolerances. A garment that already fails dimensional conformity should not be presented as a successful wearer fit simply because it looks acceptable when worn.

The fit session then follows the agreed static and activity protocol. The reviewer records the body representation or wearer group relevant to the decision, the sample size, observed deviations, wearer feedback and any condition attached to approval. The buyer's visual or commercial approval remains a separate field. So does the technician's dimensional approval. Combining them into one affirmative flag would make it impossible to know which question was answered.

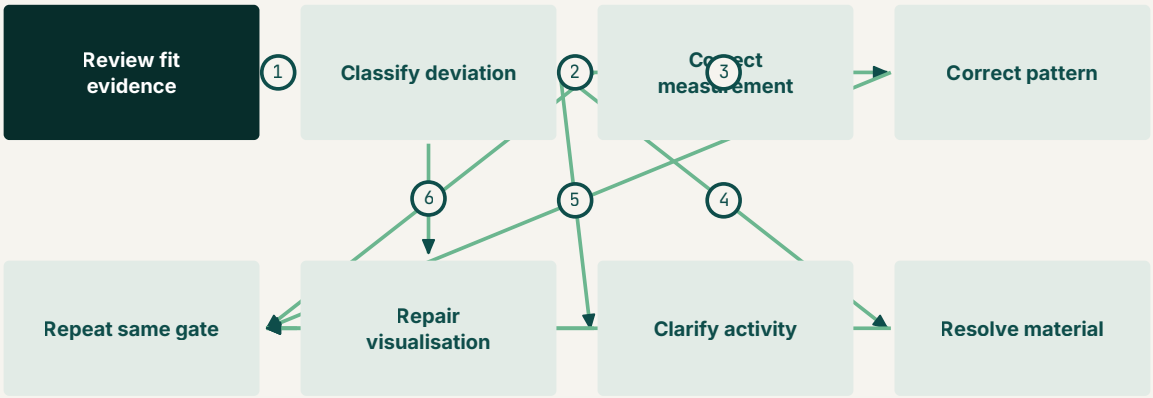
Where a physical result differs from the digital prediction, the team should retain both. The deviation is calibration evidence, not an inconvenience to delete. It may arise from body measurement, pattern geometry, fabric behaviour, construction, avatar pose, simulation settings or the measuring method. The appropriate response depends on the cause. A better-looking rerender is not a valid correction unless the underlying evidence has changed and the same test passes.

7.24 A REJECTION SHOULD PRESERVE ITS CAUSE AND CORRECTION OWNER

Exhibit 30, p. 119 proposes a return path for common failure classes. It deliberately avoids a single generic "AI failure" category. Measurement problems return to the protocol owner, pattern problems to the pattern technician, material problems to the material and fit reviewers, activity ambiguity to the product owner, and visualisation failures to the digital-system owner. Every route ends at the same evidence gate before release.

EXHIBIT 30 / DECISION FLOW

A fit rejection should preserve the cause



1. Deviation
2. Body data
3. Geometry
4. Cloth
5. Use case
6. Simulation

ILLUSTRATIVE [S19] [S20] [S27] Illustrative GFT rejection and correction flow. Categories are proposed review codes, not measured defect shares. S19 and S20 inform contextual evaluation and lifecycle records; S27 informs pose and ease review. No local trial has been performed.

The flow in **Exhibit 30**, p. 119 is not a requirement to retest everything after every correction. The record should identify the affected evidence and the conditions that remain valid. A corrected body landmark may require a revised avatar and digital run. A changed sleeve pattern may require grading, seam, material-use and physical fit checks. A corrected display setting may require only visual reconciliation if the underlying geometry and physical approval are unchanged. The accountable reviewer documents that scope.

7.25 FIT METRICS MUST RETAIN THEIR DENOMINATORS

A maker should track digital-fit predictions, physical sample outcomes and buyer approvals independently. Useful proposed measures include evaluated physical samples with complete measurement records, samples passing agreed fit conditions, deviations by cause, repeat rounds by cause and correction time by accepted style. Each rate needs its evaluated denominator, size and product coverage, and missing records must remain visible. [S19]

A successful try-on image must never auto-fill the physical fit or buyer approval field. Nor should a digital pass be counted as an avoided physical round unless the authorised release process genuinely no longer requires that round for the same product

and risk conditions. This preserves the economic analysis in the later chapter: sampling savings can be counted only after fit evidence establishes that acceptance has not been weakened.

The release condition is consequently conservative but practical. Permit bulk release only when the intended sizes have the required dimensional evidence, the agreed physical fit and activity checks are closed, deviations have owners and dispositions, and the receiving specification matches the approved revisions. [S19] [S22] Until then, generative and simulation tools remain aids to inspection and learning. They do not own the decision to cut.

Fit is established by measured bodies, garments and movement, not by the persuasiveness of the image.

CHAPTER SOURCE NOTES

S13 Authors via arXiv; SIGGRAPH 2026. 2026-04-09.

S27 Authors via arXiv; Computers & Graphics. 2026-06-28.

S19 National Institute of Standards and Technology. 2023-01-26.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S22 CLO Virtual Fashion. 2026-07-15.

S20 National Institute of Standards and Technology. 2024-07-26.

08

PART II: SPECIFICATIONS REQUIRE GEOMETRY, MATERIALS AND FIT / COLOUR AND PROVENANCE

Surface generation needs print geometry and cultural permission

A coherent digital pattern still needs repeat, colour and provenance checks before it becomes a product specification.

CHAPTER 08
DECISION TEST

A coherent digital pattern still needs repeat, colour and provenance checks before it becomes a product specification.

A generative model can synthesise a fabric motif or project a texture across a three-dimensional garment mesh in seconds, yet a convincing image is not an instruction that a printworks or cutting room can execute. [S28] Green Future Tech Limited (GFT), Hong Kong SAR, recommends treating generated surface design as a candidate concept until three boundaries are satisfied: repeat geometry and panel-placement coordinates must be defined in physical units across sewn joins, colour targets must be verified against physical swatches or strike-offs under controlled viewing conditions, and provenance records must separate model generation history, community cultural permission and physical cotton chain of custody. [S23] [S28] [S32] Within Cameroon's cotton-to-garment sector, where the inherited baseline records approximately 95 percent of cotton fibre exported raw and about 5 percent processed locally against a 50 percent local-processing target by 2030, capturing domestic value in printed or yarn-dyed apparel requires disciplined specifications rather than unverified screen imagery. [R06] [R07] Nothing in this chapter constitutes legal, intellectual-property, tax or investment advice. [S21]

8.1

The word pattern describes two different objects

In English and in generative prompts, the word *pattern* refers to two different engineering artefacts: the planar cutting geometry of garment panels and the recurring visual motif printed, woven or dyed onto cloth. [S02] [S28] Conflating those two objects under a shared label or file path means an edit to a visual motif can overwrite graded cutting lines, or a re-graded panel can distort a print repeat without triggering review. [S22] [S28]

8.2 PLANAR CUTTING GEOMETRY AND SURFACE GRAPHIC MAPS SOLVE DISTINCT PROBLEMS

Research from 2023 to 2026 separates panel geometry from surface appearance. [S02] [S28] On the geometric side, GarmentCode compiles parametric programs into planar panels with boundary edges and stitching instructions, while Design2GarmentCode and PatternGSL treat sewing panels, edge loops and stitch topology as explicit geometric structures bridging design intent and cutting. [S02] [S09] [S14] On the surface side, IMAGGarment addresses controllable synthesis across silhouette, colour and logo placement, OmniFabric synthesises textures within sewing-pattern ultraviolet (UV) coordinate space to remove lighting contamination across a reconstructed three-dimensional garment, and DressCode pairs a sewing-pattern generator with a separate texture branch. [S04] [S26] [S28]

Cutting tables and printing machines consume distinct instructions from these branches. [S22] [S28] A cutting room requires closed two-dimensional panel polygons, cut lines, seam lines, grain-direction vectors, notches and size grading rules in Standard DXF, AAMA DXF or ASTM DXF. [S22] A printer or screen engraver requires a repeatable artwork tile, physical repeat dimensions in millimetres along warp and weft, colour separations or spectral targets, and a placement rule when a graphic sits at a panel landmark. [S26] [S32]

8.3 SEPARATE IDENTIFIERS AND STATE MACHINES PREVENT SILENT OVERWRITES

To prevent ambiguity in the specification spine, GFT proposes maintaining separate primary keys, filenames and approval state machines for sewing geometry (`panel_geometry_id`) and surface artwork (`surface_artwork_id`), linked on the teaching cotton shirt only through an explicit bill-of-materials mapping (**Exhibit 31**, p. 123). [S02] [S20] [S22] [S28]

EXHIBIT 31 / SPECIFICATION FIELDS

Surface artwork and sewing geometry need different fields

Specification dimension	Surface artwork record (surface_artwork_id)	Sewing geometry record (panel_geometry_id)
Primary geometric unit	Repeat tile width and height in mm; DPI at 1:1 scale	Closed 2D panel polygons, seam lines and cut lines in mm
Directional alignment	Print orientation (one-way, two-way or tossed) to selvedge	Grainline vector parallel or angled to warp direction
Boundary continuity rule	Tile edge wrap (straight, half-drop or brick pitch)	Seam correspondence, edge-pair lengths and notch coordinates
Size scaling behaviour	Fixed rotary-screen repeat or size-graded placement rule	Parametric or rule-based size grading across POM tables
Physical verification artefact	Conditioned colour swatch, lab dip and print strike-off	Receiving-CAD DXF readback and sewn physical fit toile
Rights and provenance gate	Model log, human edit trace, cultural and mark permissions	Owned base-block lineage and technician modification log

ILLUSTRATIVE [S02] [S22] [S26] [S28] Illustrative GFT specification schema contrasting required attributes for surface print artwork against planar sewing-pattern geometry on the teaching cotton shirt. Informed by S02, S22, S26 and S28; not an audit of a deployed factory database.

Separating the records in **Exhibit 31**, p. 123 protects both pattern and print technicians. [S22] [S28] If a buyer requests a narrower stripe pitch or revised chest emblem on the illustrative cotton shirt while keeping the medium-fit block unchanged, `surface_artwork_id` advances to a new revision while `panel_geometry_id` remains locked. [S09] [S26] The cutting room need not re-verify seam allowances or armhole grading in the receiving CAD system, though marker planning changes if the print introduces a one-way restriction or stripe-matching buffer. [S22] Conversely, if the pattern technician adjusts shoulder slope on `panel_geometry_id`, the rotary repeat tile under `surface_artwork_id` needs no re-engraving, though any engineered print must be checked against the revised seam boundary. [S02] [S28] **Exhibit 31**, p. 123 is an illustrative GFT proposal and does not establish that commercial product-lifecycle software isolates these fields automatically. [S20]

8.4

Appearance control is useful within a defined artwork brief

Generative appearance models speed up early exploration when testing how colour-ways, motif scales or emblem positions alter garment balance. [R30] [S26] However, controllable synthesis only answers layout questions inside a bounded brief; it does not verify that a mark is legally clear, that a motif tiles cleanly, or that a dye house can reproduce screen pixels on cotton cloth. [S21] [S26] [S0801]

8.5 MULTI-CONDITION IMAGE SYNTHESIS SEPARATES GLOBAL COLOUR FROM LOCAL MARKS

Early generative workflows in apparel design during 2023 and 2024 relied mainly on text-to-image tools such as Midjourney and Stable Diffusion to explore style, silhouette, colour and surface pattern. [R30] Those unconstrained diffusion passes often altered garment cut whenever a prompt asked for a new print, or distorted pockets and collars when moving a logo. [S26] Later research decoupled these variables: IMAGGarment enables synthesis with separate control over silhouette, colour palette and local logo placement, while FabricDiffusion extracts distortion-rectified, flat tileable textures from a single clothing image and maps them into the UV space of three-dimensional garments for transfer and relighting. [S26] [S0802]

At concept screening, tools of the class in S26 and S0802 hold silhouette constant while varying surface treatment. [S26] [S0802] On the illustrative cotton shirt, a designer can fix the camp-collar silhouette and preview three options: solid yarn-dyed indigo, a small geometric repeat, and a solid body with a pocket motif, discarding unworkable scale ratios before commissioning production artwork. [S20] [S26]

8.6 THE ARTWORK BRIEF MUST CONSTRAIN MARKS, PALETTES AND REPEAT INTENT UPFRONT

Controlled appearance generation is wasteful if designers generate unconstrained concepts that cannot be printed or cleared for sale. [S20] [S31] Aligning with Chapter 3's constrained brief, GFT recommends attaching five artwork constraints before inference begins:

1. **Permitted marks and emblems**: vector identifiers for owned or licensed brand marks, forbidding synthetic pseudo-logos or unverified third-party insignia. [S21] [S26]
2. **Reference physical palette**: identifiers for physical yarn or dye swatches in the sample room, restricting exploration to colourways that procurement can source or dye on cotton. [S32]
3. **Repeat architecture and maximum pitch**: whether the design is an all-over rotary repeat, flatbed screen repeat or placement print, with maximum repeat dimensions in millimetres. [S0801] [S0803]
4. **Colour-screen ceiling**: the maximum discrete spot colours permitted by the target cost basis, preventing continuous gradients on a four-screen rotary line. [S20]
5. **Placement exclusion zones**: structural areas on the garment block, such as plackets, collar stands, darts or felled seams, where thick pigment paste or embroidery impairs sewing. [S02] [S26]

Setting these boundaries upfront ensures that synthesis serves manufacturing intent, while rights review and physical proofing remain independent gates. [S19] [S21] [S26] [S32]

8.7

Texture coherence differs from manufacturing repeatability

A texture map that looks continuous on a three-dimensional avatar in a viewport can still fail on a cutting table. [S28] [S0801] Computer-graphics texture synthesis optimises for perceptual coherence from camera viewpoints or across UV patches, whereas textile manufacturing requires mathematical periodicity across a flat fabric roll and physical alignment across sewn panel edges. [S22] [S28] [S0801]

8.8 VIEWPORT UV COHERENCE DOES NOT GUARANTEE A JOIN-FREE ROTARY REPEAT

Recent research highlights both the progress and limits of neural texture generation. [S28] [S0801] [S0803] OmniFabric notes that while generative models have advanced three-dimensional garment geometry reconstruction, synthesising high-quality textures remains a bottleneck because multi-view images suffer from occlusions and baked-in lighting. [S28] By synthesising textures within the sewing-pattern UV domain and removing lighting contamination, OmniFabric produces coherent digital assets suitable for rendering, yet a texture painted across UV islands of one size is not a repeatable textile roll. [S28] [S0801]

Printing yardage on woven cotton requires a stationary periodic tile that concatenates along warp and weft without boundary steps, ghost lines or drift. [S0801] [S0803] The authors of TexTile note that existing methods for tileable texture synthesis focus on general texture quality, but lack explicit analysis of intrinsic repeatability properties. [S0801] Research on structured pattern expansion with diffusion models likewise reports that on structured, stationary patterns, diffusion models are generally less reliable and less controllable, requiring noise-rolling and patch-based expansion to preserve regularity and boundary continuity. [S0803]

Consequently, a greenfield maker cannot export a raw diffusion image or UV-mapped texture directly to a print engraver. [S28] [S0803] A print technician must reconstruct the motif into a verified repeat unit whose opposite edges match at the physical resolution required by the engraving or print head. [S0801] [S0803] Detailed print-production mechanics belong to wp20; this paper governs the specification handoff preceding print execution. [S20]

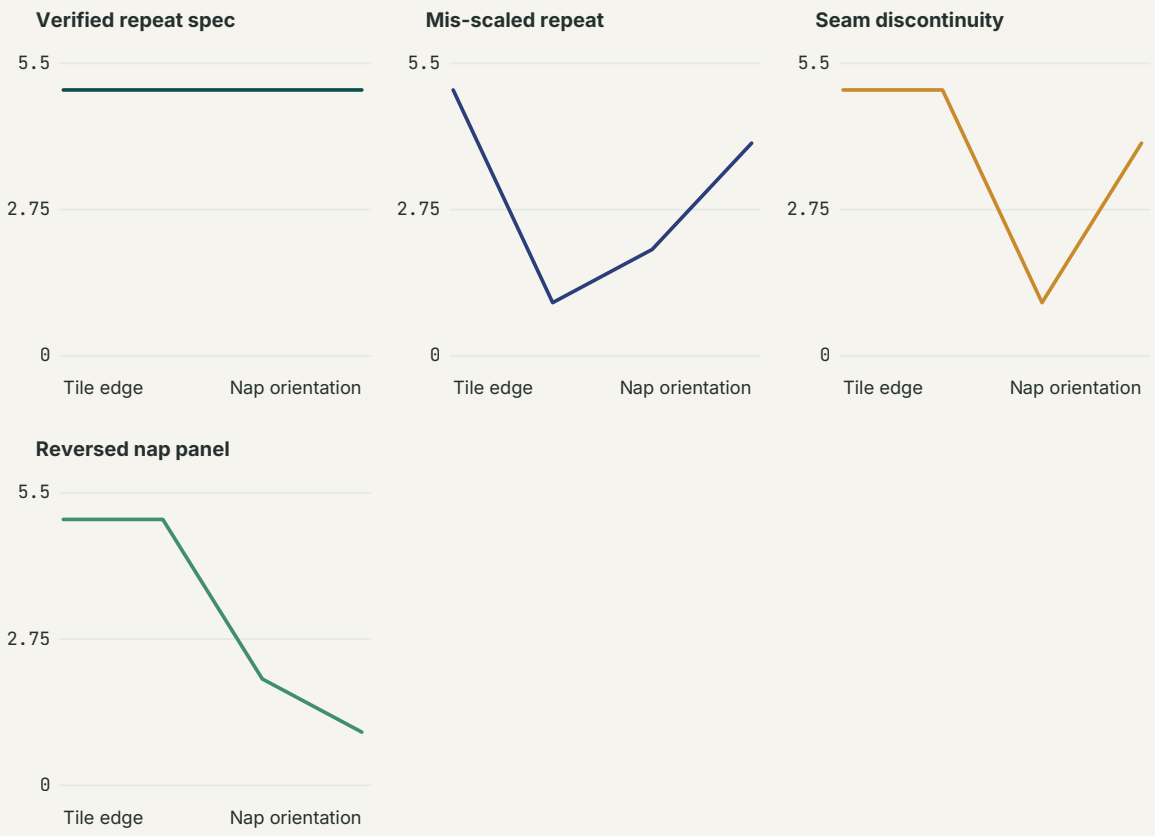
8.9 SEAM JOINS, GRAIN DIRECTION AND SIZE GRADING EXPOSE SURFACE DEFECTS

Cutting a printed fabric roll into shaped panels and sewing them together introduces geometric discontinuities that an image generator never sees. [S02] [S12] [S22] Armholes are concave, sleeve caps are convex, side seams taper, and yokes meet back panels across pleats, while each DXF pattern piece carries a grainline orientation, seam allowance width, notch set and grading rule. [S02] [S12] [S22]

To make surface failure modes explicit during review, **Exhibit 32**, p. 126 compares an illustrative baseline coherent tile against three physical assembly failures across four checkpoints on an illustrative zero-to-five readiness scale. [S12] [S22] [S28] [S0801]

EXHIBIT 32 / ILLUSTRATIVE CHECK SCORE (0 TO 5)

A digital texture can fail at a physical seam



ILLUSTRATIVE [S12] [S22] [S28] [S0801] Illustrative GFT inspection profiles comparing four surface-specification cases across four physical garment checkpoints (5 = verified pass; 1 = critical specification failure). Constructed to illustrate failure modes from S12, S22, S28 and S0801; not measured factory defect rates.

Each panel in **Exhibit 32**, p. 126 isolates a breakdown between digital appearance and physical construction:

- **Mis-scaled repeat**: a motif passes boundary checks in pixel space (`Tile edge = 5`), but was generated at 512 by 512 pixels without a declared millimetre pitch. [S0801] Mapped onto a shirt pattern without unit calibration, a 20-millimetre repeat prints at 120 millimetres (`Real scale = 1`), distorting pocket and collar proportion. [S22] [S28]
- **Seam discontinuity**: the repeat tile and scale are correct (`Tile edge = 5`, `Real scale = 5`), and a single-view render hides the side seam, or a UV generator paints across the mesh without respecting the cutting marker. [S28] When front and back panels are cut with a 10-millimetre seam allowance and sewn together, hori-

zontal bands collide out of phase across side seams (`Seam join = 1`), and size grading shifts panel lengths unless notches and marker buffers lock to the repeat pitch. [S12] [S22]

- **Reversed nap or one-way print direction**: to save fabric in marker nesting, an operator rotates alternate panels by 180 degrees along the warp grainline. [S22] If the motif is directional and the specification omits a `one_way_print` constraint on the grainline, half the sewn panels appear upside down (`Nap orientation = 1`). [S22] [S28]

The scores in **Exhibit 32**, `p. 126` are illustrative diagnostic indicators rather than empirical defect rates from a Cameroonian production line. [S19]

WORKED EXAMPLE: REPEAT-TILE PITCH, SEAM ALLOWANCE AND MARKER YIELD ON THE ILLUSTRATIVE COTTON SHIRT

Consider an illustrative calculation on the teaching short-sleeve cotton shirt. [S09] [S22] Suppose the base size medium shirt uses cotton poplin with a cuttable width of 1,480 millimetres between selvedge pinholes and a 10-millimetre seam allowance on the front, back, yoke, collar and pocket panels. [S22] [S32]

In the illustrative baseline case (a solid fabric or tossed motif requiring no seam matching), a two-way nested marker achieves an illustrative net consumption of 1.35 linear metres per shirt. [S22] Suppose the team selects an AI-assisted horizontal band repeat with a vertical warp pitch (R_y) of 160 millimetres (`0.16` metres), requiring horizontal continuity across the left placket, right placket and chest pocket, plus one-way orientation. [S22] [S26]

Enforcing those two surface rules changes the physical arithmetic at three boundaries:

1. **One-way grainline penalty**: prohibiting 180-degree panel rotation prevents tight interlocking of sleeve and body pieces, adding an illustrative `0.09` linear metres per shirt (`1.35 + 0.09 = 1.44` linear metres before repeat-phase indexing). [S22]
2. **Repeat-phase indexing buffer**: to match the horizontal band across the centre-front placket and patch pocket, each matched panel must sit at an integer multiple of the 160-millimetre vertical pitch ($k \times 160\text{ mm}$), measured at the *stitching line* rather than the raw cut edge. [S12] [S22] If a panel's raw cut height including two 10-millimetre seam allowances (`20` millimetres total) is 740 millimetres, the next matching phase start along the warp is $5 \times 160\text{ mm} = 800\text{ mm}$, leaving an illustrative 60-millimetre buffer (`800 - 740 = 60` millimetres). [S22] Across the full marker, suppose phase indexing adds an illustrative `0.14` linear metres per shirt, bringing total illustrative consumption to `1.44 + 0.14 = 1.58` linear metres per shirt, an increase of `0.23` linear metres (about `17.0` percent above the `1.35`-metre baseline). [S22]
3. **Seam-allowance offset verification**: if a script aligns the repeat tile to the outer *cut line* of the pocket rather than the folded *finished edge* (turned under by an illustrative 15 millimetres at the top hem and 10 millimetres on the sides), the printed band on the sewn pocket sits 15 millimetres out of registration against the shirt front. [S12] [S22]

All linear-metre figures in this case are illustrative scenario inputs constructed to demonstrate the geometric mechanism, not factory measurements. [S19]

8.10

Colour decisions need a controlled physical comparison

A monitor colour is an emitted red-green-blue (RGB) light signal, whereas a dyed or printed cotton shirt is a reflective surface whose perceived colour depends on ambient light spectra, colorant chemistry and yarn structure. [R06] [S28] [S32] Approving production colour from a generated image or uncalibrated screen invites delivery disputes. [S20] [S32]

8.11 SCREEN RGB VALUES ENTANGLE ILLUMINATION WITH SURFACE REFLECTANCE

In generative synthesis, pixel values combine surface reflectance (albedo) with synthetic shading, highlights and colour casts. [S28] [S0802] Both OmniFabric and FabricDiffusion note that clothing images and renders contain lighting contamination and shadows that must be removed before a flat texture map can be relit consistently. [S28] [S0802] Even after algorithmic delighting, the albedo map remains an uncalibrated RGB array with no binding to a dye recipe, colour standard or spectral curve. [S28] [S32]

Even when two monitors display the same sRGB triplet, differences in backlight ageing, white point, glare and profiles alter what designer and buyer see. [S20] Additive screen light also occupies a different gamut from reactive dyes, vat indigo or pigment pastes on cotton fibres: a saturated RGB cyan may be unattainable on carded cotton without speciality pigments, or may shift under store lighting (metamerism). [R06] [S32]

8.12 PHYSICAL COTTON SUBSTRATES AND FINISHING ALTER COLOUR AND FASTNESS

Even with a physical colour standard, interaction between colorant and cloth requires verification on the actual substrate. [S32] As supported by commercial 3D fabric documentation (S32), fabrics vary in fibre composition, yarn count, weight and thickness, governing optical appearance after dyeing or printing: [S32]

- **Substrate whiteness and preparation**: raw Cameroonian cotton lint carries natural waxes and yellow pigments requiring scouring and bleaching before printing light shades; bleaching variations shift colours warmer or duller. [R06]
- **Yarn structure and mercerisation**: mercerised combed cotton poplin reflects light more uniformly and absorbs dye differently from carded plain weave, changing depth of shade. [S32]
- **Finishing and laundering response**: resin finishes, softeners, curing temperatures and washing or sunlight exposure can cause shade change, crocking or staining. [S25]

Neither S26 (IMAGGarment) nor S28 (OmniFabric) models dye exhaustion, binder curing, rubbing fastness or shrinkage. [S26] [S28] This paper supplies no numerical colour-difference tolerance (ΔE) and claims no colourfastness result; those thresholds and tests belong to buyer agreements and wet-processing protocols in wp06 and wp20. [S19]

PHYSICAL COLOUR AND PRINT STRIKE-OFF VERIFICATION BOUNDARY

Before a surface-designed specification advances to bulk release in GFT's proposed workflow, four physical verification records must attach to `surface_artwork_id`: [S20] [S32]

1. **Master physical colour standard**: each colourway channel must reference a physical swatch identifier (`colour_standard_id`) rather than RGB values alone. [S32]
2. **Substrate-matched strike-off or lab dip**: a lab dip on approved `supplier_lot_id` yarn or greige cloth for dyed grounds, or a 1:1 repeat strike-off on the BOM base cloth for prints. [S32]
3. **Controlled viewing comparison**: visual sign-off comparing strike-off against master swatch inside a standardised viewing cabinet under agreed illuminants to screen for metamerism, with conditions logged. [S19] [S20]
4. **Post-wash and finishing check**: a retained strike-off specimen subjected to the specified finishing and care cycle to check repeat-pitch dimensional change and shade stability before bulk cutting. [S20] [S25]

8.13

A cultural reference is not a rights-free style token

Typing the name of a region, community or ceremonial textile tradition into a prompt yields stylized geometric motifs in seconds. [R30] [S23] Treating living cultural heritage as a free aesthetic vocabulary for prompt engineering is a governance failure. [S23] [S24] For a manufacturer in Cameroon or Hong Kong SAR, cultural references require explicit identification, community dialogue and documented permission before any visual derivative is generated, shown to a buyer or printed on cloth. [S23] [S24]

8.14 TRADITIONAL CULTURAL EXPRESSIONS CARRY COMMUNITY IDENTITY AND STEWARDSHIP

The World Intellectual Property Organization (WIPO) frames traditional heritage and fashion clearly. [S23] [S24] In its guidance on Traditional Cultural Expressions (TCEs), WIPO states that TCEs are integral to the cultural and social identities of indigenous and local communities, embody know-how and skills, and transmit core values and beliefs. [S24] In its Indigenous Peoples and Fashion programme, WIPO aims to create opportunities for dialogue between Indigenous Peoples and fashion companies to shape fair and equitable interaction. [S23]

A model trained on internet photographs strips away that social context. [S20] [S24] To a diffusion network, a ceremonial cloth, a royal emblem and commercial wallpaper are indistinguishable pixel distributions. [S20] [S28] Prompting for "Cameroonian traditional geometric print" or uploading a heritage textile photo into a texture-transfer pipeline such as `S0802` synthesises recombinations without knowing which symbols are restricted to specific lineages, title societies or ceremonies, and without consulting custodians. [S23] [S24] [S0802]

8.15 GRASSFIELDS NDOP AND REGIONAL WEAVING TRADITIONS REQUIRE EXPLICIT ENGAGEMENT

Across the Cameroon Grassfields and northern weaving centres, textiles such as *ndop* (indigo resist-dyed cotton cloth stitched with raffia before vat dyeing) and hand-woven prestige cloths are not arbitrary stripes. [R06] [S23] [S24] Specific geometric arrangements signify royal court architecture, judicial authority, alliances and investiture ceremonies, and genuine *ndop* production spans specialized drawing, raffia stitching and indigo vat stewardship across communities. [S24]

In GFT's visual identity system (**Warp and Signal**), indigo conceptually references Grassfields *ndop* dyeing to mark the data layer in diagrams, while ochre references northern cotton-belt soils. That palette is an editorial convention for charts; GFT claims no ownership of, licence in, or right to commercialise any traditional Cameroonian motif, nor does this paper generate or publish synthetic *ndop* prints. [S23] [S24]

If a maker or brand partner wishes to develop a collection inspired by Grassfields *ndop*, northern woven geometry or other community heritage, GFT recommends treating cultural permission as a mandatory upstream gate preceding generative ideation and sampling (**Exhibit 33**, p. 130). [S23] [S24]

EXHIBIT 33 / GOVERNANCE SEQUENCE

Cultural permission precedes commercial reuse



1. TCE present
2. Holders found
3. Terms agreed
4. Draft ready
5. Co-approved
6. Revise or halt

ILLUSTRATIVE [S20] [S23] [S24] Illustrative GFT governance sequence for any design brief referencing a traditional cultural expression (TCE). Guided by WIPO engagement principles (S23, S24) and NIST generative-AI provenance controls (S20); not a claim of a universal statutory procedure or an existing GFT community partnership.

As shown in **Exhibit 33**, p. 130, the workflow begins at brief screening (1. Screen brief for TCEs) before uploading reference photos or prompting for heritage motifs. [S20] [S23] If a traditional cultural expression is involved, the sponsor must identify the relevant artisans, traditional authorities, cooperative or institution (2. Identify community holders), discuss whether the motif is appropriate for apparel use (3. Discuss scope and meaning), and document the agreed scope, attribution, craft or benefit participation, and duration (4. Record consent and terms). [S23] [S24] After initial consent, the developed specification and placement return for joint review (5. Co-review artwork in context), because placing a respected motif on an inappropriate garment zone or distorting its proportions during grading can violate the agreement. [S23] [S24] If custodians cannot be identified, decline permission, or object during co-review, the motif is halted. [S23]

While **Exhibit 33**, p. 130 translates the institutional principles in S23 and S24 into an operational gate, it is an illustrative GFT proposal. [S23] [S24] It does not constitute legal advice on folklore or intellectual-property statutes in Cameroon, OAPI states or Hong Kong SAR (wp17), nor imply that GFT has entered into any partnership or obtained permission for any Cameroonian motif. [S23]

> **AUTOMATED SIMILARITY SCREENING CANNOT SUBSTITUTE FOR COMMUNITY CONSENT**
>
> Some vendors and procurement teams propose running reverse-image searches or cosine-similarity checks on generated prints to clear rights before production. [S20] While similarity screening can flag a direct copy of a commercial photo, it cannot establish cultural permission. [S23] [S24] A generative model may recombine elements of a traditional cultural expression so pixel similarity to any single archive photo is low while the cultural signifier remains unmistakable and unauthorized. [S24] Authority over traditional cultural expressions rests with the living community and its custodians, not an image-search database. [S23] [S24]

8.16

AI asset provenance and cotton provenance need different records

Because sustainable apparel programmes emphasize fibre-to-garment traceability, teams sometimes conflate digital design provenance with physical cotton chain of custody. [S20] [S0804] Digital asset provenance and physical material provenance track different entities, use different verification mechanisms, and fail differently; a maker must keep the two ledgers distinct while linking them at the bill of materials. [S20] [S21] [S23] [S0804]

8.17 CRYPTOGRAPHIC CONTENT CREDENTIALS VERIFY EDIT HISTORY RATHER THAN LEGAL OWNERSHIP

Initiatives such as the Coalition for Content Provenance and Authenticity (C2PA) specify technical standards (S0804) for binding origin and edit metadata to media files through cryptographic signatures (Content Credentials). [S0804] As stated in C2PA Technical Specification 2.1, each time an asset is changed, existing provenance is

preserved and each new change is added. [S0804] Supported generators or editors can embed a signed manifest recording the software agent, actions taken, ingredient hashes and timestamps. [S0804]

A C2PA manifest or internal generation log (prompt text, model identifier, version, seed, reference hashes and human revisions, as in Chapter 3 and NIST S20) is a useful audit artefact, but validates only that manifest assertions were signed by the credential holder and untampered since signing; it does not judge whether content is lawful, original or non-infringing. [S20] [S0804]

Furthermore, logging a prompt and signature does not automatically create copyright in the generated graphic. [S21] As the U.S. Copyright Office explains in Part 2 of its Artificial Intelligence Report (S21), generative AI outputs are protected by U.S. copyright only where a human author determined sufficient expressive elements, distinguishing human creative arrangements and substantive edits from prompts alone. [S21] For surface_artwork_id, GFT recommends logging three digital provenance layers:

1. **Technical generation and edit lineage**: model identifier, weights or API version, input reference IDs, prompt parameters, cryptographic hash and C2PA manifest. [S20] [S0804]
2. **Human expressive contribution record**: the designer or print technician's sketches, vector reconstructions, repeat-tile edits and colour separations documenting human contribution under applicable law. [S21]
3. **Affirmative rights and permission links**: licence records for input assets, vendor commercial-use policy snapshot (S31), and community consent record (S23, S24) where cultural references appear. [S23] [S31]

8.18 PHYSICAL COTTON CHAIN OF CUSTODY TRACKS MASS BALANCE AND LOT IDENTITY

Digital records say nothing about where the shirt's cotton was grown, ginned, spun, woven or dyed. [R06] [S32] Physical provenance tracks matter across industrial transformations: seed-cotton collection in northern Cameroon, ginning bale numbers, spinning lots, greige rolls, dye batches and marker lay logs (wp05). [R03] [R06] [S32]

Within wp10, the rule is non-substitution between the two provenance domains. [S20] Attaching a C2PA credential or AI log to a digital product passport must never be presented as evidence of sustainable or locally processed cotton (R06, R07). [S25] [S0804] Conversely, printing on 100 percent traceable Cameroonian cotton (R06) does not cure an unlicensed trademark or unconsented cultural motif. [S21] [S23] In the tech pack bill of materials, surface_artwork_id points to the digital rights envelope, while supplier_lot_id and swatch_id point to the physical cotton ledger defined in wp05 and Chapter 6. [S32]

8.19

Publication and production require separate permissions

In a fast-moving studio, a surface graphic generated for an internal mood board can slip into a buyer pitch deck, social-media teaser and sample print order within days. [S20] Because each destination carries different contractual, legal and technical prerequisites, a maker should enforce destination-specific approval states so an asset approved only for internal discussion cannot be exported to a website or printer. [S20] [S21] [S31]

8.20 VENDOR COMMERCIAL-USE TERMS ARE NARROWER THAN COMPLETE RELEASE CLEARANCE

Commercial software terms show why permission depends on tool configuration and destination. [S31] In the Adobe Firefly FAQ (S31), the vendor states that for features without the beta label, users can use Firefly-generated outputs in commercial projects. [S31] Consequently, an output generated using a feature labelled *beta* (or a research checkpoint from S26, S28, S0802 or S0803 under a non-commercial academic licence) cannot be used in a commercial garment or campaign without separate rights. [S20] [S31]

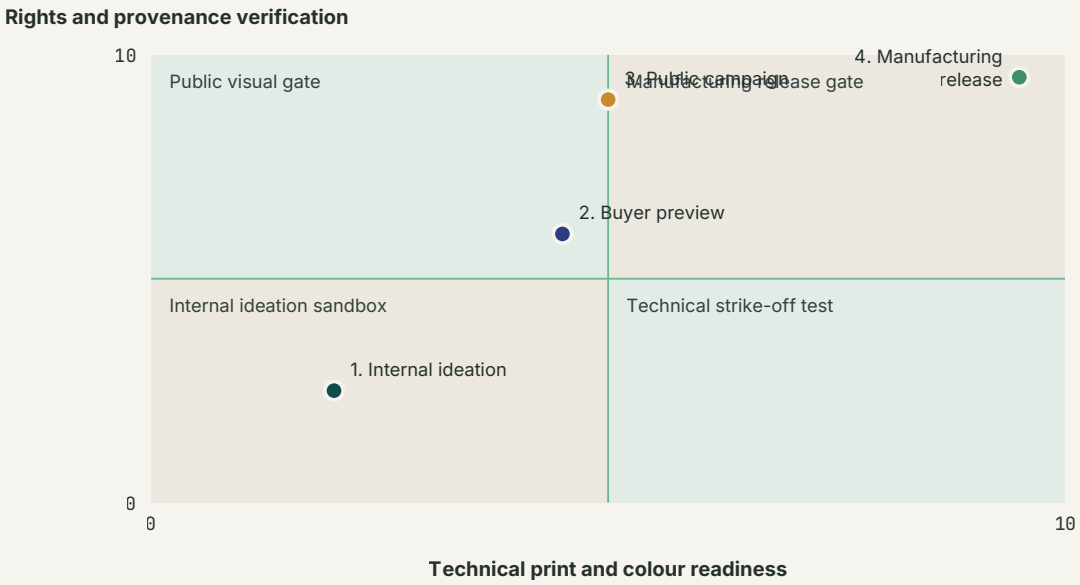
Even in a non-beta tier, vendor terms only govern the relationship between software provider and subscriber. [S31] They do not clear third-party trademarks reconstructed through image conditioning (S26), grant community permission for traditional cultural expressions (S23, S24), establish copyright exclusivity (S21), or verify that the graphic prints cleanly on cotton (S32). [S21] [S23] [S31] [S32]

8.21 FOUR DESTINATION APPROVAL STATES GOVERN WHERE AN ARTWORK ASSET MAY TRAVEL

GFT proposes partitioning the lifecycle of `surface_artwork_id` into four permission states, mapped in **Exhibit 34**, p. 134 across technical readiness and rights/provenance completeness. [S20] [S21] [S31] [S0804]

EXHIBIT 34 / ILLUSTRATIVE GATE INDEX (0 TO 10)

Asset approvals must state the permitted destination



ILLUSTRATIVE [S20] [S21] [S31] [S0804] Illustrative GFT governance matrix placing four surface-asset destinations along required technical print-geometry readiness (x-axis) and required rights/provenance verification depth (y-axis). Informed by S20, S21, S31 and S0804; positions are conceptual governance tiers, not empirical measurements.

The four destination states plotted in **Exhibit 34**, p. 134 enforce specific controls before an asset moves between systems:

1. **State 1:** `internal_ideation` **(Sandbox)**: permits designers to explore colourways and placement concepts within an isolated workspace (`x = 2`, `y = 2.5`). [S20] Input references must obey the reference manifest from Chapter 3, culturally sensitive motifs without active dialogue (S23) are excluded, and outputs cannot be sent to buyers, published online or routed to a print server. [S20] [S23]
2. **State 2:** `buyer_preview` **(Confidential pitch)**: allows a concept to be shown to a buyer in a closed review (`x = 4.5`, `y = 6`). [S20] Requires verification that the model licence permits commercial pitch use (S31), brand marks belong to the buyer or maker (S21), cultural engagement terms permit preview (S23), and the preview carries a watermark noting that repeat scale and colour await physical strike-off. [S20] [S32]
3. **State 3:** `public_campaign` **(External publication)**: governs lookbooks, website imagery or catalogues (`x = 5`, `y = 9`). [S20] Public display requires verified non-beta commercial tool status (S31), documented human expressive contribution where exclusivity is claimed (S21), signed community approval and attribution for any cultural expression (S23, S24), and a truthful caption distinguishing synthetic renders from photographs of sewn Cameroonian cloth. [S20]
4. **State 4:** `manufacturing_release` **(Production artwork)**: the sole state authorised to trigger screen engraving, bulk digital printing or marker cutting (`x = 9.5`, `y = 9.5`). [S20] [S22] Requires State 3 clearance plus verified repeat-tile periodicity

(S0801, S0803), millimetre dimensions and seam-join rules mapped to the approved DXF pattern (S22), and a signed physical strike-off on the actual cotton lot (S32). [S22] [S32]

The coordinates in **Exhibit 34**, p. 134 are illustrative ordinal positions designed to separate gate responsibilities, not quantitative risk probabilities. [S19]

8.22 CRYPTOGRAPHIC HASHES BIND THE RELEASED TECH PACK TO THE AUTHORISED ARTWORK

One vulnerability remains after a surface design passes all four gates: file drift between studio, tech pack and print floor. [S20] When a tech pack PDF holds only a thumbnail and a mutable filename such as `shirt_print_final_v2.tif`, a later edit that overwrites that filename can cause a reorder to print the wrong file without detection. [S15] [S20]

To close that loop before Chapter 9's readback gate, GFT recommends storing an immutable cryptographic digest (`sha256_hash`) of the authorised production artwork file, alongside its C2PA manifest identifier (S0804), repeat dimensions in millimetres, `colour_standard_id` list, `permission_record_id` and `approver_role` inside the versioned tech-pack record. [S20] [S0804] When the print operator or partner mill pulls the high-resolution file, the receiving system verifies that the file hash matches the tech-pack hash before ink touches cotton. [S20]

> **GFT NOTE**
>
> In GFT's Cameroon textile programme, cut, make, trim and finishing form the foundational operating scope, while upstream printing and wet processing are staged or partnered capabilities qualified against strict quality records. [R06] Configure the design workspace so ideation outputs default to `internal_ideation` only, keep `surface_artwork_id` separate from `panel_geometry_id`, require a physical strike-off on the actual cotton lot before colour sign-off, and block any culturally referenced motif unless documented community permission is attached to the specification record. [S23] [S28] [S32]

A surface motif becomes a specification only when its repeat fits the seam, its colour survives the cloth and its provenance honours the makers behind both the art and the cotton.

CHAPTER SOURCE NOTES

S28 Authors via arXiv; SIGGRAPH Asia 2026 acceptance listed. 2026-09-24.

S23 World Intellectual Property Organization. 2026.

S32 CLO Virtual Fashion, CONNECT Support. 2025-04-17.

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

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

S21 U.S. Copyright Office. 2025-01-29.

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S22 CLO Virtual Fashion. 2026-07-15.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S14 Authors via arXiv. 2026-06-23.

S04 Authors via arXiv. 2024-01-29.

S26 Authors via arXiv. 2025-04-17.

S20 National Institute of Standards and Technology. 2024-07-26.

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

S0801 Authors via arXiv. 2024-03-19.

S0802 Authors via arXiv. 2024-10-02.

S31 Adobe. 2026-05-13.

S0803 Authors via arXiv. 2024-11-12.

S19 National Institute of Standards and Technology. 2023-01-26.

S12 Authors via arXiv; WACV 2026. 2026-02-25.

S25 European Commission, Directorate-General for Environment. 2022-03-30.

S24 World Intellectual Property Organization. 2026.

S0804 Coalition for Content Provenance and Authenticity. 2024-09.

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

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

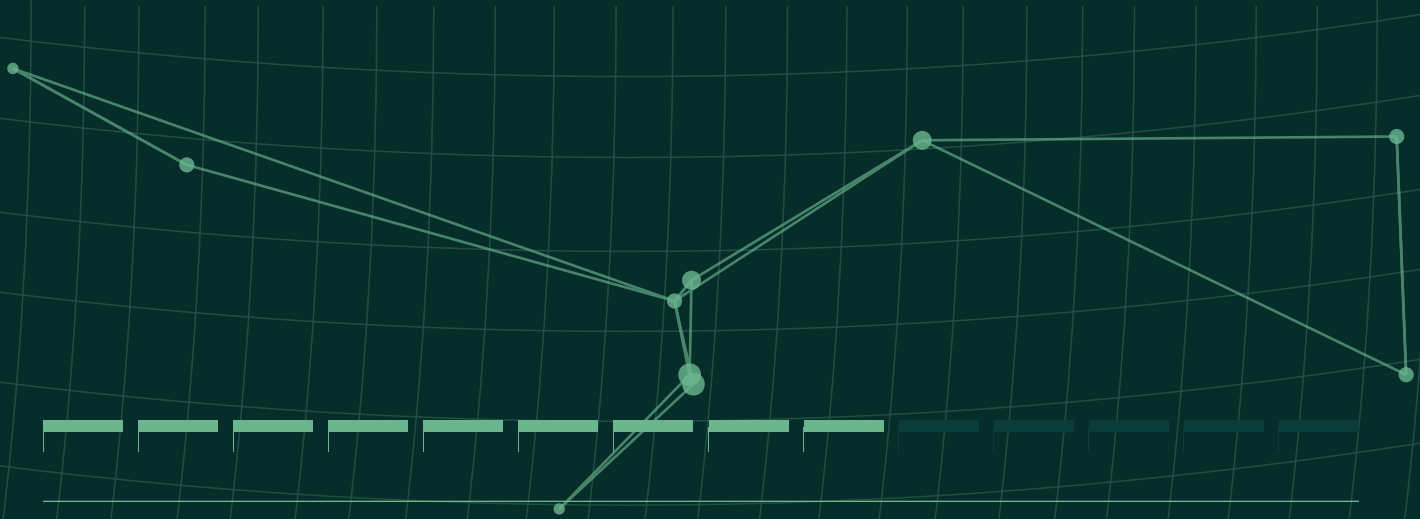
GREEN FUTURE TECH

LIMITED / HONG KONG SAR

09

A tech pack is complete only when the receiving team can use it

Generated specification fields must resolve to approved geometry, materials, instructions and revisions in the actual receiving workflow.



CHAPTER 09
DECISION TEST

Generated specification fields must resolve to approved geometry, materials, instructions and revisions in the actual receiving workflow.

A tech pack is complete only when the receiving team can retrieve the intended revision, resolve every applicable input to an approved record, reopen the pattern in its own CAD workflow and identify who may authorise the next action. A polished document or syntactically valid JSON file does not meet that test by itself. ChatGarment demonstrates that a vision-language model can emit garment descriptions and continuous parameters for a parametric sewing-pattern program, while CLO documentation shows that pattern exchange still exposes choices about format, scale, grading, notches and line interpretation. [S10] [S22] Green Future Tech Limited (GFT), Hong Kong SAR, therefore recommends a controlled information contract with explicit provenance, receiving-CAD readback, material resolution and release authority. Generated text may draft entries, but missing evidence must remain visibly unresolved. This chapter defines that contract for the paper's illustrative cotton shirt. It does not claim a completed GFT factory integration, a tested commercial CAD round trip or a production result.

9.1

The tech pack should be a controlled information contract

A useful tech pack is not a container into which every available design fact is poured. It is a controlled interface between design intent and a defined receiving activity. The product developer needs style identity and intended use. The pattern technician needs geometry, size rules and construction relationships. The sample room needs material identifiers and assembly instructions. Quality and buying teams need acceptance criteria and the current approval state. These uses overlap, but they are not interchangeable. The contract should therefore state which fields are required, optional or not applicable for each product family, and which downstream decision each field supports. [S19] [S22]

Required status must follow consequence rather than document convention. Fibre content may be mandatory for a material record, while artwork repeat dimensions are not applicable to an unprinted shirt. A construction detail may be optional during early ideation but mandatory before sampling. The applicability rule, its effective revision and the person who set it should be stored with the field. Otherwise a completeness report can reward a team for omitting difficult information or penalise it for leaving a genuinely irrelevant field blank. NIST's framework leaves risk tolerance to the organisation and its context, so this chapter does not invent a universal tech-pack threshold. [S19]

Every automatically populated field needs an origin and a review state. `Source` should identify the approved block, measured material record, buyer instruction, technician entry or model proposal from which the value came. `Owner` should identify the role accountable for resolving that value, not merely the person who typed it. `Validator` should state the check that can fail. `Approver` should be present only where the field requires a decision rather than a syntax check. A generated value can remain a proposal while an approved source value carries a reviewer and date. This distinction prevents fluency from becoming false authority. [S20]

Exhibit 35, p. 139 applies these rules to the illustrative shirt without supplying invented product measurements or supplier details. Its important feature is not the number of rows. It is the path from each field to evidence, validation and accountability.

EXHIBIT 35 / ILLUSTRATIVE FIELD CONTRACT

A specification field must point to its source and owner

Field group	Applicability	Originating record and validator	Decision owner
Style and revision	Always required	Controlled style register; unique identity and supersession check	Product developer
Technical flats	Required for sampling	Approved drawing set; front, back and detail coverage review	Designer
Measurements and tolerances	Required for fit and release	Point-of-measure table; landmarks, units, sizes and agreed limits	Pattern technician and buyer role
Pattern and grading links	Required for cutting	Approved pattern revision; receiving-CAD geometry readback	Pattern technician
Bill of materials	Required for sampling and buying	Material and trim register; identifiers, alternatives and quantity basis resolve	Procurement and technical roles
Construction instructions	Required for assembly	Operation record; seam, stitch and sequence references match panels	Sample technician
Artwork and permissions	Conditional on surface design	Controlled artwork revision; scale, placement and permitted-use checks	Designer and rights reviewer
Approval and release state	Always required	Signed state transition; no unresolved critical prerequisite	Named release authority

ILLUSTRATIVE [S19] [S20] [S22] [S32] Illustrative GFT schema for the teaching cotton shirt. Required status varies by product family and release state. No live buyer, supplier, measurement or approval data are asserted.

The table also separates validation from approval. Software can establish that a mandatory cell is populated, that an identifier exists or that two revisions differ. It cannot decide that a tolerance is commercially acceptable, that a substitute fabric preserves the intended fit or that an exception is safe to release. Those decisions remain with the designated roles. A validator should reject contradictory entries, such as a pattern revision that does not match the linked measurement table, but a successful schema check is only the start of technical review. [S19] [S22]

Revision belongs to the whole contract, not just its cover. Each linked artefact needs its own identity and revision, while the tech pack records the exact combination released together. This allows a collar drawing to change without silently replacing the

body block, and a fabric record to advance without erasing the material used for an earlier sample. The receiving team must be able to determine whether its file is current, superseded or under review before work begins.

> **GFT NOTE**
>
> Implement `awaiting evidence` as a controlled state with an owner and a return path, not as an empty cell or a plausible generated substitute. Before procurement, run one representative style through the intended receiving workflow, including CAD reopen, material resolution and revision confirmation. Treat any unresolved required field or failed critical readback as a release hold. This is GFT's proposed control, not a claim that a commercial system supplies it automatically.

9.2

Structured text can assist drafting but cannot invent missing inputs

Structured generation is useful because it can turn an approved brief and technician's notes into a first draft with consistent field names. ChatGarment's research interface generates JSON containing textual garment attributes and continuous numerical parameters, then uses that file to drive a parametric sewing-pattern model. [S10] That establishes a route from multimodal input to structured pattern parameters. It does not establish a graded bill of materials, verified supplier identifiers, buyer tolerances, operation sequences or production approval. The output boundary of the research system must not be relabelled as a complete factory tech pack.

The safe drafting pattern is retrieval before generation. The system should first receive the controlled style record, approved block, applicable point-of-measure schema, material and trim identifiers, construction library and current revision. It may then propose text for fields whose source records exist. A construction sentence should cite the operation or seam identifier it paraphrases. A measurement entry should preserve the source unit and landmark rather than infer a number from a garment description. A material description should resolve to the selected material record, not to a generic phrase such as cotton fabric. NIST's generative-AI profile identifies confabulation and information-integrity risks, which makes source-linked verification especially important when generated content enters consequential work. [S20]

Unknown inputs need a different behaviour from missing prose. If fibre content is absent, the correct output is `awaiting evidence: material record`, with procurement or the material owner assigned to resolve it. If a tolerance has not been agreed, the system should not select a familiar value from another product. If a trim code has no approved supplier reference, a plausible alphanumeric code is more dangerous than a blank because it can be mistaken for a buying instruction. These are deliberate fail-closed rules. They preserve the distinction between an absent value, an inapplicable field and a rejected value.

Extraction requires the same discipline as composition. A model asked to read an incoming document should return the exact source location, the extracted text or value, a confidence or review flag where available, and any conflict with the controlled

record. The reviewer should see the source passage beside the proposed field. Suggested actions in the NIST profile include empirical evaluation and verification of generated information, but they do not certify a particular extraction workflow. [S20] The local acceptance test must therefore include wrong units, conflicting revisions, ambiguous abbreviations and absent evidence, not only clean documents.

A practical field can carry four separate states: `proposed`, `source verified`, `technically approved` and `superseded`. The states should not collapse into a single coloured tick. Source verification establishes that the value matches its origin. Technical approval establishes that the accountable role accepts it for the intended use. Supersession records that it once applied but no longer does. This structure lets a model accelerate clerical drafting while leaving technical judgement and release authority visible. It also creates a usable audit trail when a later sample fails and the team needs to know which information was actually approved.

9.3

CAD exchange needs an explicit format and unit agreement

A filename ending in DXF does not specify enough information for a controlled handoff. CLO's documentation lists Standard DXF, AAMA and ASTM choices, together with controls for scale, grading, notches, annotations and the relationship between cutting and sewing lines. [S22] These options are useful because they make interpretation visible. They also show why the sender and receiver must agree settings before exchanging production work. Format support is a documented capability, not evidence that any two installations will preserve every required meaning.

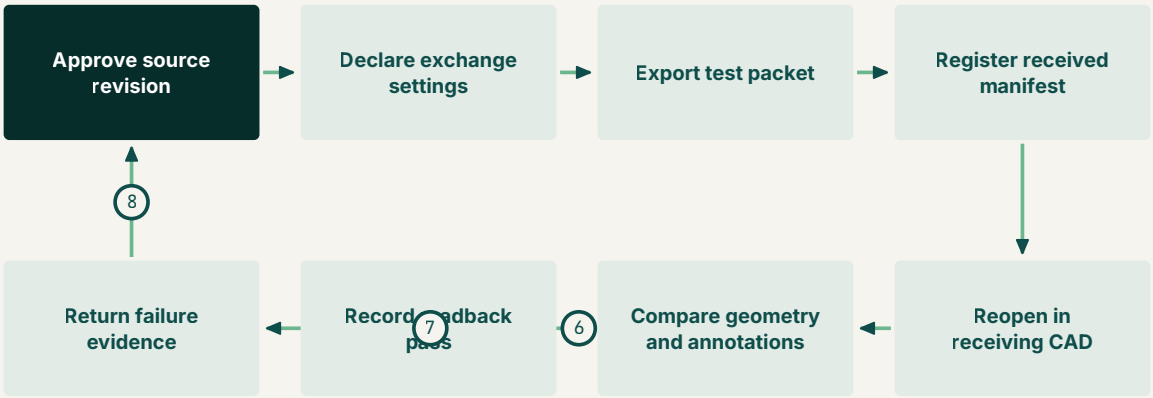
The exchange agreement should name the sending and receiving systems, their releases, the selected DXF convention, the unit, the size group and the grading method. It should state whether grading travels inside the DXF or in a separate rule file, whether annotations are expected, how notches are placed, and which line represents the cut boundary. CLO notes that cutting and sewing lines may change depending on the exporting CAD settings, and its interface provides a swap option. [S22] That makes line-role verification a required readback check rather than a visual assumption.

The agreement should also identify a representative test packet. For the illustrative shirt, the packet would include a base size and the agreed size extremes, panels with straight and curved boundaries, a sleeve and cuff with intentional easing, a collar assembly, notches, seam allowances and internal lines. The purpose is to expose interpretation problems on geometry that the team actually intends to use. A simple rectangular panel can pass while curve points, grading or annotations fail elsewhere. The packet is a GFT test design, not a universal interoperability standard.

Exhibit 36, p. 142 makes the exchange loop explicit. The loop ends only when the receiving technician has compared evidence against the approved source revision. A failed import returns to the responsible sender with settings, screenshots or measurement records attached. It does not become an unlogged manual repair at the receiving desk.

EXHIBIT 36 / CONTROLLED HANDOFF

CAD readback closes the exchange loop



6. All critical checks pass
7. Any critical check fails
8. Correct and supersede

ILLUSTRATIVE [S19] [S22] Illustrative GFT workflow informed by documented DXF controls in S22. It does not assert that a particular vendor pair has passed this loop.

RECEIVING-CAD TEST RECORD

Retain the source revision, exported file, exchange manifest, receiving-system release, import settings and comparison record. Compare panel identity and count, units, selected points of measure, curve shape, notches, cut and sewing line roles, seam allowances, grading and cut quantities. Record each check as pass, fail or approved exception, with evidence and the technician responsible. Repeat the test after a relevant software, format or pattern-library change. The record demonstrates only the tested configuration and specimen set.

Procurement should require this evidence before dependence on a workflow, not after a production deadline has made correction expensive. A vendor demonstration can establish that an export screen exists. Only the receiving test can establish that the intended team can use the resulting artefact. Even a successful test remains bounded to the versions, options and garment features examined. [S19] [S22]

9.4

Round-trip verification must inspect dimensions and annotations

Opening a file proves that the receiving parser accepted something. It does not prove that the resulting pattern is the approved pattern. Readback must compare meaning at the sewing and cutting boundary. The technician should identify every expected panel, check cut quantities and orientation, measure agreed landmarks, inspect

curves and confirm that notches and internal lines remain associated with the correct pieces. Graded sizes need their own comparison because a correct base size can conceal a missing or altered rule. [S22]

The comparison should distinguish three outcomes. `Pass` means the required meaning survived without intervention. `Corrected import` means the receiver changed a setting or repaired an artefact, which must be logged and fed back to the sender. `Fail` means the packet cannot support the intended next action. Combining corrected and untouched files into one pass rate would hide recurring exchange work. Where a technician approves an exception, the reason, affected sizes and downstream retest should be attached to the exact revision. NIST's emphasis on context-specific evaluation supports defining these outcomes before judging the workflow. [S19]

Annotations deserve the same scrutiny as coordinates. CLO provides an option to import pattern annotations and separate controls for notch placement and duplication. [S22] A panel can therefore retain its outline while losing the label that connects it to the bill of materials or construction sequence. Similarly, a notch shifted from the sewing line to an unintended outline can change assembly interpretation without changing the panel bounding box. Readback needs semantic checks, not only geometric tolerance.

Exhibit 37, p. 143 separates what the opened vendor documentation states from what remains for the local team to prove. The right-hand column is deliberately a test requirement, not a compatibility claim.

EXHIBIT 37 / DOCUMENTED CONTROL AND REQUIRED LOCAL EVIDENCE

Format support is narrower than tested interoperability

Exchange dimension	Documented control	Local evidence still required
File convention	Standard DXF, AAMA and ASTM import/export choices	Reopen the selected convention in the intended receiving release
Scale	Unit or scale setting on import and export	Compare known dimensions in the declared unit
Size group and grading	Size-group selection, grading in DXF or rule file, current-size option	Verify the complete size set and grading at agreed points
Notches	Placement, movement and duplication controls	Count and locate notches against seam intent
Cut and sewing lines	Swap option and warning about differing CAD interpretations	Confirm line roles and seam-allowance boundaries
Annotations	Pattern-annotation import option	Read labels against panel and construction records
Curves and points	Curve-point optimisation and straight-point conversion controls	Inspect shape and dimensions after import

SOURCED [S22] Sourced from CLO DXF import/export documentation for the documented-control column. The local-evidence column is GFT's proposed readback protocol. No receiving-system result or compatibility score is asserted.

The table does not rank DXF conventions. AAMA may suit one receiving workflow, ASTM another and Standard DXF a different purpose. The correct choice is the one agreed and demonstrated for the actual information required. Nor does a passed shirt packet establish readiness for every garment. New symmetry, layering, padding, special notations or grading structures can introduce different failure modes and should extend the test set deliberately.

WORKED EXAMPLE: COMPLETENESS AND READBACK ARE GATES, NOT AVERAGES

Consider an illustrative teaching pack with 24 applicable required fields. Suppose 21 resolve to approved records while fibre content, a supplier trim code and one tolerance remain awaiting evidence. Field completeness is `21 / 24 x 100 = 87.5 percent`. The pack nevertheless remains blocked because all three unresolved fields are required for the intended release. The arithmetic describes an illustrative completeness measure, not an acceptable release score or a factory observation.

Now suppose an illustrative readback checklist contains 26 applicable checks: 12 panel identity and quantity checks, eight dimensional comparisons and six notch checks. All panel checks pass. One deliberately corrupted unit trial reads a source dimension of 560 millimetres as 5,600 millimetres, a tenfold error, and one expected notch is absent. The result is `24 / 26 x 100 = 92.3 percent` of checks passed, but the readback still fails. After the unit declaration and annotation settings are corrected, all 26 checks pass in this constructed example. These values demonstrate gate arithmetic only. They are not a measured CLO compatibility rate, and readback success would still leave material, fit and release approvals to be completed.

9.5

BOM links should resolve to approved physical inputs

A bill of materials is a set of references to physical inputs, not a list of attractive nouns. Each line should resolve to a controlled material or trim identifier, its revision, the approved supplier reference where applicable, colour or finish, placement, size dependence and quantity basis. The identifier connects the tech pack to the material evidence envelope. It should not be replaced by a generated description, a thumbnail or a generic fibre family. CLO's fabric-asset guidance asks for weight, thickness, fibre content and measurement information, which illustrates the difference between a visual asset and a documented material. Those submission instructions are not a physical test standard or proof of simulation accuracy. [S32]

The receiving team needs to know what may actually be bought and used. A line stating `cotton poplin` leaves the procurement boundary unresolved. A controlled line points to the selected material record and shows whether its composition, construction, width, finish and approved use have been verified. For a button, zip, label or thread, the same logic applies to the trim record and its intended position. If an item is still being sourced, the BOM should carry `awaiting evidence` and block the relevant purchase or release state. It should not manufacture a supplier code to make the document look finished.

Approved alternatives need their own identities and conditions. An alternative is not equivalent merely because a model describes it with similar words. The record should state which characteristics may differ, which reviews are already satisfied and which changes reopen fit, costing, artwork or construction approval. A fabric substitution can alter drape, shrinkage, marker consumption and sewing behaviour, so its approval cannot be inferred from the original material's presence in the BOM. The material-measurement chapter defines the physical evidence; this chapter ensures that the released reference points to it.

Quantity basis deserves a separate field from quantity. A consumption figure may be per garment, per size, per marker, per roll allocation or per operation, with or without an explicit loss assumption. Until marker work and sampling have tested the basis, consumption should remain provisional. The receiving team must be able to tell whether it is looking at an estimate for planning or an accepted production quantity. Generated images establish neither trim availability nor landed cost, and no numerical consumption or price is supplied for this chapter's teaching shirt.

BOM validation should therefore test referential integrity before arithmetic. Every required identifier must exist, every referenced revision must be current for the pack, and every approved alternative must carry its conditions. Only then should software calculate extended quantities or roll requirements. A correct total built from unresolved inputs is still unresolved. Conversely, a missing quantity should not conceal an otherwise valid material identity; the exception should name the missing basis and its owner so that procurement and technical review can act on different tasks.

This structure also supports sample reconciliation. The sample room should record which material and trim revisions were actually used, including any approved deviation from the pack. If the physical sample differs from the BOM, the team needs a deliberate change record rather than an informal note on a printout. The accepted sample and the released BOM must eventually describe the same approved combination, or the difference must remain an explicit open condition.

9.6

Change control must travel with the specification

A tech pack is useful only if the receiver can recognise which combination of decisions is current. TailorCoPilot's TailorTrace research models sewing patterns as discrete states and records transformations as explicit operations on panels, edges, vertices and stitches. [S15] That state-history approach offers a useful design cue for specification control. It does not replace product lifecycle management, approve a buyer change or determine which production reviews must reopen. The tech pack needs its own change contract above the geometry log.

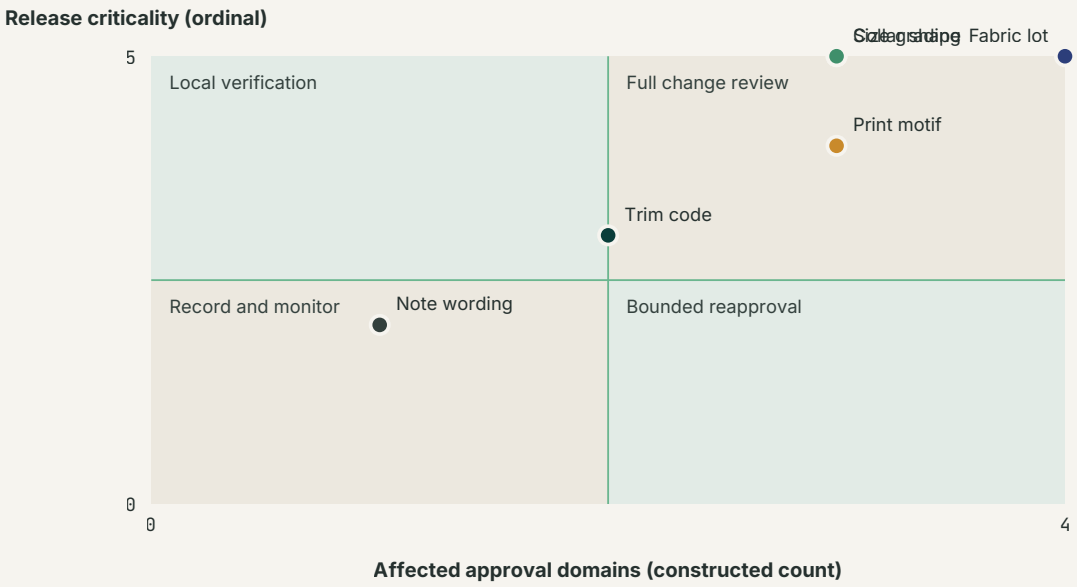
Every change request should state the reason, the originating decision, the previous and proposed revisions, affected sizes and affected material or artwork dependencies. The system should then identify the approvals whose evidence no longer covers the proposed state. A collar-geometry change may preserve the fabric record but reopen pattern readback, construction and fit. A fabric-lot change may leave

coordinates untouched while reopening material, fit, marker and costing review. The point is not to maximise bureaucracy. It is to prevent an old approval from being applied to a new physical or geometric condition.

Exhibit 38, p. 146 maps illustrative changes by the breadth of affected review and the criticality of the release decision. The positions are a planning device, not measured probabilities. Their purpose is to make reapproval explicit before the changed pack reaches the receiving team.

EXHIBIT 38 / ILLUSTRATIVE REVIEW BREADTH AND RELEASE CRITICALITY

A change can reopen several approvals



ILLUSTRATIVE [S15] [S19] [S20] [S22] [S32] Illustrative GFT change-review model. Horizontal position is the constructed count of affected approval domains from 0 to 4. Vertical position is an ordinal release-criticality judgement from 0 to 5. Neither axis is an observed frequency, probability or vendor score.

A real change board should replace these constructed positions with the actual dependency record. A wording clarification may be administrative only if it does not change meaning. If it alters a stitch instruction, it becomes a construction change despite involving no coordinates. Likewise, a new trim code may be a simple identifier correction or a different physical component. Classification should follow the consequence, supported by a reviewer where the distinction is uncertain. NIST's lifecycle guidance supports monitoring changed conditions rather than assuming that an earlier evaluation remains sufficient. [S19] [S20]

Supersession must preserve the last approved state. The previous pack, its linked pattern and material records, its readback evidence and its approval history should remain retrievable. The new revision should not overwrite them. This permits a controlled return when a change fails and allows the team to identify which samples or work orders were made under which state. A revision identifier alone is insufficient if the linked artefacts can drift independently; the release manifest must bind the exact set.

Urgent changes need the same discipline, not an informal exemption. A receiver facing an immediate production problem can hold the affected work, record the proposed deviation and obtain the defined approval for a bounded use. The record should state whether the deviation applies once, to a size group, to a material lot or to all future revisions. Silence should never mean acceptance. When the immediate issue is resolved, the team should either incorporate the change into a new controlled revision or close the temporary state and restore the prior approved pack.

9.7

Release authority must be explicit

A generated tech pack should never become a purchase or production instruction merely because it is complete in form. GFT proposes four distinguishable states: `draft`, `technically reviewed`, `sample approved` and `bulk released`. Draft permits structured development. Technical review confirms that the specified geometry, materials and instructions have been examined. Sample approval records acceptance of the defined physical evidence. Bulk release authorises a particular specification bundle for its intended downstream use. The states are not interchangeable, and a later state must retain the evidence on which it depends. [S19] [S22]

Designer, technician, quality and buyer approvals answer different questions. Design approval protects intent. Technical approval addresses pattern, construction and specification coherence. Quality approval confirms that the agreed acceptance evidence is present. Buyer approval records the external or internal product decision where that role applies. One named role should own final release after confirming that every required approval covers the same revision. A small team may combine roles, but the control should then address self-review and make the combined authority explicit rather than inventing additional staff.

The release manifest is the final unit of authority. It should identify the style, pack revision, pattern and grading files, material and trim revisions, artwork where applicable, construction instructions, point-of-measure table, approved exceptions, readback record and approvals. Persistent identifiers or file hashes can help detect substitution, but they do not establish that the content is correct. The accountable release decision must refer to the inspected artefacts, not only to a filename or automated status. [S20] [S22]

Receiving acknowledgement closes the communication loop. The receiving role should confirm that the manifest is legible, required attachments are present, the CAD packet has passed readback and no unresolved critical hold remains. If the receiver must interpret an ambiguity, the pack is not yet complete for that workflow. The question should return to the designated owner and the answer should enter a controlled revision. Informal clarification in a message thread may solve an immediate problem, but it leaves the next shift or supplier exposed to the same ambiguity.

This receiving test should be repeated at each organisational boundary that will use the pack. An internal sample room may accept a local CAD preset and a shared material library, while an external sampling partner may need explicit settings and additional references. The sender should not treat acceptance by one recipient as evidence of universal readability. Each receiving workflow needs its own manifest,

readback evidence and acknowledgement, linked to the same controlled source revision. That boundary-specific record shows where interpretation work remains and prevents a successful internal review from concealing an untested external handoff.

Release holds should be specific and actionable. A geometry hold names the failed panel, dimension or seam relationship. A material hold names the unresolved record or test. A rights hold identifies the affected artwork and missing permission. A fit hold identifies the sizes or activities not approved. A readback hold identifies the exchange setting or annotation that failed. Each hold needs an owner, evidence requirement and authorised return state. A generic `not approved` status conceals the work needed and invites local workarounds. NIST's framework supports contextual risk decisions and documentation of residual risk rather than a universal acceptance rule. [S19]

Exceptions require equal precision. An approved deviation should identify its scope, reason, affected revisions, compensating check and expiry or review condition. It should not quietly become a permanent change to the base block or material standard. Where an exception affects product safety, legal compliance or buyer obligations, the appropriate specialist review remains necessary. This chapter provides an engineering control framework, not legal, tax or investment advice.

The receiving workflow should also return outcome evidence. Specification escapes, corrected imports, sample deviations and production questions should link back to the released field and revision. Useful measures include required-field completeness, first-pass readback, corrected-import frequency, exception age, specification escapes and correction effort per accepted style. Each measure needs a stated denominator and product scope. A rising pass rate is meaningful only if the test set and critical checks remain comparable, and it must not be improved by removing difficult cases. [S19] [S20]

For a greenfield maker, the practical sequence is therefore clear. Define the field contract and unresolved states before generating documents. Agree the exchange manifest and prove readback on representative patterns. Resolve BOM references to physical inputs. Carry dependencies through every change. Release only the exact bundle that the accountable roles and receiving team have accepted. GFT's Cameroon textile programme supplies the staged operating context for this recommendation, but no local generative-specification acceptance result is claimed here.

A specification is complete when its receiver can act without guessing, and its approver can see exactly what was accepted.

CHAPTER SOURCE NOTES

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S22 CLO Virtual Fashion. 2026-07-15.

S19 National Institute of Standards and Technology. 2023-01-26.

S20 National Institute of Standards and Technology. 2024-07-26.

S32 CLO Virtual Fashion, CONNECT Support. 2025-04-17.

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

10

Technicians must own the learning loop

AI assistance should expose and preserve pattern-making judgement rather than obscure who is responsible for a garment.

CHAPTER 10
DECISION TEST

AI assistance should expose and preserve pattern-making judgement rather than obscure who is responsible for a garment.

A greenfield apparel maker should adopt generative design and pattern-assistance software only when human pattern technicians remain the accountable owners of the specification and the learning loop that improves it. [S15] Across Chapters 3 to 9, software can propose candidates, surface historical edits and run deterministic geometry checks. [S09] [S22] Software cannot take responsibility for a cutting instruction that wastes cloth or fails a buyer fit review. [S19] Treating a model as an autonomous pattern maker hides the diagnostic reasoning that junior technicians must learn and leaves nobody able to repair the specification when an unfamiliar style, local fabric lot or network outage breaks the automated path. [S15] [S20] Green Future Tech Limited (GFT), Hong Kong SAR, recommends inverting that premise: in GFT's Cameroon textile programme (prepared 31 January 2024, sources refreshed 2026), digital tools should record how experienced technicians diagnose and correct flawed outputs so every intervention strengthens both apprentice training and the version-controlled specification library. [S15] [S1005] Nothing in this chapter constitutes investment, legal, employment or tax advice.

10.1

Skills must be described as tasks rather than software brands

GFT should define technical competence through garment-engineering tasks, not fluency in a named software interface. [S1005] A technician who knows where a menu button sits in one vendor's draping tool may still approve an open panel boundary, miss a reversed sleeve notch or export a file whose millimetre scale shifts in the cutting room's marker system. [S12] [S22] Conversely, a technician who understands how collar stand curvature governs neck closure and how warp-versus-weft recovery alters chest ease can transfer that judgement across parametric programmes, tokenised pattern editors and conventional CAD workstations. [S02] [S27] Defining competence by observable tasks rather than software brands protects the factory against vendor lock-in and sets a clear curriculum before any pilot begins. [S19] [S22]

10.2 SIX HANDOFF TASKS DEFINE THE TECHNICAL CAPABILITY BASELINE

The workflow across Chapters 3 to 9 translates into six capabilities that a development team must verify: [S09] [S22]

1. **Brief interpretation and constraint setting:** translating a concept into bounded body-block permissions, fabric-weight envelopes, size ranges and excluded features before prompting. [S09] [S26]
2. **Parametric and geometric pattern editing:** inspecting two-dimensional boundaries, vertices, dart angles, grainlines and cut quantities, and modifying parameters or curves so edits propagate predictably. [S02] [S10] [S1004]
3. **Seam topology and construction specification:** verifying edge correspondence, sewing direction, easing ratios, stitch class, seam allowance and reinforcement order. [S12] [S14]
4. **Physical material reconciliation:** linking digital presets to measured weight, thickness, fibre content and directional stretch, and flagging when a lot change invalidates a simulation. [S11] [S32]

5. **Size-aware fit and graded ease evaluation:** auditing point-of-measure tables, local ease distributions and posture strain across sizes rather than trusting a smooth virtual try-on render. [S13] [S27]
6. **Exchange readback and release sign-off:** exporting patterns under Standard DXF, AAMA or ASTM conventions, reopening the file in the receiving CAD system to confirm dimensions and notches, and signing release. [S22]

Through these six tasks, differences among research representations become a training question rather than an abstract modelling debate. [S08] [S09] [S10] Whether a model emits domain-specific code (`Design2GarmentCode`), continuous parameter JSON (`ChatGarment`), discrete pattern tokens (`AIpparel`) or a hierarchical specification language (`PatternGSL` or `GarmentWeaver`), the release gate requires the same human judgement: reading the structure, verifying constraints in CAD and signing for manufacturability. [S08] [S09] [S10] [S14] [S16] Research on aligning constraint generation with design intent in parametric CAD reports that solver-aligned training fully constrains 93 percent of engineering sketches versus 34 percent for a naive supervised baseline and 8.9 percent without fine-tuning, so geometry updates predictably when a designer modifies a parameter. [S1004]

10.3 INDUSTRIAL CONTEXT REQUIRES BUILDING SUPERVISORY AND TECHNICAL DEPTH LOCALLY

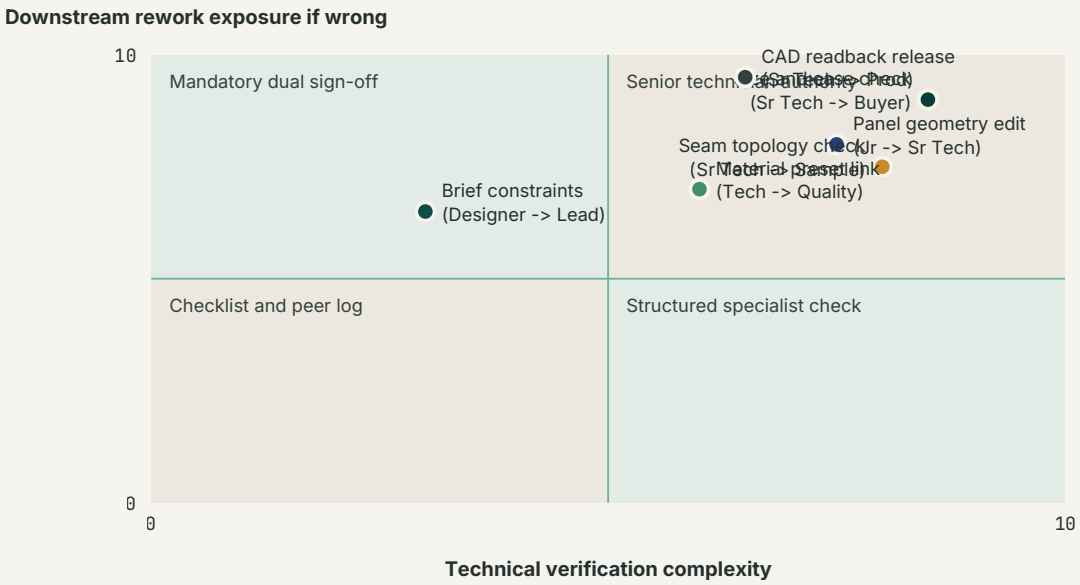
This task-based framework connects to the industrial gap motivating GFT's Cameroon textile programme. [R06] Inherited register records indicate that Cameroon exports about 95 percent of its raw cotton fibre and processes about 5 percent locally, against a 50 percent target by 2030 under the National Development Strategy (SND30). [R06] [R07] Historical reporting on the public manufacturer CICAM recorded July 2023 capacity utilisation at 33.7 percent at Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, with underutilisation costing about CFAF 2 billion in 2020. [R12] Moving from raw-fibre export or underutilised capacity to competitive garment making requires more than machinery; in GFT's feasibility study and performance-management system design (prepared 31 January 2024, sources refreshed 2026), cut, make, trim and finishing depend on local supervisors, sample machinists and pattern technicians holding dimensional tolerances from the first shift. [R06] [R12]

Published user studies evaluate small studio or university cohorts; they do not measure the supply of pattern technicians in Douala, Garoua or Dibamba, and must not be used to estimate a local skills shortage. [S15] [S1005] Instead, a greenfield maker should use the six tasks as an onboarding diagnostic on the illustrative cotton shirt before assigning release authority. [S19]

The illustrative matrix in **Exhibit 39**, p. 153 maps the six handoffs by verification complexity and downstream rework exposure when an error escapes to cutting. [S19] [S22]

EXHIBIT 39 / ILLUSTRATIVE INDEX (0 TO 10)

Roles should follow the decision that can fail



ILLUSTRATIVE [S19] [S22] Illustrative governance model grounded in GFT's performance-management system design (G02) and contextual risk prioritisation (S19, S22). Coordinates represent proposed relative verification difficulty (x) and downstream physical rework exposure if an error escapes to cutting (y); no GFT headcount, named staff or field defect rates are implied.

As **Exhibit 39**, p. 153 shows, no handoff sits in a zero-consequence zone where unreviewed software output can pass directly to cutting. [S19] Brief constraint setting carries an illustrative rework exposure score of 6.5 out of 10 because an unapproved block or fabric change invalidates downstream geometry and costing. [S09] At the upper right, panel editing, seam topology verification, graded ease approval and receiving-CAD readback require senior pattern-technician sign-off. What **Exhibit 39**, p. 153 cannot establish is how many minutes each review takes in a live plant or what headcount is needed; those depend on local order volume and style complexity. [S19]

10.4

Recorded operations can make expert corrections teachable

Senior pattern makers hold much of their diagnostic skill as undocumented procedural judgement built through physical fitting and CAD revision. [S15] [S1005] Drafting manuals record static formulas for basic blocks, yet formative research confirms that practitioners rarely draft commercial styles from scratch. [S1005] Instead, they select the closest library block, compare it against the target design, and apply a path-dependent sequence of edits while checking intermediate states to see whether balance, proportion or seam compatibility is drifting. [S15] [S1005] When a senior technician leaves, or when an AI model overwrites a pattern in one jump, that intermediate reasoning vanishes. [S15]

10.5 VERSION-CONTROLLED STATE TRACKING SEPARATES PROCESS KNOWLEDGE FROM STATIC RULES

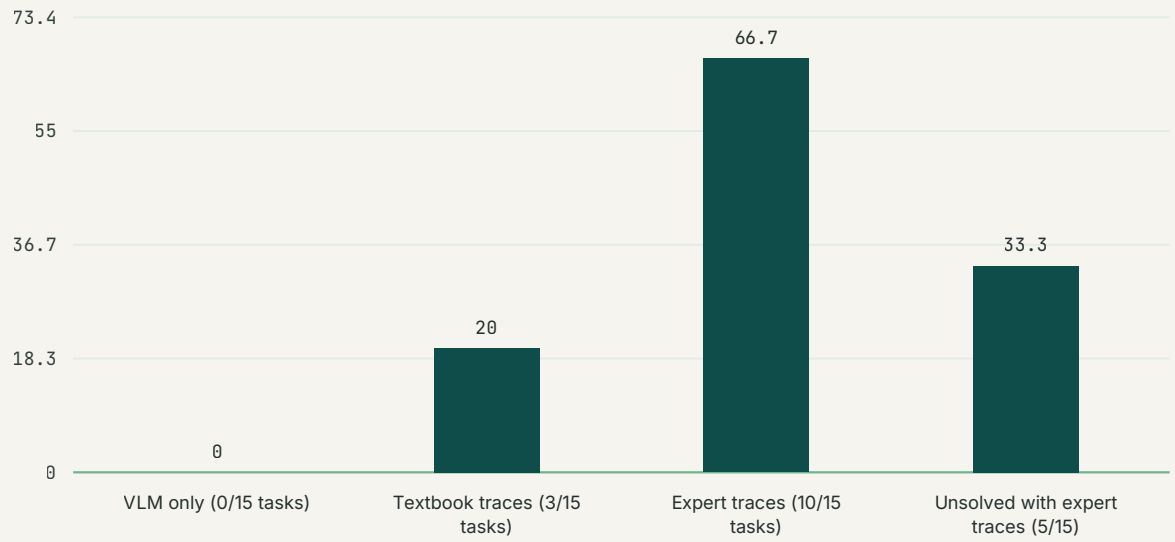
Recent human-computer interaction research shows how software can capture this procedural expertise during routine CAD work. [S15] TailorCoPilot, built on the TailorTrace backend (published August 2026), models a sewing pattern as a discrete state comprising two-dimensional panel geometry and seam topology, optional three-dimensional draped geometry and metadata. [S15] [S1005] Sitting beneath a CAD interface, TailorTrace records a senior technician's edits as semantic operations on geometric primitives (panels, edges, vertices and stitches), such as SpreadPanel() or ExtendPoint(), committing a new node to a branching history only when the state passes validation. [S1005]

The TailorCoPilot evaluation clarifies both the value and the boundary of recorded expert operations. [S15] [S1005] In a controlled study with 20 participants (10 novices and 10 advanced novices, each completing 12 trials across simple, medium and complex tasks), trace-guided assistance increased the odds of producing a valid pattern within the allotted time relative to baseline tools (OR = 2.31, 95 percent confidence interval 1.42 to 3.86, p = 0.001), reduced time (p < 0.001), lowered NASA-TLX workload (p < 0.001) and raised blinded expert ratings on a 1 to 5 scale from 1.62 to 3.92 for novices and from 2.87 to 4.27 for advanced novices. [S1005] Even more instructive is the diagnostic experiment across 15 complex tasks comparing three configurations under identical references, base patterns and inference budgets: a pretrained vision-language model (Qwen3-VL-8B) without traces, a textbook-trace variant, and the full system with expert-authored TailorTrace records. [S1005]

The sourced chart in **Exhibit 40**, p. 155 reports the share of those 15 complex tasks meeting structural validity across the three diagnostic conditions. [S15] [S1005]

EXHIBIT 40 / PERCENT OF 15 COMPLEX TASKS MEETING STRUCTURAL VALIDITY

Expert operation traces outperform static textbook rules on complex pattern revisions



SOURCED [S15] [S1005] Sourced from the TailorCoPilot diagnostic study on 15 complex pattern-making tasks (S15, S1005 Section 5.3). VLM-only (Qwen3-VL-8B without TailorTrace) succeeded in 0 of 15 tasks (0.0 percent; mean expert score 1.3/5.0); textbook-trace-only succeeded in 3 of 15 tasks (20.0 percent; mean score 1.9/5.0); expert-trace-augmented succeeded in 10 of 15 tasks (66.7 percent; mean score 4.1/5.0). Five of 15 tasks (33.3 percent) failed even with expert traces. Laboratory benchmark on 93 traces; not a GFT factory result.

Two practical conclusions follow from **Exhibit 40**, p. 155. [S1005] First, static model weights and textbook rules fail on complex multi-piece revisions: the unaugmented vision-language model succeeded on 0 of 15 tasks (0.0 percent, mean score 1.3 out of 5.0), and the textbook-trace variant succeeded on 3 of 15 tasks (20.0 percent, mean score 1.9 out of 5.0), producing local edits that broke cross-piece stitching or ruffle construction. [S1005] Adding expert-authored traces raised structural validity to 10 of 15 tasks (66.7 percent, mean score 4.1 out of 5.0), showing that the sequence in which an expert modifies panels and reconciles seams carries procedural information absent from textbooks. [S15] [S1005] Second, **Exhibit 40**, p. 155 proves why the human technician remains indispensable: 5 of 15 tasks (33.3 percent) failed even with expert traces, the corpus held only 93 traces (about 2,500 derived training pairs) under one default fabric preset, and the system failed on three-dimensional manipulations such as a twisted bust knot. [S1005] Neither S15 nor S1005 has been deployed at a GFT facility or used to train Cameroonian staff.

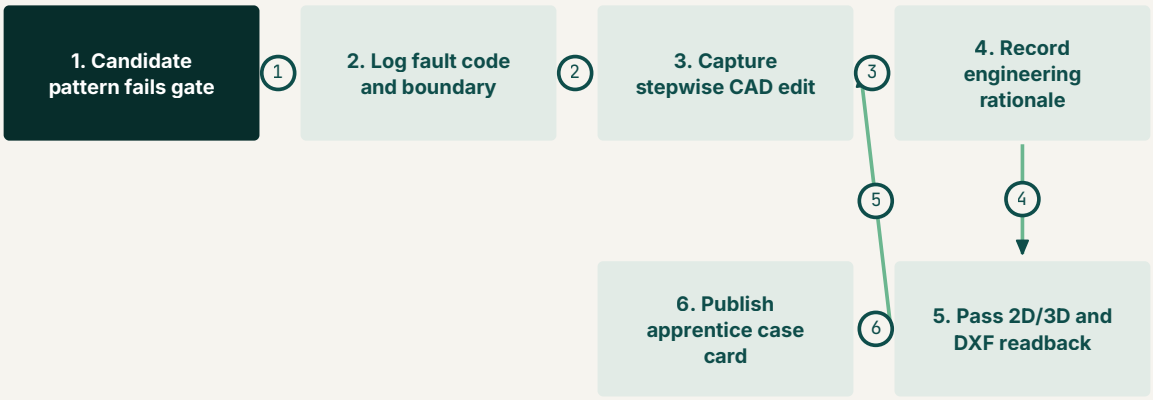
10.6 COORDINATE DELTAS NEED A RECORDED ENGINEERING RATIONALE

To turn daily review into a training asset, GFT proposes extending state-tracking logs so an expert correction records not only *what* geometry changed, but *why* the previous state failed. [S15] [S20] On our illustrative cotton shirt, suppose a generative module widens a collar stand by 8 millimetres (illustrative) to match a sketch, but misses the front neck curve and leaves sleeve cap ease at zero millimetres (illustrative) for woven poplin. If the senior technician drags vertices in CAD and overwrites the file, a junior technician sees only the finished coordinates.

Under GFT's proposed protocol in **Exhibit 41**, p. 156, every expert intervention on a failed candidate follows a six-stage loop binding the geometric operation sequence to a fault code and a concise rationale. [S15] [S20]

EXHIBIT 41 /

An expert correction should become a reusable lesson



1. Flagged panel or seam ID
2. Open editable base state
3. Operation trace + delta
4. Linked cause and rule
5. Residual fault
6. Verified state committed

ILLUSTRATIVE [S15] [S20] Illustrative GFT workflow proposal informed by version-controlled pattern-state tracking (S15) and NIST generative-AI lifecycle monitoring and human-review guidance (S20). No automated retraining or deployed GFT software pipeline is implied.

In **Exhibit 41**, p. 156, a failed candidate pattern (Stage 1) is never patched without trace. [S15] [S20] The technician logs the fault code and boundary IDs (Stage 2), repairs the base state while recording modified primitives (Stage 3), and attaches a rationale note (Stage 4), such as explaining that a collar stand extension on the shirt requires a matching neckline redraw and notch shift. [S15] Only after passing two-dimensional boundary checks, three-dimensional drape checks and DXF readback (Stage 5) is the state pair committed to the apprenticeship library as a case card (Stage 6). [S15] [S22] While **Exhibit 41**, p. 156 structures local capture, it does not guarantee that a model trained on those traces will generalise to new garment families; its primary benefit is giving human trainees an archive of real plant corrections. [S1005]

10.7

Human-AI ideation should preserve design authorship decisions

The learning loop starts upstream of pattern geometry, when a designer uses image or sketch generation to explore silhouette, collar, pocket or colourway variations. [S07] [S26] Treating ideation as an unrecorded prompt lottery loses two assets: junior designers never learn why plausible images were rejected for production, and the enterprise loses the documentary record of human creative choices. [S21] [S1001]

10.8 HYBRID SKETCH WORKFLOWS REQUIRE DECISION LOGS FOR ACCEPTED AND REJECTED CANDIDATES

Research on human-AI fashion ideation shows how modular systems keep designers in control. [S07] The HAIGEN system (submitted August 2024, revised July 2026) divides assisted ideation into four modules: a cloud text-to-image generator (T2IM), and three local modules that convert images into a sketch library (I2SM), recommend related sketches for refinement (SRM), and apply colour and style transfer (STM). [S07] Keeping sketch refinement local avoids uploading proprietary sketches to the cloud, though as noted in Chapters 2 and 11, that architecture is a research design choice rather than an audited security guarantee. [S07] [S20]

For staff development, the critical work occurs between those modules. [S1001] Research on the metacognitive demands of generative AI (S1001) shows that generative tools shift heavy monitoring and control demands onto users across prompting, evaluating outputs for errors, and integrating tools into workflows. [S1001] A junior designer who accepts a sketch of the illustrative shirt without noticing uneven button spacing or a missing back yoke seam pushes hidden rework onto the pattern technician. [S09] [S1001]

GFT recommends that the ideation log retain three records per style brief:

- **The constrained intent record:** the initial human sketch or base-block reference, the permitted modification envelope on the illustrative shirt, and negative constraints. [S07] [S09]
- **The rejection rationale log:** at least two rejected candidates per accepted concept, annotated with the manufacturing, material or rights reason each was turned down. [S01] [S20]
- **The human modification trace:** the manual redrawing, vector flat edits and construction callouts performed to create an unambiguous technical flat for Chapter 4 pattern drafting. [S07] [S09]

10.9 RECORDING HUMAN CONTRIBUTION SUPPORTS GOVERNANCE WITHOUT GUARANTEEING COPYRIGHT

Logging human selections and manual sketch modifications also supports the rights governance in Chapter 11. [S21] In January 2025, the U.S. Copyright Office released Part 2 of its Artificial Intelligence Report, concluding that generative AI outputs are protected by copyright in the United States only where a human author determined sufficient expressive elements, distinguishing human creative control from merely providing prompts. [S21]

A maker in Cameroon or Hong Kong SAR must read that policy within its territorial boundary. [S21] The U.S. Copyright Office report states United States practice; it does not settle copyrightability or industrial design protection under the laws of Cameroon, the Organisation Africaine de la Propriété Intellectuelle (OAPI) or Hong Kong SAR, and logging prompting time does not create legal authorship. [S21] GFT recommends recording original sketches, manual edits and flat revisions because those records teach junior staff disciplined design control while preserving factual evidence of human contribution for legal counsel. [S20] [S21]

10.10

A training exercise needs observable acceptance criteria

When a factory introduces AI-assisted pattern tools, management may judge training by how fast novices finish a trial or how confident they feel. [S15] [S1005] Both measures are unsafe in production engineering. [S19] The TailorCoPilot authors note that their 20-participant study measured short-term performance across 12 trials in a controlled setting and could not establish whether trace-based support builds durable independent skill over time. [S1005] A trainee who finishes a shirt revision ten minutes faster, yet misses that the exported DXF file dropped seam allowances or mis-graded the collar neckband, has learned to click through software rather than engineer a garment. [S19] [S22]

10.11 COMPETENCY TESTING MUST INJECT REALISTIC SPECIFICATION FAULTS

Following the NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0), which states that an organisation must define its own risk tolerance and evaluate performance in context (MEASURE 2.3), GFT proposes qualifying technicians through fault-injected, unseen tasks on the illustrative cotton shirt: [S19] [S20]

1. **Invalid panel and topology repair:** repair an unclosed cuff polygon, a self-intersecting collar peak and a reversed sleeve-to-armhole seam direction in CAD, and explain why cutting or sewing would fail. [S12] [S14] [S30]
2. **Material substitution reconciliation:** compare a simulation run on a lightweight cotton preset against a laboratory datasheet for a heavier, lower-stretch woven lot, flag the mismatch and trigger drape and shrinkage re-checks. [S11] [S32]
3. **Size-mismatch and graded-ease diagnosis:** inspect a smooth virtual try-on render paired with a graded table where the extra-large chest has negative ease, reject the render and restore prescribed ease across sizes. [S13] [S27]
4. **Receiving-CAD DXF round-trip readback:** export a graded pattern using documented DXF options (Standard DXF, AAMA or ASTM) and reopen it in the receiving CAD system to catch a seeded unit-scale shift or dropped notch before release. [S22]

The illustrative table in **Exhibit 42**, p. 159 sets out these four tasks, their seeded faults, required trainee evidence and automatic fail triggers. [S19] [S22]

EXHIBIT 42 / ASSESSMENT PROTOCOL FIELDS

A training assessment needs unseen tasks

Unseen task	Seeded fault to detect	Required trainee evidence	Pass criterion	Automatic fail trigger
1. Panel and seam repair	Open cuff boundary; reversed sleeve seam pair	Annotated 2D panel log and repaired topology	100% seeded faults repaired and explained	Releasing an open or reversed seam to cutting
2. Material substitution	Preset weight/stretch differs from swatch record	Updated property table and drape re-check request	Mismatch flagged; physical fit check triggered	Approving simulation on unverified preset
3. Graded fit diagnosis	Fit-blind VTO render hides XL chest pinch	Point-of-measure audit and ease correction	Negative ease detected across all graded sizes	Signing off fit from VTO appearance alone
4. DXF export readback	10x unit scale shift and dropped collar notches	Receiving-CAD readback table vs tech pack	All dimensions, notches and allowances match	Treating file open success as valid readback

ILLUSTRATIVE [S19] [S22] Illustrative GFT technician qualification table applying NIST AI RMF contextual measurement principles (S19) and documented CAD exchange controls (S22) to the chapters 4 to 9 handoffs. No trainee scores or Cameroonian test results have been collected.

As **Exhibit 42**, p. 159 shows, every pass criterion requires both a repaired state and a written fault explanation. [S19] Because a single sitting cannot prove retention, GFT recommends repeating unseen fault checks when a technician qualifies for a new garment family or when a model update alters behaviour. [S19] [S20]

WORKED EXAMPLE: SCORING A TRAINEE COHORT ON UNSEEN SPECIFICATION FAULTS

Consider an illustrative evaluation of 8 junior technicians (N = 8 illustrative) attempting the 4 unseen shirt tasks in **Exhibit 42**, p. 159, giving 8 x 4 = 32 illustrative trials. Each task has 3 seeded defects (1 critical release-blocking defect and 2 major defects), for 4 x 3 = 12 defects per trainee and 32 x 3 = 96 across the cohort (illustrative).

Suppose the cohort is split into two illustrative groups of 4 trainees:

- **Group A (Speed-prioritised AI users, 4 trainees, 16 trials):** averages 18 minutes per task against a 35 minute manual baseline (48.6 percent time reduction, illustrative), but catches only 7 of 12 defects per trainee (58.3 percent capture, illustrative) and releases 1 critical defect in every task, or 4 per trainee.
- **Group B (Verification-disciplined users, 4 trainees, 16 trials):** averages 29 minutes per task (17.1 percent time reduction, illustrative), catches 11 of 12 defects per trainee (91.7 percent capture, illustrative), catches 100 percent of critical defects (4 / 4 per trainee, illustrative), and safely escalates the 1 remaining collar-pleat fault.

Under GFT's defect-gated rule, any released critical defect sets the trial score to 0, giving Group A 0 qualified tasks (0 / 16 = 0 percent pass rate, illustrative) and 16 escaped cutting errors that, at Chapter 12's illustrative cost of XAF 30,000 per ruined sample round, waste 16 x XAF 30,000 = XAF 480,000. Group B achieves 12 direct passes and 4 safe escalations (100 percent safe release governance, illustrative) while saving 4 x (35 - 29) = 24 illustrative minutes per trainee. All figures are hypothetical teaching arithmetic.

10.12

The interface should show uncertainty and correction paths

Even a trained technician can be pushed into approving a faulty pattern if the software interface projects false certainty. [S20] [S1002] Many generative tools display a smooth three-dimensional render or pre-filled table as the primary view while hiding geometric warnings, assumed material defaults and unverified seam pairings behind secondary menus. [S13] [S20] Under schedule pressure, an interface that makes clicking *Approve* effortless while making verification slow induces overreliance. [S1002]

10.13 LOWERING THE COGNITIVE COST OF VERIFICATION REDUCES OVERRELIANCE

Human-computer interaction research explains why merely adding AI explanations often fails to stop human errors. [S1002] Across 5 experiments (N = 731 participants) studying overreliance in AI-assisted decision-making, Vasconcelos et al. (S1002) modelled engagement with AI explanations as a strategic cost-benefit choice: users weigh the cognitive effort of reading and verifying an explanation against the effort of trusting the AI and the value of being right. [S1002] Higher task or explanation difficulty increased overreliance on wrong predictions, whereas explanations reduced overreliance when they lowered verification cost. [S1002] Tankelevitch et al. (S1001) similarly show that generative systems impose heavy metacognitive monitoring demands when users must evaluate outputs without structural cues. [S1001]

In apparel specification software, a prose paragraph from a language model is a high-cost, low-precision explanation. [S20] [S1002] By contrast, a visual overlay of before-and-after two-dimensional panel outlines, a table of point-of-measure deltas in millimetres, and a checklist of failed boundary or seam rules lower verification cost so a technician can spot a defect immediately. [S15] [S1002]

10.14 FIVE INTERFACE CONTROLS KEEP AUTHORITY WITH THE TECHNICIAN

Guided by the NIST Generative AI Profile (`NIST AI 600-1`) and the verification-cost principle, GFT proposes five interface requirements for any generative specification tool evaluated in a pilot: [S20] [S1002]

1. **Side-by-side state diffing before commit:** display the approved base state beside the proposed state, highlighting changed panel boundaries, vertex shifts, seam-length deltas and modified bill-of-materials fields. [S15] [S1002]
2. **Explicit provenance badges on every field:** label each value as `MEASURED` (linked to a laboratory, swatch or CAD readback record), `INFERRED` (model-estimated, as in `S11`) or `MISSING` (requiring human entry). [S11] [S20]
3. **Hard-stop validation banners:** if deterministic checks find an open panel contour, self-intersection, unpaired seam, missing cultural permission or failed DXF read-back, disable *Release* until repaired or formally overridden by a senior technician. [S14] [S19] [S22]
4. **Three explicit action paths at each step:** expose `Accept with check log`, `Reject and revert to previous state` and `Escalate to senior technician with note`, feeding into **Exhibit 41**, `p. 156`. [S15] [S20]
5. **Uninterrupted manual CAD fallback:** in line with GFT's digital strategy (prepared 31 January 2024, sources refreshed 2026), which treats power and network reliability as design constraints, the local workstation must open, edit and export all two-dimensional patterns and tech packs offline when cloud APIs are unreachable. [S22]

> **GFT NOTE**
>
> Never adopt a generative design or pattern interface that auto-fills buyer approval fields, hides estimated fabric properties behind default presets or requires a live cloud connection to open existing two-dimensional pattern files. [S11] [S22] During vendor due diligence (Chapters 11 and 13), require the supplier to demonstrate how a technician inspects a failed seam match, rejects an incorrect step, records the rationale and completes DXF export manually with the AI service switched off. [S15] [S20] [S22]

10.15

Role separation supports review without inventing headcount

A learning loop prevents escaped errors only when the person who edits a candidate pattern is not the sole person who approves its release to cutting. [S19] [S20] In GFT's performance-management system design (prepared 31 January 2024), operational governance assigns every key performance indicator and quality gate to a defined

functional role within a common system of record. Extending that rule to generative specification separates creation, verification and release across six functional roles, without inventing a fixed headcount before an order book exists. [S19]

10.16 SIX FUNCTIONAL ROLES GOVERN THE SPECIFICATION LIFECYCLE

GFT proposes assigning accountability across six roles covering the handoffs from Chapter 3 to Chapter 9: [S19]

- **Design intent owner (Designer / Product Developer):** authors the constrained brief, checks reference permissions, curates candidates, logs rejections and delivers the technical flat. [S07] [S09]
- **Pattern geometry owner (Junior or Intermediate Pattern Technician):** selects the base block, executes parametric and assisted edits, resolves two-dimensional geometry and prepares initial grading. [S02] [S15]
- **Technical specification verifier (Senior Pattern Technician):** audits panel boundaries, seam topology, ease allowances, grading and receiving-CAD DXF readback, and approves additions to the trace library. [S12] [S15] [S22]
- **Material and quality verifier (Quality / Laboratory Lead):** validates physical swatch tests, reconciles simulation presets against supplier lots and signs off sample construction. [S11] [S32]
- **Rights and provenance reviewer (Procurement / Compliance Lead):** verifies input licences, vendor commercial terms, traditional cultural expression records and body-data consent. [S20] [S21] [S23]
- **Release authority (Technical Development Manager / Plant Operations Lead):** confirms that all evidence fields, physical fit approvals and readback checks are complete and signs final release. [S19]

In a lean launch team, one person may hold multiple roles if conflict controls are documented: the `Pattern geometry owner` must never serve as the sole `Technical specification verifier` or `Release authority` for that same revision. [S19] [S22]

METHOD

Method: conflict-of-interest rule for combined roles in a lean launch team

1. **Create six sign-off fields:** maintain separate tech-pack signature fields for Design Intent, Pattern Geometry, Technical Verification, Material/Quality, Rights/Provenance and Final Release. [S19] [S20]
2. **Allow non-conflicting combinations in Phase 1:** permit one qualified staff member to combine `Design intent owner` + `Rights initial preparer`, or `Technical specification verifier` + `Release authority`, once qualified on **Exhibit 42**, p. 159. [S19]
3. **Enforce two-person verification:** block release if the same user ID appears as both `Pattern geometry owner` and `Technical specification verifier`. [S19] [S20]
4. **Log provisional overrides:** if an urgent sample requires a single senior technician to edit and verify a pattern, tag the file `SINGLE-REVIEWER PROVISIONAL`, require Quality Lead sign-off on the first physical prototype, and audit the file at the weekly review. [S19]

10.17 WORKFORCE PLANNING MUST RESPECT STATUTORY LABOUR BASELINES AND AVOID UNMEASURED RATIO CLAIMS

Vendor marketing sometimes claims that one technician with generative software can replace several manual pattern makers. GFT rejects importing any such staffing-reduction multiplier into a greenfield feasibility model. [S15] [S20] As shown in **Exhibit 40**, p. 155, even the strongest published trace-based system failed on 33.3 percent of complex revisions in a diagnostic study and has no longitudinal factory evidence. [S1005] Reviewing, repairing and logging rejected candidates also adds verification work that must be timed locally (Chapter 12). [S1001] [S1002]

Any local staffing plan for a Cameroonian apparel facility must rest on measured workloads and lawful employment terms. [R14] [R16] Inherited register records note that Cameroon's statutory minimum wage (SMIG) rose from XAF 41,875 a month in 2023 to XAF 43,969 a month from 23 February 2024 (about US$76, up 5 percent), in an economy where informal employment is about 90 percent and the standard working week is 40 hours (48 hours maximum, overtime at 1.2 times), while multi-sector average salary estimates are about US$359 a month and Cambodia's garment minimum wage stood at US$204 a month in 2024. [R14] [R15] [R16] [R17] Those figures provide macro-level context from the register, not a pay scale for skilled CAD technicians. [R14] [R16] Setting technician pay, shift rosters and training budgets belongs to `wp09` and `supervisor-pipeline`, informed by local labour law review and measured pilot workloads. [R14] [R16]

10.18

Local evaluation must test transfer rather than assume it

The final step in the learning loop is verifying whether assisted work and trace-based training transfer to new garments, new fabrics and unassisted problem solving. [S1003] [S1005] In machine learning for garment reconstruction, researchers explicitly test

whether models generalise to unseen topologies: NeuralTailor (S1003) evaluated whether two-dimensional pattern reconstruction generalised to garment categories with pattern topologies unseen during training, while GarmentCodeData (S30) documented how synthetic datasets depend on their programmed distributions and quality filters. [S30] [S1003] The same rule governs human learning. [S1005] A trainee who replays an expert trace on a V-neck dress (S1005) has not yet proved that they can engineer a pleated bush-shirt pocket, adjust ease for heavy twill or diagnose a grading error with the assistant turned off. [S19] [S1005]

10.19 TRANSFER TESTING REQUIRES HELD-OUT GARMENT VARIANTS AND UNASSISTED TRIALS

To prevent TailorCoPilot's short-term user-study gains (OR = 2.31 across 12 trials) from being mistaken for permanent local capability, GFT proposes a three-tier transfer check during the Phase 3 pilot in Chapter 13: [S15] [S1005]

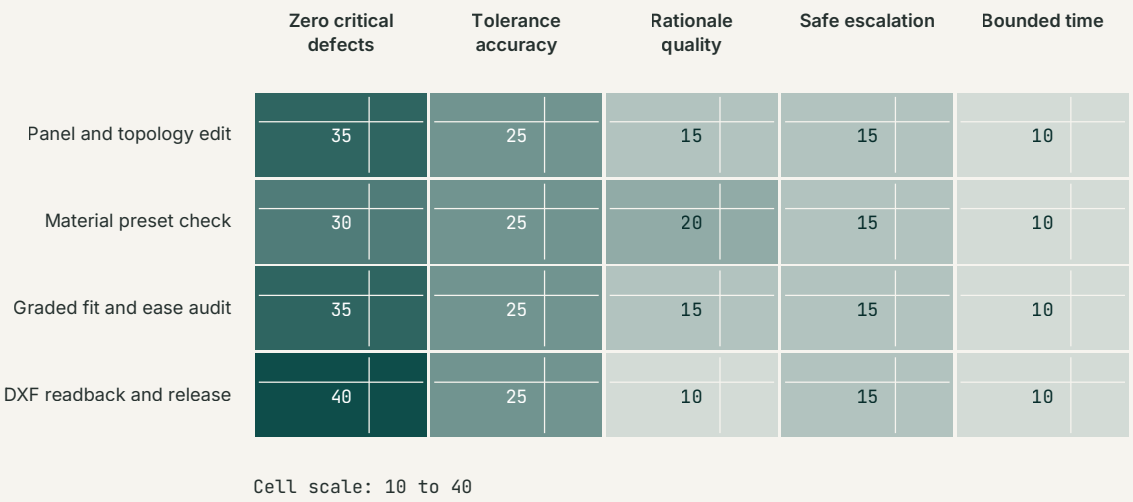
1. **Assisted performance on familiar blocks:** measure time, residual defects and DXF readback accuracy on minor variants of the base cotton shirt. [S15] [S22]
2. **Assisted transfer to held-out variations:** test technicians on unseen structural variations excluded from onboarding, such as a two-piece sleeve with a functional placket or a fabric lot with different warp/weft shrinkage. [S19] [S30]
3. **Unassisted CAD retention check:** require every technician to complete at least one held-out modification and full DXF readback in manual CAD with the assistant disabled, verifying that tool use builds rather than erodes independent judgement. [S20] [S22] [S1005]

10.20 TIME AND QUALITY MUST BE REVIEWED TOGETHER ACROSS EVERY TASK

When management reviews technician progress, speed and quality must be assessed together so faster completion never hides defect escapes. [S19] [S20] The illustrative weave-matrix in **Exhibit 43**, p. 165 sets out GFT's proposed scoring weights across the four core assessment tasks, showing how structural validity, dimensional tolerance, rationale quality and escalation discipline outweigh speed. [S15] [S19]

EXHIBIT 43 / ILLUSTRATIVE RUBRIC WEIGHT (PERCENT OF 100 PER TASK)

Time and quality must be reviewed together



ILLUSTRATIVE [S15] [S19] Illustrative GFT evaluation rubric weights (each row sums to 100 percent) grounded in NIST AI RMF contextual measurement (S19) and TailorCoPilot quality dimensions (sewability, structural balance, reference conformity and workload in S15, S1005). Completion time receives at most 10 percent weight and is scored only after critical defect count equals zero. No trainee results are fabricated.

In **Exhibit 43**, p. 165, each row sums to 100 illustrative rubric points across five columns, and the four quality and governance columns carry 90 percent of the weight on every task. [S19] Bounded completion time carries only 10 percent weight and is strictly gated: releasing a critical defect in panel geometry, material matching, graded ease or DXF readback fails the trial regardless of speed. [S19] [S22] Conversely, a technician who spots that an agentic loop (S18) or pattern generator (S16) produced an unsewable layered detail and escalates it with a diagnostic note earns full credit on escalation. [S16] [S18] These weights link to `kpi-cascade` and `supervisor-pipeline` as illustrative proposals to be calibrated with the technical lead. [S19]

> **WARNING**
>
> **Unresolved constraint: short-term assisted success does not prove long-term skill retention**
>
> Published human-AI pattern-making studies test short sessions on bounded garment sets (S15, S1005) and cannot show whether months of assisted work strengthen or dull a junior technician's independent diagnostic skill when software fails. [S1005] A greenfield maker must run periodic unassisted CAD checks on held-out styles during Phases 2 and 3, logging error rates by task in wp14 before expanding generative tools or changing training budgets in wp09. [S19] [S20]

10.21 COMPANION PAPERS OWN THE SURROUNDING WORKFORCE BUDGET AND DATA SPINE

Clear boundaries across the white-paper library keep accountability sharp. [S19] This chapter specifies the six technical capabilities, the operation-trace loop (**Exhibit 41**, p. 156), the unseen fault examination (**Exhibit 42**, p. 159), the verification-centred interface rules and the six-role sign-off separation. [S15] [S19] [S20] Companion paper wp09 owns recruitment, vocational partnerships, wages and training budgets; wp14

owns the version-controlled data spine and model registry storing operation traces; and wp20 owns three-dimensional digital sampling and print execution. [S15] [S20] Weaving human diagnostic judgement into every recorded revision ensures that the factory's people grow more capable with every style they engineer.

A model can propose a shape, but only a technician who understands why a seam fails can sign for the garment and teach the next shift.

CHAPTER SOURCE NOTES

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S22 CLO Virtual Fashion. 2026-07-15.

S19 National Institute of Standards and Technology. 2023-01-26.

S20 National Institute of Standards and Technology. 2024-07-26.

S1005 Authors via arXiv HTML; UIST 2026. 2026-08-30.

S12 Authors via arXiv; WACV 2026. 2026-02-25.

S02 Authors via arXiv; ACM Transactions on Graphics. 2023-06-06.

S27 Authors via arXiv; Computers & Graphics. 2026-06-28.

S26 Authors via arXiv. 2025-04-17.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S1004 Authors via arXiv. 2025-04-17.

S14 Authors via arXiv. 2026-06-23.

S11 Authors via arXiv. 2026-01-14.

S32 CLO Virtual Fashion, CONNECT Support. 2025-04-17.

S13 Authors via arXiv; SIGGRAPH 2026. 2026-04-09.

S08 Authors via arXiv. 2024-12-05.

S16 Authors via arXiv. 2026-08-31.

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

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

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

S07 Authors via arXiv. 2024-08-01.

S21 U.S. Copyright Office. 2025-01-29.

S1001 Authors via arXiv; ACM CHI 2024. 2023-12-18.

S01 Authors via arXiv. 2025-09-04.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S1002 Authors via arXiv; Proceedings of the ACM on Human-Computer Interaction. 2022-12-13.

S23 World Intellectual Property Organization. 2026.

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

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

R15 wage.is. Publication date unrecorded; inherited source-text check.

R17 just-style. Publication date unrecorded; inherited source-text check.

S1003 Authors via arXiv; ACM Transactions on Graphics. 2022-01-31.

S18 Authors via arXiv. 2026-10-11.

GREEN FUTURE TECH

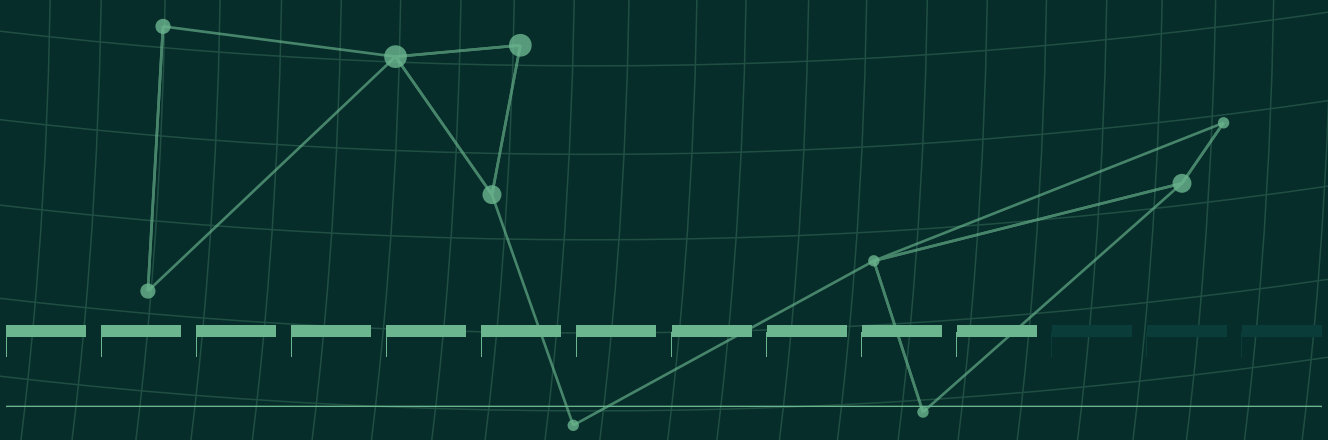
LIMITED / HONG KONG SAR

11

PART III: ADOPTION REQUIRES PEOPLE, RIGHTS AND MEASURED ECONOMICS / GOVERNANCE BEFORE SCALE

Rights and model risks need versioned records

Procurement must establish what may be uploaded, generated, retained and used before sensitive design work enters a model.



CHAPTER 11
DECISION TEST

Procurement must establish what may be uploaded, generated, retained and used before sensitive design work enters a model.

A maker should release a generated design only when its permission, transformation history and technical acceptance refer to the same identifiable revision. Within GFT's Cameroon textile programme, prepared in January 2024 with source notes refreshed in 2026, the industrial objective is to move from a sector baseline where about 95 percent of cotton fibre is exported raw and about 5 percent is processed locally towards higher domestic value retention, alongside the stated national target of 50 percent local processing by 2030. [R06] [R07] Capturing that value requires garment specifications that buyers and factories can verify. A software subscription, an attractive image and a successful export answer different questions. GFT proposes a versioned evidence envelope that connects input references, authorised destination, model configuration, human decisions and release conditions. [S20] Unknown permissions remain unresolved decisions, rather than becoming implied consent as files move between teams. The controls below are proposed operating practices, not an account of an installed GFT system. This chapter is not investment, legal or tax advice; applicable obligations and particular designs require jurisdiction-specific review.

11.1

A permission to use software differs from rights in an output

Adobe's Firefly FAQ, updated in May 2026 and opened for this chapter, states that for features without the beta label, users can use generated outputs in commercial projects. [S31] That statement concerns the vendor's contractual conditions for a particular service tier and feature set. It does not establish that a buyer supplied an input reference lawfully, that a generated garment graphic avoids another party's intellectual property rights, or that an output receives statutory copyright protection in the territories where the garment will be made or sold. [S31] The actual feature, model route and contract still need to be identified before any design enters a commercial tech pack.

GFT recommends separating the rights review into five distinct decisions: software permission, input permission, model-training permission, output-use conditions and the substantive infringement assessment for the intended product. Each review should name the exact object being assessed and its authorised destination. An image may be approved as an internal mood-board reference without being approved for upload to an external model, inclusion in public advertising or rotary screen printing on a garment. A file-level label such as "licensed" is therefore insufficient unless its scope, territory and expiry travel with the asset record. This is a proposed control structure, not a declaration about a particular supplier licence.

For the running cotton-shirt workflow, the proposed record would distinguish a buyer's silhouette reference, a maker's base pattern block, a surface artwork tile and the generated composite presentation. The reviewer should identify the relevant permission for each input component and then assess the proposed combination. A permission granted for internal design exploration should not silently become permission for commercial garment reproduction. When the intended activity changes from internal review to buyer sampling or bulk manufacture, the accountable owner should

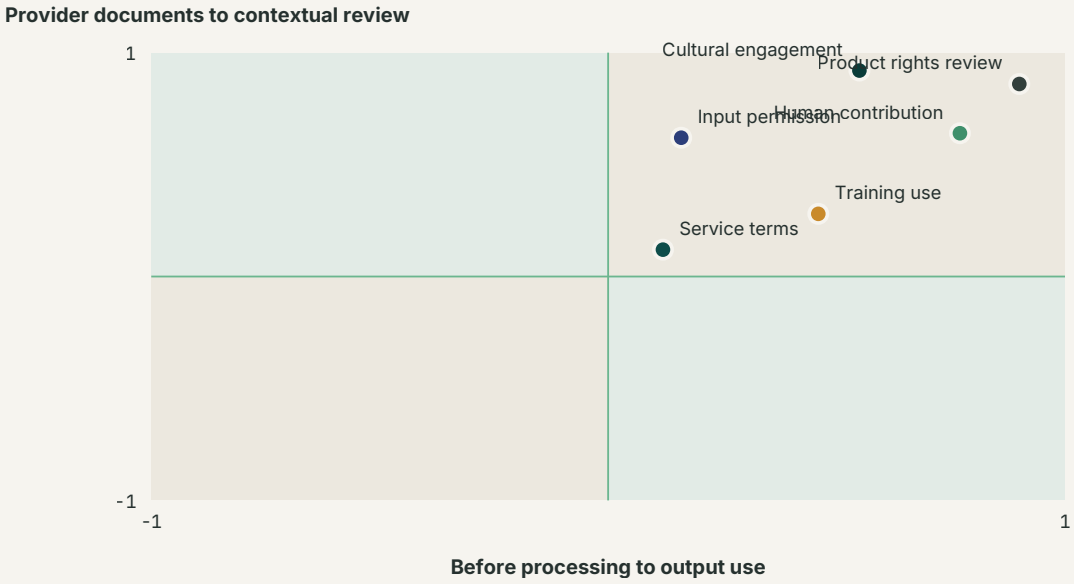
reopen the decision and record the reason, relevant correspondence and affected specification revisions. Retaining an earlier approval preserves history; it does not extend its scope.

Official policy documents show why a maker cannot rely on one universal rule for copyrightability or disclosure. In the United States, the Copyright Office's March 2023 registration guidance states that applicants have a duty to disclose artificial-intelligence-generated content in registration applications. [S1101] Its January 2025 Part 2 report concludes that generative outputs receive copyright protection only where a human author has determined sufficient expressive elements, distinguishing creative arrangements or modifications from supplying prompts. [S21] By contrast, the Hong Kong SAR Commerce and Economic Development Bureau noted in its July 2024 consultation announcement that under the existing Copyright Ordinance (Cap. 528), AI-generated works are already protected by copyright under computer-generated work provisions, while infringement liability and a text-and-data-mining exception were examined separately. [S1102] Neither document determines the legal position in Cameroon or an export market, and neither clears a specific shirt design. [S21] [S1102] Copyrightability, disclosure, training exceptions and service permissions must remain separate entries in the decision record.

The proposed rights map in **Exhibit 44**, p. 169 separates where evidence comes from and when it is reviewed across the specification workflow. Its coordinates only arrange the questions for reading. They are not legal scores, risk probabilities or an assurance that any question has been resolved.

EXHIBIT 44 / ILLUSTRATIVE SCHEMATIC POSITIONS, NOT SCORES

A model contract answers only part of the rights question



ILLUSTRATIVE [S21] [S23] [S31] GFT proposed review map. Horizontal placement moves from before processing towards intended output use. Vertical placement moves from provider documents towards contextual maker review. Coordinates are arbitrary layout positions, not measured risk. The sources support the distinctions, not the coordinates or legal clearance.

A useful approval statement should name the asset revision, permitted activity, approving role, territorial restrictions and evidence location. GFT recommends distinct workflow states for permission requested, evidence received, reviewed, approved with conditions, refused and expired. An unanswered request remains pending and blocks downstream release. The workflow should let a designer continue with independently authorised alternatives while holding the affected asset. This gives uncertainty a visible consequence without treating every incomplete record as a reason to halt unrelated styles.

Procurement should also distinguish the right to use an output from any contractual defence offered by a software supplier. GFT proposes recording the actual contract clause, eligible model version, account tier, exclusions, required user filters and notification procedure for raising a concern. [S31] A marketing summary should be filed as a claim awaiting confirmation where the contract has not been inspected. No indemnity, exclusive ownership or training-data warranty is assumed in this chapter. The reviewable result is an identified agreement and a bounded decision, rather than a general declaration that a tool is commercially safe.

11.2

Human contribution should be recorded without promising protection

GFT proposes keeping a versioned decision history that explains what a person actually created, selected, modified and approved. Useful evidence includes the initial hand sketch, an annotated design alternative, a revised surface repeat tile, a manually edited pattern panel and the accepted construction specification. The record should preserve the links between these objects across revisions. An export timestamp alone cannot explain whether the technician redrew a curve, chose a parametric option or accepted a machine suggestion without modification. Those actions represent different technical and creative contributions even when they produce the same final file format.

The human-contribution entry should describe visible changes in ordinary technical language. Recording that the technician redrew the two-piece collar stand curve and shifted the chest pocket placement is far more inspectable than stating that the team improved the output. The author should connect that description to before-and-after artefacts, rather than claim an unmeasured share of authorship. For the illustrative shirt, the maker can retain the chosen sketch, rejected collar options, generated program and edited CAD drawing. [S09] This is a proposed method, not a guarantee that an edit crosses a legal threshold. [S21]

A prompt log has a narrower purpose. It helps explain instructions, negative constraints and run behaviour, so GFT recommends keeping it alongside the expressive and geometric artefacts. Under the US Copyright Office's March 2023 guidance and January 2025 report, applicants must disclose AI-generated material and cannot rely on prompts alone for human authorship over machine-determined expression. [S1101] [S21] Even where a jurisdiction such as Hong Kong SAR provides statutory rules for computer-generated works, distinguishing human-authored drawings from machine

outputs remains essential for contractual disclosures and cross-border sales. [S1102] Prompt length or iteration count should never serve as a proxy for protectable authorship.

Recent research and technical standards show how discrete transformation records can be structured. TailorCoPilot pairs an agentic interface with TailorTrace, a version-controlled mechanism storing discrete pattern states and editing operations to support apprenticeship and auditability. [S15] In media provenance, Version 2.1 of the Coalition for Content Provenance and Authenticity (C2PA) specification, published in September 2024, defines a manifest structure where each time an asset is changed, existing provenance is preserved and the new change is added. [S1103] Neither mechanism settles third-party infringement, but both illustrate GFT's recommended principle: preserve each state transition as an explicit, linked record rather than overwriting files in place. [S15] [S1103]

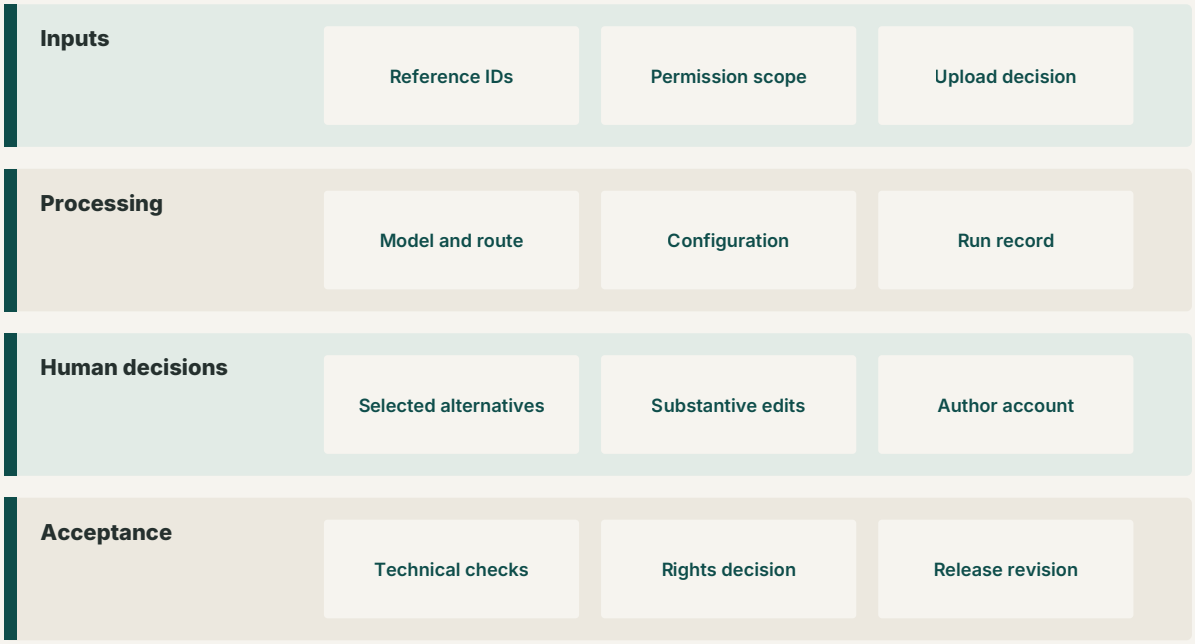
RECONCILING HUMAN EDITS, PATTERN STATES AND PROVENANCE MANIFESTS

GFT proposes a three-part check before a design revision enters rights and technical review: 1. **Input-to-candidate link:** Verify that every generated image, parametric script or JSON parameter file cites the identifier and permission state of each reference sketch, block or swatch supplied. [S09] [S10] 2. **State-diff capture:** Record human modifications as explicit diffs between raw model output and technician-edited revisions, retaining both files for inspection. [S15] [S1103] 3. **Manifest-to-pack binding:** Bind the hash or version identifier of approved visual and geometric assets into the tech pack header so any later edit invalidates approval until reviewed. [S20] [S1103]

The provenance envelope in **Exhibit 45**, p. 172 is a proposed way to connect these records without merging their meanings. It draws on the NIST generative-AI profile's attention to content provenance and lifecycle documentation. [S20] The specific layers, fields and division of responsibilities are GFT's design recommendations, not a prescribed NIST apparel system.

EXHIBIT 45 / ILLUSTRATIVE RECORD LAYERS

The design record needs a provenance envelope



ILLUSTRATIVE [S20] [S1103] GFT proposed evidence structure informed by NIST provenance guidance and C2PA manifest chaining principles. Each layer links to the same identified design revision. Retention and access must be defined separately; collecting provenance does not authorise storage of restricted material. No installed system or certification is asserted.

GFT recommends that a release record point to immutable copies or controlled versions of the accepted artefacts. A later edit should create a numbered successor revision rather than overwrite the evidence behind an earlier decision. [S1103] File digests can be used to check whether the reviewed bytes match the delivered bytes. The digest should be treated as an integrity check within the proposed system, not as evidence that the content was licensed, dimensionally accurate or authored by the person named. Those questions require their own supporting records.

The evidence envelope should preserve uncertainty in contribution. Where a reviewer cannot distinguish a manually drawn element from generated material, the record should state that ambiguity explicitly. [S1101] Where a collaborator supplies only a flattened image, the maker should request the editable source or document its absence. A decision to proceed should state which evidence was missing and why the intended activity remains acceptable under the relevant review. Silence should not become an invented account of creative work.

Access to authorship evidence needs boundaries. GFT recommends a concise review copy identifying the design decisions, with restricted links to confidential source files. Not every downstream manufacturer needs access to rejected buyer references or private correspondence. The aim is to let an authorised reviewer reconstruct the decision with sufficient evidence. Retaining every intermediate file indefinitely is not assumed; the retention rule must reflect purpose, permission and approved obligations.

11.3

Sensitive data needs a declared destination

The NIST generative-AI profile identifies data privacy and information security among the primary lifecycle risks. [S20] GFT's proposed response is to classify inputs before they enter a service and record the actual processing destination. The review boundary should include the selected model route and connected plug-ins, rather than stopping at the visible application's name. This is a procurement and workflow recommendation, not a finding that any named vendor exposes confidential data. [S20]

GFT recommends separate upload decisions for four apparel data classes: public references, proprietary pattern blocks, buyer designs and wearer body data. Public availability should trigger a rights check rather than automatic upload approval. [S21] [S23] Proprietary blocks and grading tables should be routed only through a configuration approved for confidential geometry. [S09] [S22] Buyer designs should retain their contractual restrictions. Body measurements and images should have a defined purpose, access rule and retention decision before processing. [S20] [S30] These are proposed classes for the apparel workflow, not statements of applicable privacy law.

The upload authorisation should identify the receiving account, model provider, processing region where established, retention terms, training-use terms, human-access conditions and deletion mechanism. [S20] [S31] An unknown field should be marked unknown. GFT recommends testing the documented behaviour with permitted synthetic material, such as open synthetic pattern specimens from a public dataset, before authorising sensitive inputs. [S30] The record should distinguish a policy claim, an observed test result and a contractual commitment. A test confirming that a visible history can be deleted should not be recorded as proof that all provider-held copies have disappeared.

HAIGEN's revised research abstract (first submitted in August 2024 and read in its July 2026 revision) describes a four-module system combining cloud inspiration generation with local sketch-library, recommendation and colouring modules, noting a privacy rationale for that split. [S07] It is a research example of separating tasks by destination. This chapter has not audited its code, network traffic or storage behaviour, and its current abstract reflects later revisions rather than the original 2024 submission. [S07] GFT should not treat the authors' privacy claim as an independently verified guarantee.

For the shirt workflow, GFT proposes sending only an approved generic description to an external ideation service while keeping a restricted base block and physical fabric records within the maker's authorised environment. [S07] [S22] [S32] Any generated concept or parameter file returns as an unapproved candidate that must pass validation before opening in CAD. [S09] [S10] This arrangement should itself be tested: an assistant that retrieves the block or serialises coordinates through an integration changes the boundary even if the user never attaches a file manually. [S10] [S20] The control should describe permitted flows and check them in the intended configuration, rather than rely on a label such as "local".

A provenance system can create its own concentration of sensitive material. GFT recommends storing identifiers, hashes and permission decisions in the ordinary design record while restricting access to body images, contracts and buyer files. [S20]

[S1103] Logs should omit unnecessary source content and should never become a second uncontrolled archive. The relevant reviewer needs to know where evidence can be inspected and under whose authority; other users need only the release state and conditions.

> **GFT NOTE**
>
> Approve the upload route before uploading the design. If retention, training use or onward processing cannot be established for the intended model and account, use an authorised alternative or retain the work locally. Record the unresolved field and its owner. A generic vendor assurance should not become an approval for sensitive garment data.

GFT also recommends an exit check before adoption. The maker should establish which approved artefacts, pattern geometries, material records and account histories can be exported in standard formats such as DXF, which formats remain usable without the service, and what deletion evidence the provider can supply. [S22] [S31] Restricted information should not be copied into personal accounts to bypass an incomplete procurement review. When access ends, the design record should retain the decision trail allowed by the relevant retention rule and identify anything that could not be independently verified.

11.4

Community engagement cannot be replaced by similarity screening

The World Intellectual Property Organization (WIPO) describes traditional cultural expressions as integral to the cultural and social identities of indigenous and local communities, embodying know-how and skills and transmitting core values and beliefs. [S24] Its fashion guidance promotes dialogue between Indigenous Peoples and fashion companies to shape fair and equitable interaction and warns against harmful uses without consent. [S23] These institutional explanations support careful engagement; they do not identify a representative authorised to approve a particular Cameroonian motif, such as Grassfields ndop resist-dyed geometry. The fashion page is dated 2 November 2023; the traditional-cultural-expressions page displays no publication date. Both were reopened on 30 October 2026; neither is a new statute. [S23] [S24]

GFT proposes a cultural-reference record whenever the design brief intentionally draws on an identifiable tradition or community expression. It should describe the reference, its source, the intended garment application and the proposed public account of its origin. The maker should then establish who should be consulted and what evidence supports that identification. [S23] [S24] A photograph credit, a retailer's description or a model's fluent summary should not be substituted for that investigation. Where authority remains contested or unclear, the unresolved position should be visible to the release reviewer.

Similarity screening has a limited role in this proposed workflow. It can be used as a search aid to identify visual references for further review. [S20] GFT should not assign it authority to decide cultural permission. A low similarity result should not close an engagement question, and a high similarity result should prompt examination of the actual expression, context and intended use. [S23] [S24] The review should preserve the search scope and limitations so that an incomplete search is not later described as comprehensive clearance.

For a culturally referenced shirt surface, GFT recommends linking any consultation record directly to the specific artwork revision and placement map shown during dialogue. [S23] [S1103] A general discussion about a collection should not automatically approve every generated variation or colourway. The record should distinguish permission to study a reference, approval to adapt it, approval to manufacture, and agreement about public attribution. [S23] Proposed benefit arrangements, restrictions and review opportunities should be retained where relevant. These are recommended questions for a responsible process, not invented terms attributed to any community.

GFT's proposed engagement record should also make room for refusal, limited permission and a request for further discussion. A design programme should not record only favourable outcomes or treat participation in a meeting as endorsement. The responsible role should document the response accurately and identify which changes (such as altering motif proportions, crossing seam joins or changing garment category) require renewed contact. [S23] [S24] Reusing an approval requires matching its scope.

The same record should govern the language accompanying the garment. GFT recommends reviewing claims such as authentic, traditional or community-approved against the actual agreement and evidence. [S20] [S23] A designer's interest in an expression is not evidence of provenance for the product. The maker should preserve the distinction between inspiration, authorised collaboration and independently verified origin. If no adequate account can be supported, the public description should be revised or the cultural reference removed through the documented review process.

Community engagement also needs appropriate information handling. The proposed record should capture the decision and authorised terms without placing restricted stories, images or correspondence in an unrestricted asset library. [S20] [S24] GFT recommends agreeing what may be retained and shown to technical partners, then carrying those limits into subsequent generation and publishing workflows. A later supplier should receive only the material and permissions required for its role.

11.5

Model changes require a regression decision

NIST's Artificial Intelligence Risk Management Framework (AI RMF 1.0) notes that while the framework can be used to prioritise risk, it does not prescribe risk tolerance, and calls for documented evaluation under conditions similar to intended deployment. [S19] Its generative-AI profile cautions against extrapolating capabilities from narrow assessments and supports lifecycle monitoring. [S20] GFT proposes applying that

principle at the boundary between design intent and accepted specification. Neither document sets an apparel pass rate or certifies the maker's workflow; local acceptance conditions must be defined before testing. [S19] [S20]

In generative apparel pipelines, model updates can affect multimodal language models, parameter estimators, parametric pattern programs or stitch-prediction networks. [S09] [S10] [S12] The proposed version record should identify the model, provider route, available version identifier, prompt template, configuration, reference set and export tools. [S20] [S22] When a provider does not expose an immutable version, the record should state that limitation and retain the execution timestamp and disclosed service identity. Reproducing the exact pixels is not assumed; the maker should retain the accepted output and test whether the revised workflow still meets the defined conditions for its task.

Research benchmarks illustrate why regression testing must check whole-specification outcomes alongside component metrics. In geometric stitch prediction, AutoSew reports a 96 percent F1-score on edge relations alongside a 73.3 percent error-free assembly rate on test garments. [S12] Those two figures use different denominators: high edge-level accuracy can coexist with failed garment assembly when a critical seam pair fails. [S12] GFT therefore recommends a retained evaluation set containing authorised reference blocks, known failure cases and held-out variations that test the complete handoff from permissions to CAD readback. [S19] [S22] [S30] Test cases should be labelled by what they challenge, including instruction compliance, missing material data, seam topology and DXF unit scaling. [S12] [S22] [S32] Cases used repeatedly to adjust prompts should be identified as development cases, while held-out cases remain separate until the designated evaluation.

WORKED EXAMPLE: AN IMPROVED TOTAL CAN CONCEAL A BLOCKING REGRESSION

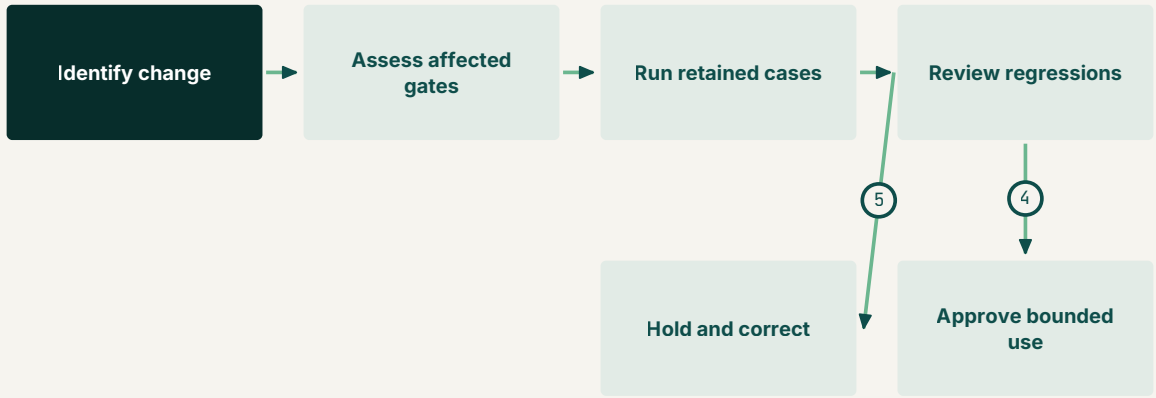
An illustrative shirt-workflow suite contains 24 cases. Assume the incumbent passes 22 and the candidate passes 23. The assumed paired results are 21 cases passed by both, 1 passed only by the incumbent, and 2 passed only by the candidate. Thus the totals are 21 + 1 = 22 and 21 + 2 = 23, with 21 + 1 + 2 = 24 distinct cases. The illustrative pass rates are 22 / 24 = 91.7 percent and 23 / 24 = 95.8 percent, rounded to a decimal place.

Suppose the 24 illustrative cases comprise 8 mandatory governance and geometry gates and 16 advisory styling preferences. The incumbent passes all 8 mandatory gates (8 / 8 = 100.0 percent) and 14 advisory cases (14 / 16 = 87.5 percent). The candidate passes all 16 advisory cases (16 / 16 = 100.0 percent) but fails 1 mandatory gate (7 / 8 = 87.5 percent) by routing a restricted block to an unapproved endpoint. Under a proposed hard condition requiring 100.0 percent on mandatory gates, the candidate remains on hold despite its 4.1 percentage-point higher aggregate pass rate. The result measures this constructed suite only. It is not an observed vendor result, a probability of safe operation or a recommendation to use this suite size.

The proposed decision path in **Exhibit 46**, p. 177 makes that distinction operational. A change opens an impact review; the tests selected depend on the changed boundary. The record should connect findings to the candidate configuration and show why the release authority accepted, restricted or rejected its use. [S19] [S20]

EXHIBIT 46 / ILLUSTRATIVE DECISION SEQUENCE

A model update should reopen the relevant test gate



4. Conditions met
5. Blocking issue

ILLUSTRATIVE [S19] [S20] GFT proposed change-control process informed by NIST evaluation guidance. A hold preserves the last authorised workflow only if it remains available and suitable. A rollback decision cannot retract an earlier disclosure. No benchmark result or deployment is claimed.

GFT recommends preserving the old and new findings side by side, including unresolved cases and reviewer disagreement. [S15] [S20] A release decision should state the permitted product envelope and any extra human checks. If the supplier removes the previous version, rollback may be unavailable; the operational alternative should then be an authorised manual CAD workflow or another qualified route. [S09] [S22] A changed retention policy can reopen upload permission even when garment outputs appear unchanged. [S20] [S31] Functional regression and contractual change need linked but distinct decisions.

11.6

Vendor due diligence must produce reviewable artefacts

CLO's exchange documentation lists Standard DXF, AAMA and ASTM options that include units, grading, notch handling and cutting or sewing lines, while its CONNECT fabric documentation requires entering weight and thickness alongside visual maps. [S22] [S32] These pages establish documented interface fields, not a verified exchange between the maker's chosen software versions or laboratory-validated fabric physics. GFT therefore recommends a sample export, material metadata check and receiving-system inspection as procurement evidence. [S22] [S32] This chapter defines what evidence procurement should retain and who must inspect it.

The proposed procurement pack in **Exhibit 47**, p. 178 converts broad assurances into inspectable requests. Its reviewer labels are responsibilities, not GFT appointments or headcount. Its hold conditions are proposed gates for the intended workflow, not allegations about any supplier.

EXHIBIT 47 / ILLUSTRATIVE EVIDENCE REQUESTS AND DECISIONS

Procurement should demand an inspectable evidence pack

Evidence request	Review owner	Hold condition
Model-specific terms	Procurement and legal	Use scope unresolved
Data-policy record	Data custodian	Destination unknown
Sample editable export	Pattern technician	Readback fails
Material metadata schema	Quality and sourcing	Basis undisclosed
System requirements	Technical owner	Environment untested
Support boundaries	Workflow owner	Failure path absent
Exit and deletion terms	Procurement and data	Recovery unproven

ILLUSTRATIVE [S22] [S31] [S32] GFT proposed procurement checklist. Actual contracts, policies, software versions and sample files must be obtained and reviewed at adoption. No current price comparison, vendor score or accepted supplier is established.

GFT recommends retaining the supplier's response in its original dated form (including exported DXF files, contract schedules and policy snapshots) alongside the maker's assessment. [S20] [S22] A demonstration recording should identify the configuration and permitted sample material. An answer that references a beta feature, a different model or another account class should remain pending for the intended use. [S31] Procurement should distinguish evidence received from evidence inspected and evidence accepted. A complete folder is not an accepted workflow if nobody with the relevant responsibility has reviewed its contents.

The technical inspection should ask whether the receiving team can open the sample, identify units, recover required annotations and edit the intended objects, while confirming that material records separate measured values from estimates. [S22] [S32] GFT proposes recording discrepancies against the supplier's claimed capability and the maker's acceptance conditions. A visually convincing three-dimensional preview should not close a two-dimensional DXF editability requirement. [S10] [S22] If a proprietary format is used during generation, the exit review should establish what usable industrial exports remain when access ends and whether those files have actually been reopened. [S22]

Commercial review should examine the relationship between policies, the agreed contract and operational settings. GFT recommends asking how model substitutions are disclosed, which support activities may expose uploaded material, and how an incident should be raised. [S20] [S31] Pricing remains unresearched here and requires

formal quotes before Chapter 12 and Chapter 13 pilot evaluations. The procurement record should not imply that a trial account's features or permissions apply to a future production arrangement. [S31]

Finally, the maker should assign an owner to each pending request and connect it to a decision consequence. Missing export evidence may block technical qualification while allowing a limited ideation exercise with authorised public material. [S22] [S31] Unknown upload terms should block sensitive uploads regardless of a successful export demonstration. [S20] Recording those boundaries explicitly ensures that enthusiasm for a demonstration cannot expand the permission granted.

11.7

Incidents should lead to containment and correction

GFT proposes a response record for a rights concern, an unintended confidential upload, a corrupted pattern export, a misleading material statement or an escaped specification error. [S20] [S22] [S23] The NIST generative-AI profile includes lifecycle monitoring, incident-related actions and information-security considerations. [S20] The response design here is GFT's apparel-specific proposal; it is not a certification, regulatory acceptance or claim that such an incident has occurred in GFT operations.

The first record should preserve the reported concern, time identified, reporting route, affected style and revision identifiers, and current release state. A report should trigger an immediate trace of the provenance envelope without converting an initial allegation into an established legal finding. [S20] [S1103] GFT recommends distinguishing suspected cause, confirmed observations, containment action and unresolved scope. The responsible role should identify the relevant generated outputs, human-edited descendants, specification revisions and external recipients. [S15] [S1103] Where the dependency trail is incomplete, that limitation should inform the hold rather than disappear from the report.

> **UNVERSIONED OVERWRITES DESTROY INCIDENT CONTAINMENT**
>
> If designers or pattern technicians overwrite shared pattern files, print tiles or tech pack tables in place without incrementing revision identifiers, an incident cannot be bounded to a specific batch. [S15] [S1103] When a rights objection, unit-scale defect or unverified fabric preset is found in an unversioned folder, the maker cannot prove which cutting orders used the defective state and which used an earlier valid state. GFT recommends enforcing append-only revision histories and hash-linked release packets before any generative tool touches shared product folders, so containment targets exact revisions. [S20] [S1103]

Containment should match the event. The proposed workflow suspends the affected upload route, holds the disputed artwork, quarantines the corrupted DXF export or stops use of the affected specification as appropriate. [S21] [S22] [S23] It identifies which unaffected work can continue and why. For an unintended external disclosure, the team must establish the destination and contractual remediation process; restor-

ing an earlier local model version cannot undo an external transmission. [S20] [S31] Required notifications and legal duties need project-specific authority rather than assumptions here.

Correction must produce a new identifiable revision. [S15] [S1103] GFT recommends retaining the original evidence in a restricted archive under approved retention rules, recording the corrective change and identifying the test needed to reopen the gate. Replacing a file in place leaves downstream users unable to distinguish superseded instructions from current ones. An export or seam defect requires clean receiving-system readback of the new revision, while a rights concern requires documented resolution or an independently created replacement rather than superficial re-prompting. [S12] [S21] [S22] [S23]

Learning should then return to the workflow. The maker should decide whether to add a permanent regression case, change an upload class, revise a supplier condition or refine technician training from Chapter 10. [S15] [S19] [S20] Repeated exceptions should be tracked in GFT's proposed `risk-matrix` tool so recurring control failures remain visible across styles. A small team can allocate these duties across existing roles, provided the record separates who proposes a correction from who accepts its release.

The chapter's boundary remains the accepted design and specification. Broader model operations belong to companion paper WP14 and applicable legal duties belong to companion paper WP17, connected through the same revision identifiers. [S19] [S20] Within WP10, these versioned records form a prerequisite for Chapter 12's economic model and Chapter 13's pilot qualification. GFT's recommendation is practical: make the permission decision, processing route and technical release reconstructable before expanding generated design work.

A design is ready when its permissions and its production instructions describe the same revision.

CHAPTER SOURCE NOTES

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

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

S20 National Institute of Standards and Technology. 2024-07-26.

S31 Adobe. 2026-05-13.

S1101 U.S. Copyright Office. 2023-03-16.

S21 U.S. Copyright Office. 2025-01-29.

S1102 Commerce and Economic Development Bureau, Hong Kong Special Administrative Region Government. 2024-07-08.

S23 World Intellectual Property Organization. 2023-11-02.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

S1103 Coalition for Content Provenance and Authenticity. 2024-09.

S10 Authors via arXiv; CVPR 2025. 2024-12-23.

S22 CLO Virtual Fashion. 2026-07-15.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S07 Authors via arXiv. 2024-08-01.

S32 CLO Virtual Fashion, CONNECT Support. 2025-04-17.

S24 World Intellectual Property Organization. n.d..

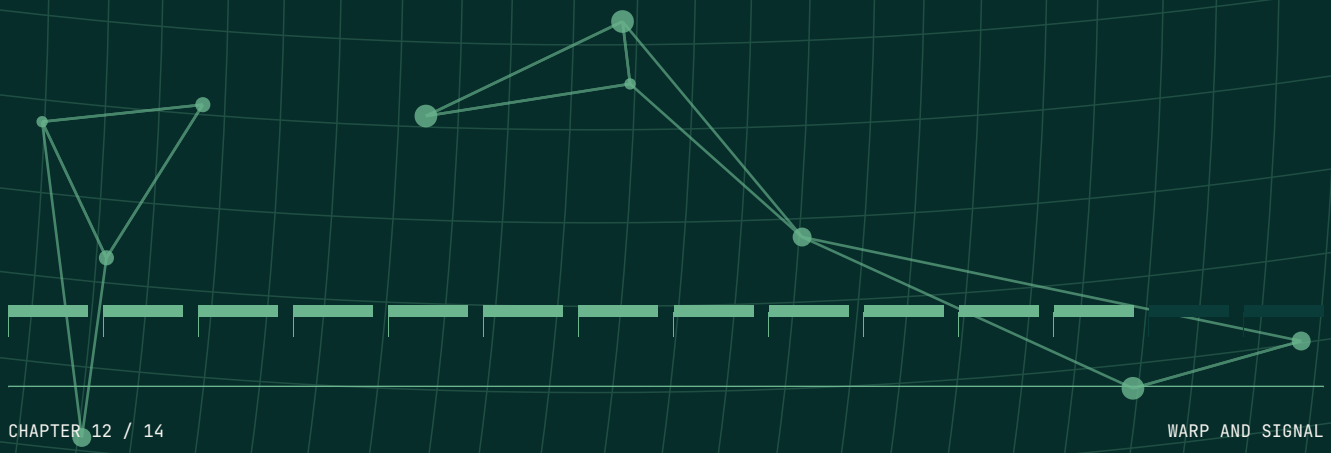
S19 National Institute of Standards and Technology. 2023-01-26.

S12 Authors via arXiv; WACV 2026. 2026-02-25.

12

Economics must count accepted styles and retained physical checks

Savings exist only where the cost of review and correction is lower than the work genuinely avoided.



CHAPTER 12
DECISION TEST

Savings exist only where the cost of review and correction is lower than the work genuinely avoided.

GFT should evaluate generative design by the full cost of each accepted style, not by images produced, prompts completed or time saved on an isolated task. An accepted style has a usable specification, the required physical approvals and an authorised release state. The assisted workflow must reach that same state as the manual alternative. Its economic case is the cost genuinely avoided, less software, integration, review, correction, retained sampling and every other incremental cost. This boundary belongs inside GFT's Cameroon textile programme as a proposed measurement discipline. It is not a report of savings already achieved. No current vendor price, local hourly rate or sample-avoidance rate has been established for this chapter. The numerical cases below are therefore illustrative, not forecasts. NIST's risk framework leaves risk tolerance to the organisation, while its generative profile calls for contextual evaluation across the lifecycle. [S19] [S20] GFT should likewise set acceptance and economic thresholds before seeing pilot results. This chapter is not investment, legal or tax advice.

12.1

The counterfactual must describe the same accepted product

The comparison begins with two complete routes to the same product outcome. The manual route and the assisted route should receive the same brief, permitted base block, material information, size range, quality requirements and buyer conditions. Both must produce editable pattern geometry, construction instructions, a bill of materials, fit evidence, rights records and the required physical samples. Both must pass receiving-CAD readback and obtain release from the accountable role. CLO's documentation shows that exchange settings can include units, grading, notches, seam allowances and cutting or sewing lines, so a file that merely opens is not an adequate common endpoint. [S22] Acceptance must mean that the receiving team can use the specification, not that the generating team has exported it.

A generated image can be an intermediate artefact in either route, but it is not the output being valued. The same applies to a syntactically valid parameter file, a visually plausible simulation or an unreviewed tech pack. Each may reduce some work while creating other work. Counting them as finished styles would move the denominator to an easier task and conceal the remaining technical effort. GFT's proposed unit is narrower: one style whose specified release conditions are complete at an identified revision. A style accepted with conditions should retain those conditions in the record. It should not be grouped silently with unconditional approvals.

The counterfactual also needs matched complexity. A collar variation on an approved cotton-shirt block is not equivalent to a new multi-panel garment with unfamiliar construction. Fabric behaviour, size coverage, print requirements and available technician expertise can change the work substantially. GFT should pair cases by product family and record the features that make each case difficult. Where exact pairing is impractical, the evaluation should stratify the results and disclose the differences rather than present one average as if every style required the same effort. NIST's emphasis on contextual measurement supports this discipline, although it does not prescribe an apparel test design. [S19]

Both routes must include failure. A rejected generated candidate consumes generation and review effort even though it never enters the accepted-style denominator. A manual concept abandoned after technical review also consumes effort. The evaluation should retain both, with reason codes that distinguish creative rejection, invalid geometry, material mismatch, fit failure, rights holds and exchange failure. Excluding failed runs would overstate assisted performance. Excluding manual dead ends would make the baseline unfairly expensive. The honest comparison includes all work required to obtain the accepted set.

Physical checks require the same rule. A sample round can be counted as avoided only when the responsible acceptance plan no longer requires it and the final approval remains equivalent. A digital fit image cannot erase a physical fit requirement by assertion. Nor can a cheaper digital iteration be credited with avoiding fabric, labour and freight if the physical sample was still made. Retained samples belong in both the activity record and the cost model. Their presence does not make assistance worthless. It simply prevents digital work from being mistaken for physical work that did not occur.

> **DEFINE THE COMPARISON BEFORE COLLECTING TIME RECORDS**
>
> GFT should freeze a comparison sheet for each style family. Name the acceptance state, applicable physical checks, receiving-CAD procedure and responsible approvers. Assign the same requirements to manual and assisted cases. Record every proposed candidate, including rejects, and link each accepted style to its final revision. Report cost per accepted style, physical rounds per accepted style and correction hours per accepted style together. A change to any denominator requires a new analysis note, not a silent edit of the historical result.

The resulting counterfactual is deliberately conservative. It does not ask whether assistance can produce more visual options. It asks whether the organisation can obtain the same accepted outcome with fewer total resources, while preserving independent release authority. That question remains answerable even when some benefits are only potential. Unpriced learning, reusable records and additional design choice can be reported separately without being converted into invented cash savings.

12.2

A complete cost model includes rejected outputs

A usable model separates costs by behaviour rather than by accounting label alone. Some costs recur with every candidate, such as generation, initial screening and failed-run handling. Some attach to an accepted style, such as final CAD correction, material review and physical approval. Others are periodic, such as subscriptions, integration maintenance, model regression testing and training. One-off implementation work includes process design, data preparation, permission review, export tests and connection to the design record. The classification matters because higher accepted volume may dilute a periodic cost while doing nothing to reduce review effort on each rejected candidate.

GFT should start with an incremental model. Include costs that change because the assisted route is introduced and costs that disappear because it is introduced. Existing pattern work that continues unchanged is not a new benefit or a new charge. A technician's hours should be counted when they move between activities, using observed time and an approved cost basis. Released capacity is not automatically a cash saving. It becomes one only when the organisation actually avoids expenditure or uses the capacity for valuable work that would otherwise require additional resources. Until then, report hours released and cash avoided as different outcomes.

The physical sample cost needs its own bill of activities. Depending on the actual workflow, it may include fabric and trims, cutting and sewing labour, sample-room overhead, finishing, measurement, courier services, travel, duties where applicable and disposal. This chapter supplies no rates for those items. Finance and product-development teams should agree which components are incremental and which are already absorbed in existing operations. A single cost per round is convenient for scenario analysis only after its inclusions and price date are stated. Otherwise, an apparently precise avoided-cost total may conceal inconsistent boundaries between baseline and assisted cases.

Digital costs deserve equal care. A subscription quote may omit usage charges, integration support, additional storage or the effort required to preserve approved artefacts. Export and readback functions may need configuration and testing in the intended CAD environment. [S22] Model or provider changes can require another evaluation cycle, while failures can create correction work outside the generating interface. [S20] Rights review and data-classification work also belong in the model when they are incremental. Omitting these activities would make the assisted route appear cheaper by transferring effort to uncounted teams.

Rejected outputs are not overhead to be hidden in a general rate. Their number and handling time reveal whether the workflow is improving. GFT should record the generation cost where measurable, designer screening time, technician review time and downstream correction for every candidate cohort. A cheap failed generation may still be expensive if an expert must inspect it carefully. Conversely, a relatively costly run may be useful if it reliably prevents a physical iteration. Only linked records can distinguish those cases.

The measurement register in **Exhibit 48**, p. 186 assigns each proposed measure a denominator and an evidence source. It is a collection design, not a table of observed results. Blank values should remain blank during a pilot until the underlying records exist.

EXHIBIT 48 / GFT PROPOSED MEASUREMENT REGISTER

An economic case needs measured denominators

Measure	Denominator or boundary	Required record
Accepted styles	Styles reaching frozen release state	Approval and revision ledger
Candidate acceptance	Accepted styles divided by proposed candidates	Complete candidate cohort
Review effort	Hours per proposed and accepted style	Role-level time log
Correction effort	Hours per accepted style by failure cause	Edit and reason log
Physical sampling	Rounds per accepted style	Sample and approval record
Avoided rounds	Matched baseline rounds less assisted rounds	Paired style comparison
Digital cost	Periodic, candidate and integration boundary	Quotes, usage and support records
Net annual effect	All measured avoided and incremental costs	Reconciled scenario workbook

ILLUSTRATIVE [S19] [S20] [S22] Illustrative collection structure. No observed values are supplied. Accepted style means the frozen acceptance definition has been met. Candidate and review records must include rejected work. The sources support contextual evaluation and exchange checks, not these particular management fields.

The register prevents two common distortions. First, accepted styles and proposed candidates remain separate, so a high rejection load cannot disappear. Second, effort and physical rounds remain visible beside money, so a favourable exchange rate or a broad overhead allocation cannot create an apparent operational improvement. The table cannot establish that any benefit exists. It only makes the eventual benefit claim inspectable and shows which missing records currently block it.

> **GFT NOTE**
>
> Do not approve a commercial case from a subscription quote and an assumed number of avoided samples. Require a matched baseline, a complete candidate cohort, technician time records, an explicit physical-sample cost boundary and current integration costs. Report missing values as missing. Where a physical approval is retained, charge it to the assisted route and never count it as avoided work.

Finance should reconcile the model to actual expenditure after the operational review. The technical team can explain why rounds changed, but it should not mark its own assumed savings as realised cash. The finance reviewer should check price dates, currency treatment and whether expenditure was genuinely removed. A model expressed in CFA francs (XAF) should retain that currency through the calculation rather than introduce an unnecessary conversion. Where another currency is required for a supplier quote, the conversion basis and date become additional assumptions to disclose.

12.3

A transparent example can test the arithmetic without predicting savings

A small scenario shows the method without claiming market knowledge. Suppose, for an entirely illustrative case, that a maker accepts 40 styles a year. The manual route requires 3 physical sample rounds per accepted style and the assisted route requires 2. Assume an all-in incremental cost of XAF 30,000 per physical round and XAF 900,000 a year for incremental assistance. The assistance figure is a hypothetical bundle, not a vendor quote. It could represent software and related support, but no actual supplier or price is implied. All other costs and final acceptance requirements are held unchanged in this first case.

In this illustrative case, baseline physical sampling expenditure is 40 styles multiplied by 3 rounds multiplied by XAF 30,000, giving XAF 3,600,000. The assisted physical sampling expenditure is 40 multiplied by 2 multiplied by XAF 30,000, giving XAF 2,400,000. The difference is XAF 1,200,000. Expressed as activity, the same calculation gives 40 multiplied by the difference between 3 and 2 rounds, or 40 avoided rounds. These are constructed inputs and GFT arithmetic, not an observation of a factory or a recommended sampling level.

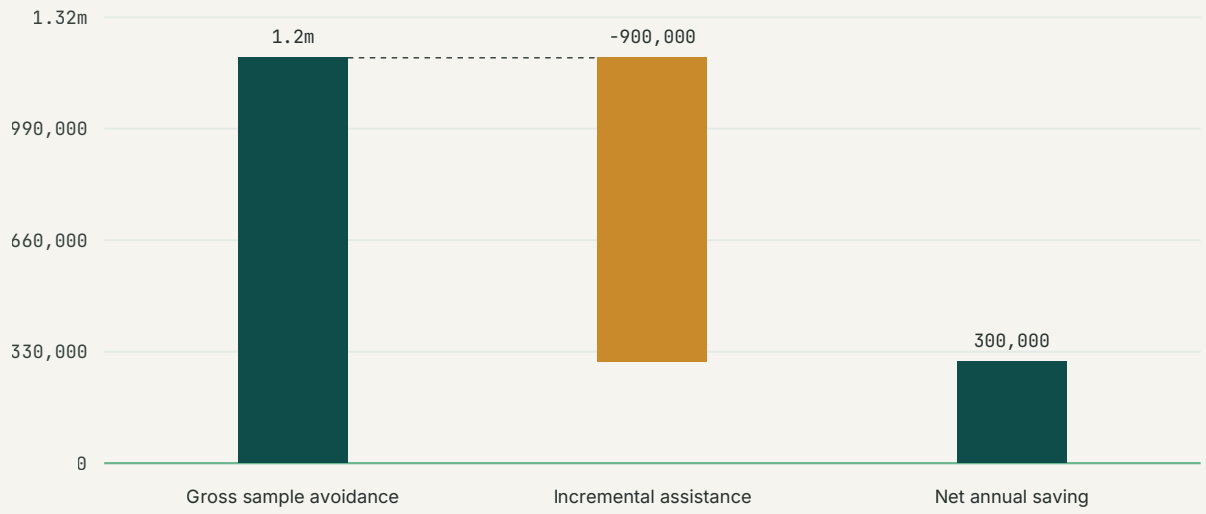
WORKED EXAMPLE: NET ANNUAL EFFECT AFTER INCREMENTAL ASSISTANCE COST

Use only the illustrative inputs above. Avoided rounds equal 40 x (3 - 2) = 40 rounds. Gross avoidance equals 40 rounds x XAF 30,000 = XAF 1,200,000 a year. Add the hypothetical XAF 900,000 incremental assistance cost to the assisted physical sampling expenditure: XAF 2,400,000 + XAF 900,000 = XAF 3,300,000. Compare that with the illustrative baseline of XAF 3,600,000. Net annual saving is XAF 3,600,000 - XAF 3,300,000 = XAF 300,000. Equivalently, XAF 1,200,000 gross avoidance less XAF 900,000 assistance cost gives XAF 300,000. The result is valid only inside this constructed scenario. It excludes additional review, correction, integration and other changed costs, and assumes equivalent final acceptance.

The waterfall in **Exhibit 49**, p. 188 shows why gross avoidance is not the economic result. The annual assistance charge absorbs most of the constructed gross benefit. This presentation also keeps the unit constant: every bar is XAF per year, not a mixture of sample rounds, percentages and money.

EXHIBIT 49 / XAF PER YEAR

Illustrative net saving is smaller than gross sample avoidance



ILLUSTRATIVE Fully illustrative scenario, not a forecast or market rate. Inputs: 40 accepted styles, 3 baseline physical rounds and 2 assisted rounds per style, XAF 30,000 per round, and XAF 900,000 annual incremental assistance cost. Gross avoidance = 40 x (3 - 2) x XAF 30,000 = XAF 1,200,000. Net = XAF 1,200,000 - XAF 900,000 = XAF 300,000. Equivalent final acceptance and unchanged other costs are assumed.

The illustrative example deliberately does not calculate a return on investment, payback period or net present value. Those measures would require an investment boundary, timing, asset life, financing assumptions and a view of costs beyond the first annual charge. None is established here. The arithmetic answers only whether this particular set of hypothetical annual flows produces a positive difference. It says nothing about whether the assumed reduction from 3 rounds to 2 is achievable, whether XAF 30,000 represents a real sample round or whether XAF 900,000 resembles a procurable service.

The acceptance condition is the most important assumption. In this illustrative scenario, if the assisted route still needs 3 physical rounds to obtain the same approval, there are no avoided rounds in the model. Gross avoidance becomes zero and the constructed annual effect is a cost of XAF 900,000 before any other changes. If it needs additional correction to reach the same release state, that correction must be added even where the round reduction remains. A positive answer under ideal assumptions is a reason to collect evidence, not a reason to scale.

A second useful check concerns the illustrative cost boundary. If XAF 30,000 covers only courier and material charges while sample-room labour remains unchanged, the avoided amount should not include that labour. If labour truly changes, the time and expenditure records should demonstrate it. The scenario's simple unit cost is therefore a placeholder for a reconciled cost build-up. It should be replaced by component-level observations rather than treated as a general apparel benchmark.

12.4

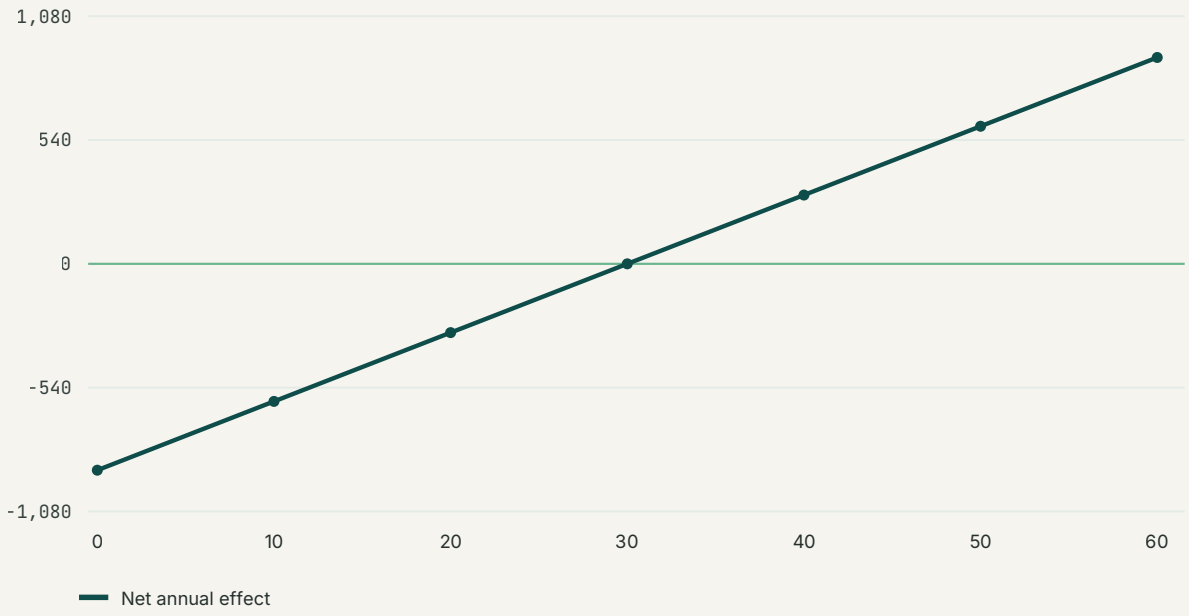
Breakeven depends on usable volume and actual avoidance

Breakeven asks how much accepted output is needed for gross avoidance to cover incremental cost. In the simple illustrative case, each accepted style avoids 1 round at XAF 30,000. Annual incremental assistance cost is XAF 900,000. The contribution per accepted style is therefore XAF 30,000, and breakeven volume is XAF 900,000 divided by XAF 30,000, which equals 30 accepted styles a year. Below that volume, the constructed net annual effect is negative. Above it, the effect is positive only while every original assumption continues to hold.

The line in **Exhibit 50**, p. 189 makes that relationship visible. It is not a demand forecast. The horizontal axis counts styles that reach the same accepted state, not concepts generated or orders hoped for. The vertical axis reports net annual effect after the hypothetical assistance charge, with all other changed costs still excluded.

EXHIBIT 50 / THOUSAND XAF NET ANNUAL EFFECT

Illustrative breakeven requires enough accepted styles



ILLUSTRATIVE Illustrative arithmetic from hypothetical inputs. Net annual effect = N accepted styles x 1 avoided round x XAF 30,000 - XAF 900,000, divided by 1,000 for display. N ranges from 0 to 60 accepted styles. Breakeven is N = 30. Equivalent acceptance, constant unit cost and no other changed costs are assumed. This is not a forecast of volume or savings.

The zero-avoidance case breaks the illustrative line's premise rather than merely shifting it. If assisted and manual routes both require the same physical rounds, gross avoidance is zero at every volume. More accepted styles then do not cure the XAF 900,000 incremental annual charge through sampling savings. A commercial case that depends on scale should therefore identify which variable scales: physical avoidance, review effort, periodic charges or something else. Volume alone is not a universal remedy.

Additional review cost changes the threshold directly. Suppose, as another illustrative variation, that assisted work adds XAF 300,000 of annual review and correction cost beyond the original XAF 900,000. Total incremental annual cost becomes XAF 1,200,000. With the same XAF 30,000 contribution per accepted style, breakeven becomes XAF 1,200,000 divided by XAF 30,000, or 40 accepted styles. At the original 40-style volume, the earlier XAF 300,000 net saving is fully absorbed. The example shows why review cannot be left as a qualitative footnote.

A per-style review charge has a slightly different effect. If the illustrative assisted route adds XAF 10,000 of correction for every accepted style, net contribution falls from XAF 30,000 to XAF 20,000. Breakeven then becomes XAF 900,000 divided by XAF 20,000, or 45 accepted styles. This is still a constructed case. In practice, correction may attach to rejected candidates as well as accepted styles, so a complete model needs both candidate volume and acceptance yield.

That distinction matters when many proposals fail. Assume, for an illustrative case, that 50 proposed candidates produce 40 accepted styles. All 50 may incur generation and screening costs, while only 40 enter the accepted-style denominator. The candidate acceptance ratio is 40 divided by 50, or 80 percent in this constructed example. An evaluation that prices only the 40 successful cases would omit the work attached to the other 10. The proper allocation depends on when costs occur, which is why the register separates candidate-level and accepted-style records.

Breakeven should therefore be presented as a family of conditional cases. Vary accepted volume, avoided rounds, cost per genuinely avoided round, candidate acceptance, review effort and annual digital cost. Hold quality requirements fixed within each comparison. Do not choose a favourable combination after seeing pilot outcomes and call it the base case. GFT should preserve the original assumptions, record observed replacements and show the effect of each change. This turns breakeven into a diagnostic tool rather than a sales argument.

12.5

A faster task is not automatically a shorter calendar

Task time and elapsed lead time answer different questions. A technician may complete a pattern edit more quickly while the style still waits for a fabric decision, a buyer response, a sample-room slot or physical delivery. GFT should therefore record touch time, queue time, material waiting, approval waiting and rework separately. Each needs a start event, an end event and an owner. The primary calendar measure should run from the agreed brief state to the authorised specification state, with pauses and their causes visible. A second measure can isolate active technician work. Publishing only the second would make a local efficiency look like an end-to-end improvement.

TailorCoPilot's research abstract describes a user study with novices and advanced novices and presents the system as a route for capturing practice-based expertise and supporting apprenticeship. [S15] That research context is not a factory lead-time result, a staffing reduction or proof that an assisted Cameroonian apparel workflow would finish sooner. A local pilot may measure whether task time changes, but no

research time effect is imported into this chapter's model. GFT must retain its own task definitions, product mix and acceptance conditions rather than multiply an external study result by an assumed annual style volume.

Queue behaviour also changes the interpretation of released hours. If assistance moves work earlier into a queue that is already full, total calendar time may not improve. If it reduces a genuine bottleneck and downstream resources are available, the accepted style may finish earlier. Even then, earlier completion has no automatic cash value. The maker should identify the realised consequence, such as avoided expediting or additional accepted work performed with existing capacity. That consequence must not also be counted as labour saved if the same hours support both claims.

Rework needs a separate clock. A fast first output followed by repeated correction may have more total touch time than a slower manual route. Record time by cause and revision, including work performed by people outside the design team. The record should also distinguish learning during a pilot from steady operation, without deleting early failures. A later improvement can be shown as a trend only after the case definitions remain stable.

> **RECONCILE TIME BEFORE ASSIGNING IT VALUE**
>
> For every matched style, retain active work by role, waiting by cause, physical-sample transit and approval time. Report the full distribution of cases rather than only the best assisted example. Finance should value time only when an actual expenditure or productive-capacity consequence is identified. Otherwise, report hours and elapsed days without monetising them. This keeps operational learning useful even when the cash effect is unresolved.

12.6

Environmental claims need boundaries and physical data

A lower sampling bill does not establish a lower environmental impact. Money can fall because a price changes while material use stays constant, and an avoided sample round can coincide with additional computing or extra design variants. The European Commission's textiles strategy explicitly considers the entire lifecycle of textile products. [S25] The GHG Protocol Product Standard likewise addresses full lifecycle emissions and reduction opportunities. [S29] These sources support a broader accounting boundary. They do not show that generative design reduces impacts in this workflow.

The first requirement is a common functional unit. For a development comparison, GFT can use one accepted style meeting the same technical and physical requirements. That unit should not be confused with the lifecycle of the garments later produced under the style. A sample-development account may cover materials, sample-room energy, transport, computing and disposal associated with reaching acceptance. A full product account would extend to raw materials, manufacturing, distribution, use and end of life. Claims must name which boundary they use and avoid transferring a development-stage result to the whole garment.

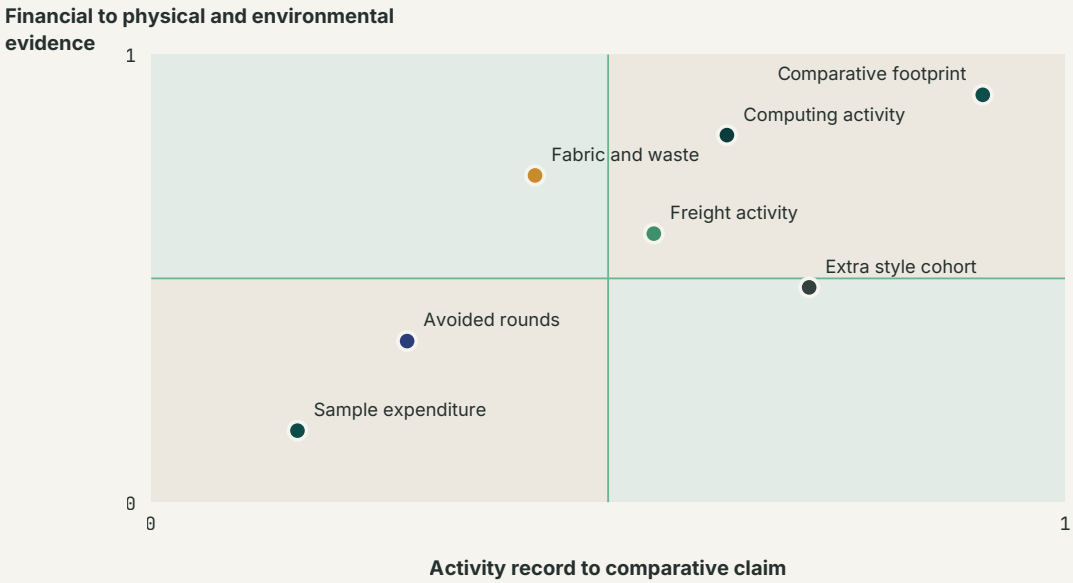
Physical activity data come before emissions estimates. For each matched style, record fabric and trim quantities, sample rounds, waste destinations where known, transport legs, energy or fuel data where available, and the digital workloads that changed. The record should distinguish measured quantities from allocations and assumptions. Computing effects may require provider or infrastructure data that are not available. NIST's generative profile identifies environmental impacts associated with resource-intensive training and inference, while also noting variation by task and configuration. [S20] That uncertainty is a reason to disclose the gap, not to assign an invented factor.

Rebound belongs in the same analysis. Assistance may increase the number of concepts reviewed, colourways sampled or styles accepted. Those additional activities can consume material and computing even if each individual iteration becomes cheaper. Conversely, better selection may reduce unnecessary development. Neither direction should be assumed. Compare the complete cohort and state whether the analysis covers a fixed accepted-style set or an expanded development programme.

The evidence map in **Exhibit 51**, p. 192 separates financial records from physical and environmental evidence. Its coordinates arrange the questions for reading. They are not scores, measurements or proof that the necessary data exist.

EXHIBIT 51 / ILLUSTRATIVE SCHEMATIC POSITIONS, NOT SCORES

Environmental and financial claims need separate evidence



ILLUSTRATIVE [S20] [S25] [S29] GFT proposed evidence map. Horizontal positions move from activity records towards comparative claims. Vertical positions move from financial records towards physical and environmental evidence. Coordinates are arbitrary layout positions. Sources support lifecycle and computing-boundary distinctions, not the positions or any quantified benefit. No emissions factors or measured savings are supplied.

The map shows why sample expenditure and a comparative footprint sit far apart. The first can be reconciled to invoices. The second needs consistent lifecycle boundaries, physical activity data, suitable factors and a treatment of uncertainty. Avoided rounds are useful activity evidence, but they do not supply the missing conversion. The chapter therefore claims no carbon, water or waste reduction from its illustrative XAF

case. Product footprint calculation belongs to the companion work on carbon accounting, with product-specific rules and legal duties reviewed separately. [S25] [S29]

12.7

The investment decision should survive a downside case

GFT should test the decision against adverse but plausible operating conditions before expanding a pilot. The minimum set is no avoided physical rounds, higher correction effort, fewer accepted styles, lower candidate acceptance and a changed provider price. Each scenario should retain the same release requirements. A favourable result produced by relaxing fit, rights or CAD checks is not an economic upside. It is a different and unacceptable product boundary. NIST's framework leaves the tolerance decision to the organisation, so the accountable team must define both quality holds and financial limits before the test. [S19]

The downside review should distinguish a stop condition from a poor return. Critical geometry, permission or fit failures can hold release even when the cost model appears favourable. Positive annual arithmetic cannot compensate for an unresolved release requirement. Separately, a technically acceptable route may be rejected because accepted volume is too low or total correction effort too high. Keeping these decisions distinct prevents a single financial score from concealing the reason for action.

Procurement should replace scenario placeholders with dated quotes, usage terms and an integration scope. Receiving-CAD tests should demonstrate the actual export and reopening path before the integration cost is treated as complete. [S22] The evaluation should also retain failed runs, missing records and expert disagreement rather than reporting only successful demonstrations. [S20] If the provider changes its model or charging basis, the affected technical and economic cases should be reopened. No current price or contract concession is assumed here.

A decision workbook should preserve each scenario as a separate, dated view. It should name the observed input, the remaining assumption, the source record and the reviewer who approved any replacement. When an observation is missing, the cell should remain unresolved rather than inherit a favourable value from another case. This makes the difference between evidence and judgement visible to the next reviewer. It also allows a later pilot to update the model without rewriting the earlier baseline.

Annualisation needs the same restraint. A short pilot may reveal candidate-level effort and failure causes, but periodic charges, seasonality and product mix can change the annual result. GFT should annualise only from a stated operating plan and show which

volumes are authorised rather than merely possible. Where the evidence supports only a per-style observation, report it at that scale. A larger annual total is not more informative when it depends on an untested expansion assumption.

A pilot decision can then take one of three forms: stop, repair and retest, or expand within the measured envelope. Expansion requires both accepted technical outcomes and a downside case that the accountable team can bear. Capital structure, tax treatment and financing sit outside this chapter. The annual illustrations are not valuations or investment recommendations. Their purpose is narrower and more useful: show exactly which observations would justify further work, and which would end it.

Count the accepted garment, keep the physical evidence, and let savings follow the record.

CHAPTER SOURCE NOTES

S19 National Institute of Standards and Technology. 2023-01-26.

S20 National Institute of Standards and Technology. 2024-07-26.

S22 CLO Virtual Fashion. 2026-07-15.

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

S25 European Commission, Directorate-General for Environment. 2022-03-30.

S29 GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2026.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

13

PART IV: GFT SHOULD SCALE ONLY WHAT PASSES ITS GATES / BUY, TEST AND RETAIN CONTROL

A local pilot must test the complete specification handoff

A greenfield maker should buy proven components and qualify a bounded assisted workflow before increasing dependence on it.

CHAPTER 13
DECISION TEST

A greenfield maker should buy proven components and qualify a bounded assisted workflow before increasing dependence on it.

A greenfield maker should procure a usable specification workflow, not a collection of impressive demonstrations. GFT's Cameroon textile programme, prepared in January 2024 with source notes refreshed in 2026, provides the adoption sequence: establish controlled records and responsible technicians, measure the manual process, then qualify a bounded partner-led generative-specification pilot. The pilot must carry the same authorised brief through concept generation, editable geometry, material and fit review, receiving-CAD readback and accountable release. It must also test failure recovery and exit. A successful image or an exported file answers only a fragment of that question. This chapter defines the product envelope, matched comparison, release conditions and evidence report needed before dependence increases. Every design is a proposed GFT control, not an account of a completed pilot, an operating factory or a realised saving.

13.1

The procurement unit should be a usable workflow

Buying ideation assistance, buying CAD capability, buying material services and commissioning a research integration are different decisions. Ideation assistance may help a designer explore collar or pocket alternatives. CAD capability may create and exchange editable panels. Material services may supply physical evidence and simulation inputs. A research integration may connect these activities through generated programs or structured parameters. None of them, purchased alone, establishes that the maker can release a garment. Procurement should therefore describe the complete path from an authorised brief to an accepted specification, then identify which components already work and which require qualification.

Design2GarmentCode illustrates why the distinction matters. Its research route translates multimodal design concepts into parametric pattern-making programs, making structure available for subsequent inspection. [S09] That is a useful candidate mechanism, not a procurement outcome. The buyer still needs to know who can edit the program, how deterministic checks handle invalid geometry, whether the selected material record is measured, and whether the receiving team can reopen the result. A subscription or model name cannot answer those workflow questions.

GFT recommends a partner-led pilot only after approved base blocks, material records and responsible technicians exist. These prerequisites give the trial a stable reference. Without an approved block, the team cannot tell whether an unfamiliar panel is a permitted design variant or an error. Without a material identity, a visually plausible drape has no confirmed physical referent. Without a technician who owns correction and escalation, a failed output becomes an informal experiment rather than controlled evidence. The partner can supply methods and integration effort, but the maker must retain product and release authority.

This sequence does not imply training a foundation model from the factory's sales history. The first question is whether existing components can operate within local constraints using authorised blocks, briefs and material records. Fine-tuning, new dataset construction and bespoke model development are separate investment decisions with their own data and evaluation burdens. They should not be smuggled

into a pilot whose purpose is to test a bounded handoff. Early exploration with approved public inputs can inform procurement, but it remains outside production release.

The procurement record should name the intended use, excluded uses and acceptance evidence for each component. For generation, retain the model and integration version, input classes, output representation and failure behaviour. For CAD, retain supported formats and the exact receiving configuration. For material services, retain the physical-to-digital identity link and measurement provenance. For integration, retain execution logs, permitted operations and the manual fallback. Unknown items stay pending. They do not become assumed capabilities because a demonstration appeared convincing.

A usable workflow also includes workload. The buyer should ask who prepares inputs, reviews every candidate, corrects rejected work, operates physical checks and maintains records after the partner leaves. A component that reduces drawing effort but increases unexplained review work may not improve the accepted-style outcome. Chapter 12 supplies the economic boundary. The present procurement task is narrower: make every dependency and responsibility visible before the trial begins, so the measured comparison can include the work that the demonstration omits.

> **GFT NOTE**
>
> Write the purchase specification around an accepted handoff. Require the partner to supply editable artefacts, exact configurations, failure records and an exit route, not only a successful visual demonstration. Keep block ownership, material acceptance and final release with the maker. Defer bespoke model development until a bounded evaluation shows which limitation, if any, existing components cannot resolve.

13.2

The first pilot needs a deliberately narrow product envelope

The evaluation set should be small enough for complete technical review and varied enough to expose transfer failures. GFT proposes an illustrative planning design of 12 styles across 3 product families, with 4 styles in each family. These are deliberate teaching choices, not a statistically powered acceptance standard. The actual maker must select families for which it owns or may use the blocks, can obtain the materials, has responsible reviewers and understands the intended wearer and buyer conditions. A number alone cannot make an unsuitable product mix representative.

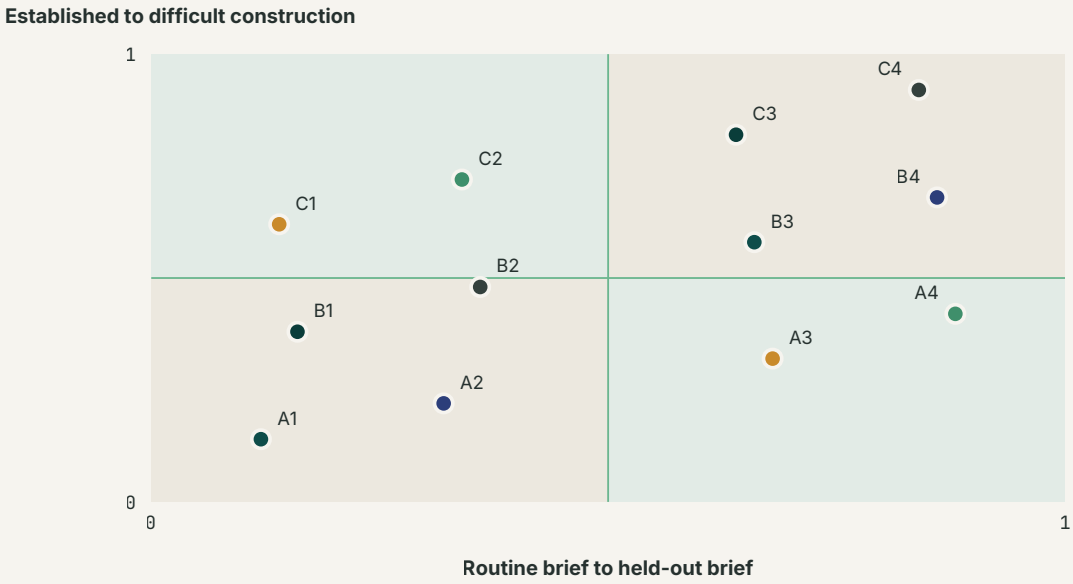
Each family should contain two routine cases and two held-out cases. Routine cases vary an approved block in ways the technician already understands, such as a permitted collar or pocket configuration. Held-out cases introduce a difficult but permissible combination that was not used to tune prompts, templates or integration rules. The difficulty might concern an unfamiliar feature combination, a size and grading interaction, or a material change that requires fresh physical review. It should not be an impossible brief designed to make the system fail. The purpose is to test the stated use envelope, not to create theatre.

GarmentCodeData shows why a synthetic collection does not settle this selection. Its generation process defines design distributions, body-model assumptions and quality controls that exclude inadequate samples. [S30] Those choices make the research dataset inspectable within its own domain. They do not make it a record of the maker's garments, wearers or production history. Local routine and held-out cases therefore need their own inclusion reasons, permissions and acceptance conditions. A research example can inform case design without becoming evidence of local transfer.

Exhibit 52, p. 198 arranges the illustrative sample by brief familiarity and construction challenge. Family labels A, B and C preserve places for the maker's chosen products. Cases 1 and 2 are routine; cases 3 and 4 are held out from adjustment. Coordinates are layout positions, not difficulty measurements or expected scores.

EXHIBIT 52 / ILLUSTRATIVE CASE POSITIONS, NOT PERFORMANCE SCORES

A pilot needs both routine and held-out cases



ILLUSTRATIVE [S19] [S30] GFT planning design: 12 styles = 3 families x 4 styles. Cases 1 and 2 in each family are routine approved-block variants; cases 3 and 4 are held out from prompt, template and integration adjustment. Coordinates arrange brief familiarity and construction challenge only. Product choices, permissions, materials, wearer needs and buyer conditions remain local inputs. This sample has no claimed statistical power.

The exhibit's main value is coverage discipline. An illustrative pilot containing only A1, B1 and C1 would show whether the route handles familiar work. It would say little about transfer. Conversely, a pilot containing only difficult cases would not establish whether assistance helps routine development. Each case needs a frozen brief, block and material revision, size range, permissions and intended acceptance route before work starts. If the team adjusts a held-out case after seeing the output, it should reclassify that case as development evidence and create a new independent test rather than retain a misleading label.

WORKED EXAMPLE: PAIRED ATTEMPTS ARE NOT INDEPENDENT GARMENTS

The illustrative 12-style plan gives 3 families x 4 styles = 12 styles. Running both manual and assisted paths gives 12 × 2 = 24 workflow attempts, while each family contributes 4 × 2 = 8 attempts. These are paired observations on the same styles, not 24 independent garments.

To balance order, 6 styles can take the manual path first and the other 6 can take the assisted path first, with the order reversed for the second attempt. Within each 4-style family, 2 styles follow each initial order. This allocation is an illustrative planning choice. An aborted run remains an attempted workflow with a recorded outcome, while a missing measurement remains missing. Neither should be silently removed to improve the reported result.

The case register should also record exclusions before generation. A brief without input permission may need replacement before it enters the pilot. That is a screening decision, not an assisted failure. By contrast, an authorised case whose generated program cannot be parsed is an attempted case with a generation or validation failure. Keeping those categories separate prevents procurement preparation from being confused with model performance. The final report should show both the screened pool and the evaluated set, with reasons and stable identifiers.

13.3

A matched baseline makes the test interpretable

The manual and assisted paths should answer the same product question under the same acceptance conditions. Each pair needs one frozen brief, the same block and material references, the same required size range and the same physical and receiving-CAD checks. The reviewer should not demand a complete tech pack from manual work while accepting a rendered concept from assisted work. Equally, the assisted path should not receive extra procurement information or expert correction that the manual path is denied. Any necessary difference in support is part of the workflow and must be recorded as such.

NIST's AI Risk Management Framework calls for a human baseline when AI augments or replaces human activity. It also warns that controlled measurements can differ from operational settings. [S19] For this pilot, the baseline is not a generic industry lead time. It is the maker's own observed work on matched styles, with product mix and approval conditions retained. The report should preserve routine and held-out strata so a favourable average on easy variants cannot conceal weak performance on the cases that motivated qualification.

Order effects need a deliberate control. A technician who completes the manual path first may understand the brief better when starting the assisted path. A technician who sees a generated alternative first may anchor the subsequent manual design. Balancing the initial order across styles reduces the risk that one workflow systematically benefits from learning. It does not remove learning within a pair. The record

should therefore retain order, elapsed time between attempts and any reused artefact. Where genuine independence is impossible, the limitation belongs in the interpretation rather than being hidden by the word matched.

The protocol should freeze model, integration and review rules before the evaluation batch. A partner may repair a broken connector, but the repair creates a new configuration whose results should not be pooled uncritically with the earlier one. Prompt or template adjustment on a routine development case is legitimate preparation. The same adjustment on a held-out evaluation case turns that case into part of development. Configuration identifiers and change dates let the reviewer reconstruct which system produced each result.

Every candidate and failed run belongs in the attempt log. Record the brief and output identifiers, model route, run status, validation outcome, expert help, correction reason and final disposition. Include runs that produce no usable image, programs that do not parse, exports that cannot be reopened and candidates abandoned after review. Selecting only visually impressive outcomes would remove precisely the review burden and technical failures the pilot must measure. NIST's generative-AI profile supports empirical capability checks, source verification and lifecycle monitoring, which reinforces the need for inspectable evidence rather than a curated demonstration. [S20]

Time capture should separate generation, technician touch time, expert assistance, queue time, physical sampling, buyer waiting and rework. A short generation event can coexist with a long approval calendar. Correction time should include diagnosis and rechecking, not only the visible edit. Partner effort supplied without an invoice still changes the resource boundary and should be identified. The aim is not stopwatch surveillance. It is to understand which activities are required to obtain an accepted specification and who performs them.

TailorCoPilot's TailorTrace research records discrete pattern states and explicit transformation operations. [S15] GFT can apply the same general principle to the pilot log: preserve the candidate, the technician's edit, the reason and the retest. This makes a manual rescue visible as work and as learning. It also prevents a corrected file from being presented as an unaided model success. The source reports a research system and user study, not a local training or productivity result, so the pilot must still observe its own corrections.

A matched comparison needs rules for missing evidence. If a timestamp fails, report the affected duration as unknown rather than replacing it with an assumption. If physical review is pending, the case has not completed the acceptance route. If a reviewer disagrees on a critical failure, retain both assessments and escalate under the release rule. These are evidence states, not inconveniences to be averaged away. The dataset should permit a reader to move from every summary value back to the case, revision and decision that produced it.

13.4

Hard release conditions must be set before results are seen

The acceptance sheet should be signed before the partner presents the pilot results. GFT proposes six critical conditions: no unresolved geometry, rights, material or fit issue; every applicable required field present; receiving-CAD readback complete; and an identified release authority. These are minimum decision categories, not universal apparel tolerances. The accountable product team must add buyer-specific dimensions, permitted deviations and residual-risk decisions before testing. NIST explicitly does not prescribe an organisation's risk tolerance, so no generic AI framework score can supply the garment threshold. [S19]

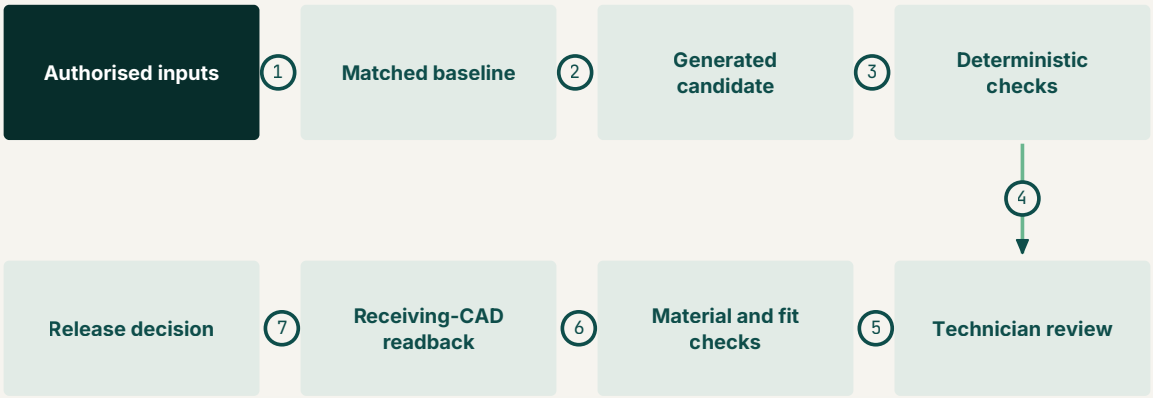
A critical gate cannot be traded against a favourable average. A style with unresolved input permission is not rescued by correct panel geometry. A successful visual fit review is not rescued by, and cannot compensate for, an absent physical material record. A file that opens but loses notches or grading has not completed receiving-CAD readback. The release owner needs to see each condition independently. Composite scores would allow strong but irrelevant performance to conceal the precise reason a specification must remain on hold.

The team should distinguish critical release conditions from advisory improvement measures. Correction hours, candidate volume and elapsed calendar time inform value and manageability. They do not justify overriding a critical failure. Some conditions may be not applicable to a product family, but the reason must be recorded before that family is evaluated. An exception created after an unfavourable result is a change to the protocol. It needs independent review and a new decision, not silent editing of the original checklist.

Exhibit 53, p. 202 shows the proposed route from permissions to the expansion decision. Every arrow represents a passed handoff, not an automatic transition. A failure returns to the responsible correction point with its evidence attached, while the last approved specification remains unchanged.

EXHIBIT 53 / ILLUSTRATIVE EVIDENCE AND DECISION SEQUENCE

Expansion requires every critical handoff to pass



1. Pass
2. Pass
3. Candidate retained
4. Pass
5. Geometry accepted
6. Pass
7. Pass

ILLUSTRATIVE [S19] [S20] [S22] GFT proposed qualification route informed by NIST contextual evaluation and lifecycle monitoring, plus vendor-documented CAD exchange controls. Arrows show required pass conditions. A failure returns to the relevant owner; it does not advance. Physical material and fit requirements are product-specific. No threshold, local result or automatic scale-up is asserted.

The flow also clarifies where evidence originates. Deterministic checks can detect defined schema or geometry failures, but they cannot establish cultural permission, buyer intent or physical comfort. Technician review can correct a pattern, but the correction must be logged rather than credited to generation. Physical checks can support material and fit acceptance within stated conditions, but they do not prove CAD exchange. Readback closes the information route, after which the accountable role decides whether the complete record permits release.

CLO's documentation lists Standard DXF, AAMA and ASTM options and controls for units, grading, notches, seam allowance and cutting or sewing lines. [S22] Those options make an exchange agreement concrete. They do not establish interoperability between the pilot's actual systems. The acceptance sheet should name the source and receiving releases, format, unit, size group and annotation conventions. The reviewer should compare reopened dimensions and required annotations with the approved source, recording pass, fail or a reasoned technician exception.

An exception needs boundaries. It should identify the affected field or feature, why the difference is acceptable, who approved it and whether the same exception may recur. Repeated manual repair of the same exchange defect is not a collection of independ-

ent passes. It is evidence of an unresolved workflow limitation. The pilot report should group such exceptions by cause and include the effort required to keep the route usable.

Release authority should remain explicit even when roles are combined in a small team. The person operating the model may also prepare a specification, but a separate review or documented compensating check should protect the critical decision. A partner can recommend, never self-authorise, production use. The final state should distinguish draft, technically reviewed, sample approved and released for the intended downstream purpose. No generated document should enter cutting merely because it has a complete filename or a polished cover.

13.5

Connectivity failures must preserve the approved work

The programme treats power and network reliability as design constraints, so the pilot should test degradation rather than assume continuous service. That programme statement is not a measurement of factory outages. Household electricity access and national internet use would not fill the evidence gap because they describe different boundaries. The local trial needs its own controlled interruption records and, later, operating observations from the intended site. Wider power and network architecture belongs to companion papers; this chapter concerns continuity of the specification handoff.

GFT recommends durable local copies of approved patterns, material records, tech packs and release histories. A controlled local repository should identify revision, approval state and the file or record digest used for integrity checking. The team must be able to open the last authorised specification without contacting the model provider. Read access alone is insufficient if the technician cannot continue a permitted manual edit, export through the qualified route or reconstruct why a revision was approved.

The recovery design should distinguish four interruption points: before input upload, during generation, after generation but before local validation, and during export or receiving-CAD readback. At each point, the record should show which artefacts exist locally, which external state is uncertain and what action is safe. A repeated submission after a timeout may create duplicate candidates or ambiguous revisions. A resumable transfer needs an attempt identifier and reconciliation rule, not merely a retry button.

A controlled recovery rehearsal can begin with an authorised style whose approved artefacts are already retained locally. The future pilot team should interrupt the external route at a planned step, verify that no unapproved candidate enters the release folder, reopen the last approved pattern and continue the manual path. When service resumes, the team should reconcile queued requests, model and integration versions, output identifiers and successor revisions before accepting any new candidate. The test result should record lost work, duplicates, unresolved states and reviewer actions without assigning invented recovery-time targets.

Version checks prevent a recovered service from silently changing the evaluation. If the provider model, connector or validation library changes during interruption, subsequent attempts belong to a new configuration until the impact is reviewed. NIST's generative-AI profile identifies value-chain and component-integration risks and supports lifecycle monitoring. [S20] For the maker, the practical response is to bind each output to the exact route that produced it and to rerun retained cases when a relevant component changes.

The manual fallback should preserve quality rather than create an emergency bypass. Core orders continue under the existing release rules, with the same physical checks and buyer approvals. Pilot work can pause when its evidence cannot be completed. Local continuity does not require running a foundation model on site, nor does it require owning every component. It requires the approved work, editable tools and decision records needed to operate without the unqualified dependency. That is the resilience condition procurement should test.

13.6

Buyers should retain exit and data-portability options

Exit is a workflow test, not a contractual slogan. The buyer should be able to retrieve approved geometry, grading, material links, artwork, construction instructions, configuration records and decision history in forms its own team can use. A compressed folder of proprietary files may satisfy a superficial export requirement while leaving the maker unable to reconstruct a garment. Portability therefore needs reopen evidence, semantic checks and an identified receiving owner for each artefact class.

The exit rehearsal should use retained approved work before dependence deepens. A reviewer receives the export in a controlled local environment, opens each artefact with the intended alternative or manual route, and compares required content with the release record. The team documents missing fields, unsupported features, conversion effort and any permission or retention constraint. This rehearsal does not require cancelling a service or executing a migration. It establishes whether the approved work remains usable if the relationship or component changes.

Exhibit 54, p. 205 defines the proposed evidence pack. The entries are acceptance requirements to be rehearsed, not claims that a vendor currently supplies every format or that this lane has completed an export. CLO's documentation supports the need to specify DXF options and inspect the receiving result; it does not demonstrate this particular exit route. [S22]

EXHIBIT 54 / ILLUSTRATIVE PORTABILITY ACCEPTANCE CONTRACT

A vendor exit test starts with retained approved artefacts

Artefact	Required retained form	Reopen and reconciliation test	Owner
Base pattern	Editable source and agreed DXF	Compare panels, units and boundaries	Pattern technician
Grading	Size group and rule records	Check sizes, notches and key measures	Pattern technician
Material	IDs, evidence and presets	Relink swatch, test record and revision	Material reviewer
Artwork	Authorised source and repeat data	Check scale, placement and permission	Design and rights
Construction	Tech pack, BOM and instructions	Resolve applicable required fields	Specification owner
Decision history	States, edits and approvals	Reconstruct release and correction reasons	Release authority
Configuration	Model, connector and run records	Identify version and retained test route	Digital owner

ILLUSTRATIVE [S15] [S20] [S22] GFT proposed exit evidence. Formats and receiving tools must be agreed with the actual buyer and partner. Reopen tests compare required content with the approved revision, not merely file presence. DXF documentation supports pattern-exchange fields only. No provider migration, deletion, purchase or completed local readback is claimed.

Pattern and grading checks should compare meaning as well as coordinates. Units, panel identities, cut quantities, seam relationships, notches and size rules need to resolve to the same approved revision. Where the receiving system cannot preserve a feature, the exit record should identify the loss and the controlled reconstruction method. A technician's manual repair is a portability cost and a limitation. It should not be hidden by marking the file transfer successful.

Material and artwork portability requires links to evidence, not just embedded names. A material preset detached from its physical swatch and measurement record cannot support the same release decision. Artwork needs its authorised source, repeat geometry, placement and permission scope. If a service cannot export a generated intermediate representation, the buyer should know whether the remaining editable CAD and source records are sufficient for continued work. That answer depends on the intended future route, not on the quantity of downloaded files.

Decision and configuration history allow a new team to distinguish an approved correction from an unreviewed model output. TailorTrace illustrates the value of explicit pattern states and operation sequences in a research setting. [S15] GFT's exit contract extends the requirement across approvals and component versions. The record should retain who accepted each critical condition, which revision they saw and what later change would reopen their decision. Access controls and retention limits still apply; portability does not justify unrestricted copying of sensitive inputs.

Commercial review should cover export assistance, format limits, retention, deletion, support after termination and the treatment of uploaded content under the actual agreement. Those terms require current documents and jurisdiction-specific review where necessary. This chapter does not infer them from a product demonstration. The buyer's practical objective is narrower and testable: preserve an authorised manual route, reconstruct accepted specifications and avoid discovering irreplaceable records only after the service becomes unavailable.

13.7

The pilot report must distinguish observed results from pending work

The report should let a decision-maker reconstruct what happened without attending the demonstration. It should begin with the protocol actually executed, including product envelope, case selection, matched pairs, model and integration versions, reviewer roles, gate sheet and deviations. A deviation is not automatically fatal. It must be visible enough for the reader to judge its effect. The report should separate cases screened before generation, attempted workflows, completed evaluations and records still pending.

Per-case outcomes come before aggregate ratios. Each style should show routine or held-out status, manual and assisted order, candidate and revision identifiers, critical-gate results, physical checks, CAD readback, release state and unresolved issues. Summary measures then use explicit denominators from those records. An unattempted readback is a coverage gap, not a pass. An unevaluated fit sample is pending, not a rejection. A missing duration is unknown, not zero. These distinctions prevent administrative completion from masquerading as technical acceptance.

Quality results should retain failure taxonomy and correction effort. Parsing, deterministic validation, technician correction, simulated assembly, physical sampling and receiving-CAD review are different events. The report should identify where each failure occurred, what evidence supported diagnosis, who corrected it and whether the retest passed. NIST's generative-AI profile calls for empirical capability checks and monitoring across the lifecycle. [S20] A local report honours that principle by preserving negative results and configuration changes rather than presenting only the final corrected garment.

Time and cost should follow the chapter 12 model and the same accepted-style boundary. Include input preparation, generation, review, correction, expert help, integration, physical samples, logistics and repeat failures. State whether partner effort was donated, included in a quotation or absorbed by another activity. Report touch time separately from queue and waiting time. The pilot can support a downside case using observed inputs; it cannot promise savings from avoided samples unless the same final acceptance was retained and all changed costs were counted.

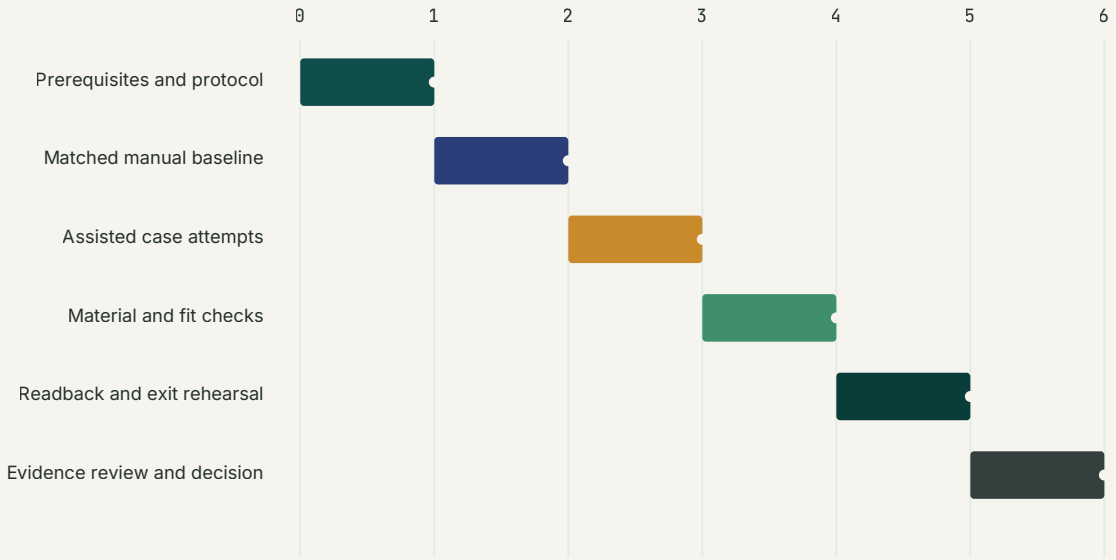
The review should also state what the sample cannot establish. Twelve selected styles, if that illustrative plan is used, would provide bounded case evidence rather than a universal success rate or a statistically powered estimate for all garments. Routine results would not qualify held-out families. A passed fabric and size range

would not qualify new materials or wearers. A single model configuration would not establish performance after an update. Generalisation requires new evidence at the boundary that changes.

Exhibit 55, p. 207 sequences the work by prerequisites. Numeric positions show dependency order only. They are not weeks, durations or resource commitments, and the physical and readback stages should not begin with unapproved candidates merely to protect a planned date.

EXHIBIT 55 / ILLUSTRATIVE DEPENDENCY POSITIONS, NOT ELAPSED TIME

Pilot work should be sequenced by dependencies



ILLUSTRATIVE [S19] [S20] [S22] GFT proposed sequence. Prerequisites include authorised inputs, approved blocks, material records, technicians and a frozen protocol. Baseline and assisted work use matched conditions. Physical checks and receiving-CAD readback depend on suitable candidates. Numeric positions show order only; team capacity, durations and calendar dates are unknown. Review can stop, repair or bound expansion.

The sequence contains return paths even though the exhibit shows the forward dependency chain. A missing permission returns the case to screening. A geometry failure returns it to technical correction. A material failure reopens procurement and physical review. A readback failure reopens exchange configuration. Each return retains the failed attempt and prevents the case from being counted as complete. A revised protocol or component may also require earlier cases to be labelled with their original configuration rather than pooled into one conclusion.

The final decision should be one of stop, repair, hold or bounded expansion. Stop is appropriate when a critical failure makes the route unacceptable, permissions cannot be resolved, or the manual fallback cannot be preserved. Repair applies when an identified defect has a feasible correction and a defined retest. Hold applies when evidence is insufficient or buyer conditions remain unresolved. Bounded expansion is available only when every critical gate passes, the observed workload and downside economics are acceptable, and the authorised envelope is explicit.

Bounded expansion should name the product families, blocks, size range, materials, permissions, model and integration versions, CAD route and retained human checks covered by the evidence. It should assign owners for monitoring and define which change reopens qualification. Additional volume inside that envelope is a different decision from adding new families or removing physical review. The report should record both the permitted next step and the evidence that would reverse it.

Procurement closes the loop by comparing the observed route with the original purchase specification. Components that passed remain candidates for controlled use. Components that failed return to correction or rejection. Pending claims stay pending. The result is a qualified operating boundary rather than a declaration that generative design works everywhere. No local result is supplied by this chapter, and no investment, legal or tax conclusion follows from the proposed protocol.

Qualify the whole handoff, then scale only what the evidence can hold.

CHAPTER SOURCE NOTES

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S19 National Institute of Standards and Technology. 2023-01-26.

S20 National Institute of Standards and Technology. 2024-07-26.

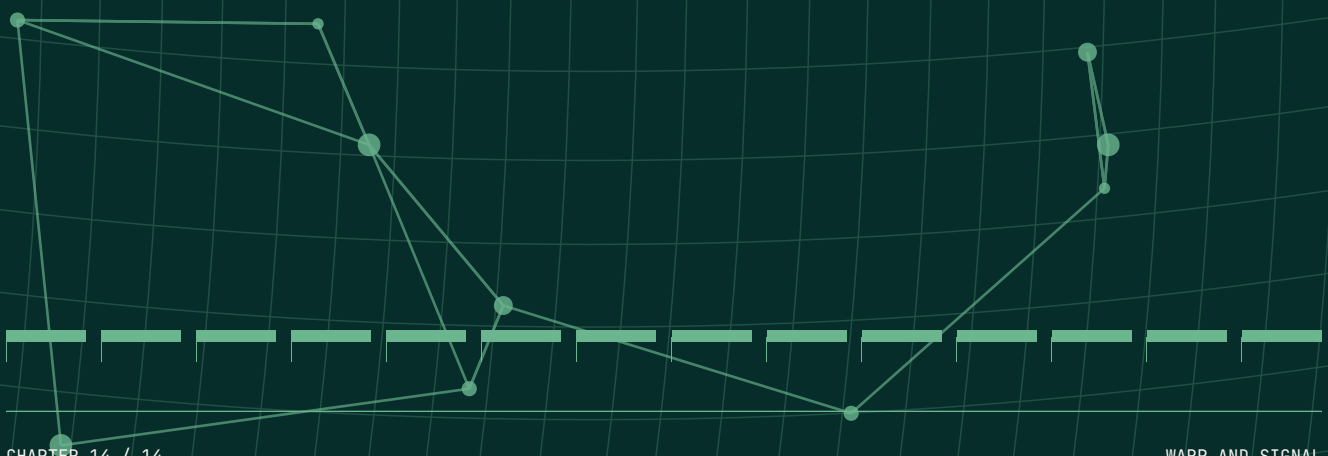
S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

S22 CLO Virtual Fashion. 2026-07-15.

14

GFT should scale measured readiness rather than generated output

Adopt creative assistance conditionally, preserve production authority and expand only when the specification workflow demonstrates value.



CHAPTER 14
DECISION TEST

Adopt creative assistance conditionally, preserve production authority and expand only when the specification workflow demonstrates value.

GFT should adopt generative design conditionally and scale only the parts of the workflow that produce accepted specifications with measured correction effort. Generated images, programs and documents are candidates for review, not evidence of production value. This distinction matters for GFT's Cameroon textile programme, prepared in January 2024 with source notes refreshed in 2026. The inherited sector baseline records about 95 percent of cotton fibre exported raw and about 5 percent processed locally, while the stated national target is 50 percent local processing by 2030. [R06] [R07] Those figures describe an industrial ambition, not an artificial-intelligence return. Capturing more value requires specifications that people can inspect, samples that satisfy agreed conditions and production records that support correction. The recommendation is therefore to retain technician release authority, physical checks, receiving-CAD readback and versioned evidence, then expand use only within a tested envelope of products, materials and model configurations. This chapter proposes that operating sequence. It does not report a deployed GFT capability, a factory trial or a realised saving.

14.1

The programme should retain its staged sequence

The source programme's setup, launch, scale and optimise sequence remains the right adoption structure. Setup establishes records and responsibilities. Launch establishes operating measurement. Scale prepares and qualifies new capabilities against that baseline. Optimisation extends only what has passed its conditions. Within that structure, generative specification remains a partnered pilot in the programme's third phase. It is not an asserted installed capability and should not become a hidden dependency for factory launch.

This sequence is not a reason to postpone all creative exploration. GFT proposes a limited refinement: low-risk evaluation using approved public inputs may begin earlier, provided it stays outside the production release path. Such work can test prompt recording, visual review and the usability of generated artefacts without uploading proprietary blocks, buyer designs or wearer data. Its output is an internal learning record, not an approved garment, a procurement decision or evidence that the future factory will perform better. Keeping this boundary permits useful learning without rewriting the programme's staged controls.

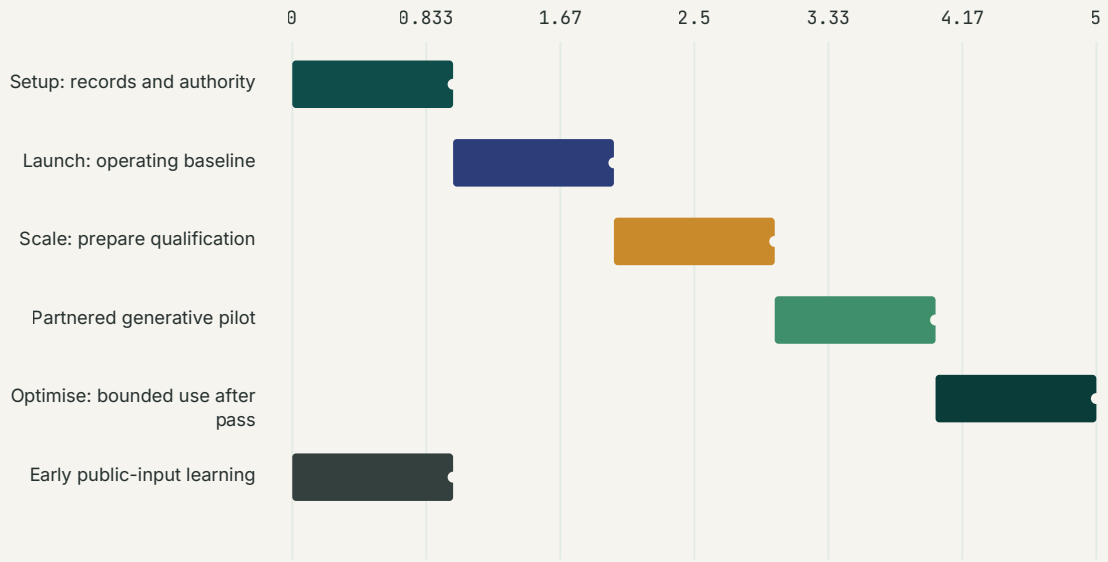
Each stage needs an entry condition, an evidence product and an exit decision. A date alone should not move the programme forward. Setup exits when the maker can reconstruct who approved a manual specification and which inputs it used. Launch exits when the baseline measures are defined consistently and their missing records are visible. Scale exits when the bounded partner evaluation has been designed, partner evidence inspected, routine and difficult cases prepared, and failure and correction recording made ready. The partnered pilot then tests the complete handoff. Bounded optimisation begins only with an explicit decision that names the authorised envelope and remaining limitations.

NIST's AI Risk Management Framework makes risk tolerance contextual and does not prescribe a universal tolerance for the organisation. [S19] GFT should apply that principle through local gate decisions rather than adopt a generic maturity score. The accountable team should state what would block release, what would require repair and what degree of uncertainty it is prepared to retain before seeing favourable pilot results. A high volume of attractive outputs cannot compensate for an unresolved condition that the team has designated critical.

Exhibit 56, p. 211 expresses this logic as a gated roadmap. The positions are sequence markers, not months, deadlines or resource commitments. Early public-input exploration is shown as a separate learning activity because it neither replaces the baseline nor grants access to the production workflow.

EXHIBIT 56 / ILLUSTRATIVE STAGE POSITIONS, NOT ELAPSED TIME

The roadmap is gated by readiness rather than a promised date



ILLUSTRATIVE [S19] GFT proposed sequence based on the programme's setup, launch, scale and optimise structure. Numeric positions show order only. Durations, team capacity and calendar dates are unknown. Early exploration is restricted to approved public inputs and does not authorise production use. The partnered generative-specification pilot remains within the programme's third phase.

The roadmap also makes reversal legitimate. A failed exchange test can return the work to setup if common units and revision rules are absent. An inconsistent baseline can return it to launch. A critical rights or geometry failure can stop the pilot without weakening the core manual workflow. The record should preserve what was learned, which gate failed and what evidence would justify another attempt. Stop and repair are management decisions, not unsuccessful exceptions to a predetermined adoption story.

GFT NOTE

Do not convert phase labels into calendar promises. Authorise the next stage only when its entry evidence is complete enough for the intended decision, and name the role accepting any remaining uncertainty. Keep early exploration separate from sensitive inputs and production release. A stop decision that prevents an unqualified specification reaching cutting is a useful control outcome.

14.2

Setup should build reusable records and decisions

Setup should produce a small, coherent record system before a design model is integrated. GFT recommends six connected registers: owned base blocks, physical materials, design briefs, input permissions, specification revisions and release decisions. Each register needs stable identifiers and an accountable role. The objective is not to accumulate documentation. It is to ensure that a selected concept can be traced to an authorised brief, an editable pattern, identified fabric, applicable permissions and a named approval state.

The base-block record should identify the pattern revision, units, size range, panel and seam identifiers, permitted variants and the location of the controlled editable file. Research such as Design2GarmentCode shows that multimodal concepts can be translated into parametric pattern-making programs, but that research capability does not choose the maker's block policy or approve the resulting geometry. [S09] GFT should begin with blocks the technician can inspect and reproduce manually. Generated variants should enter as successors to an identified block, never as anonymous files detached from their construction history.

Material identity requires the same discipline. A physical swatch, supplier reference, test record and digital simulation preset should link to one material record while retaining their own revisions. A visual texture or model-generated fibre description must not populate a measured-property field. Where a value is absent, the record should say that it is awaiting evidence. This allows procurement and technical review to proceed on other parts of the brief without inventing a material specification for the running cotton-shirt example.

The brief template should separate confirmed requirements from assumptions and unresolved buyer inputs. It should name the intended use, permitted design changes, required size range, available material identifiers, construction restrictions, artwork constraints and approval route. An unknown tolerance remains unknown. The model may propose options, but it cannot turn an absent buyer requirement into a fact. A reviewer should be able to see which constraints were supplied to the model and which were imposed only during later human review.

Permission records should follow the input and the intended destination. A public reference used for early exploration needs a different review from a proprietary block uploaded to an external service or an artwork approved for manufacture. GFT recommends recording origin, permitted activity, destination, restrictions, evidence location

and review state. This is an operational control, not a conclusion about copyright or infringement. An incomplete permission remains visible to the release owner rather than becoming an implied approval.

Revision rules should connect all of these records. A draft, a technically reviewed specification, a sample-approved specification and a bulk-released specification are different states. A later edit creates a successor revision with a reason and identifies which approvals must reopen. CLO's documentation of DXF exchange includes format, unit, grading, notch and cutting or sewing line options, which illustrates why a file name alone is not a sufficient exchange contract. [S22] Setup should record the agreed export settings and the receiving-system readback evidence for each qualified route.

Release authority also needs explicit role separation. Design intent, geometry, material, fit, rights and final release may sit with a small team, but the person proposing a generated change should not silently become its approver. Where roles are combined, the review record should state the compensating check. These are responsibilities, not invented GFT positions or staffing numbers.

A setup exit test should use a manually prepared specification before any generative integration. The team should take an authorised brief through block selection, material identification, technical review, export and receiving-CAD readback. If the team cannot reconstruct that manual decision, adding generated content will create more ambiguity rather than solve it. The research base contains no live-factory readiness assessment, so every prerequisite must be marked present, validated, missing or not applicable in the maker's own inventory, with an owner for each gap.

14.3

Launch should establish the operational baseline

Launch should make the existing specification workflow measurable before assisted work is compared with it. GFT recommends a baseline covering accepted specifications, lead time, rework causes, physical sample rounds and quality outcomes. The programme's capture and management-cadence design provides the operating context, while its inspection-qualification priority reinforces the need to validate measurement before relying on it. Inherited inspection figures are not local performance and should not be reused as a promise for the new workflow.

A baseline begins with a denominator dictionary. An accepted style means a specification that has met the stated release conditions, not a concept selected for review. An attempted CAD readback means an actual import and inspection in the receiving workflow, not an exported file. An evaluated fit sample means a physical sample assessed against the agreed conditions. A correction hour includes diagnosis, editing, rechecking and expert assistance spent after a candidate enters technical review. These definitions should be stable enough for the manual and assisted comparisons to answer the same question.

The event record can remain compact. It should link the style and brief revision, actor role, activity, start and completion records, outcome, failure category, successor revision and evidence location. Touch time should be separated from queue time, material waiting, buyer waiting and rework. Otherwise a faster editing task can be misreported as a shorter approved-tech-pack lead time even though the specification remains idle downstream. Capture should serve process learning rather than employee surveillance, using role-level analysis wherever that answers the management question.

NIST identifies the need for a human baseline when an AI system augments or replaces human activity and warns that controlled measurements can differ from operational settings. [S19] For GFT, this means the baseline should include the actual product mix and approval conditions the pilot intends to test. Routine approved-block variants and more difficult construction cases should be identifiable. Results should retain their spread and missingness rather than collapse into one average that conceals a different mix of work.

The baseline should also distinguish process failure from evidence failure. A missing timestamp is a capture problem. A rejected collar curve is a geometry outcome. A delayed fabric delivery is a material-waiting event. An absent buyer tolerance is an unresolved requirement. Combining these into a generic delay category would hide the corrective owner. Each review should ask whether the record supports a decision, not merely whether a field contains a value.

Synthetic research data cannot fill this gap. GarmentCodeData uses explicit design distributions, body-model assumptions and quality filtering to generate synthetic garments and sewing patterns. [S30] That makes it useful for understanding a research test domain, not a substitute for observations from the maker's products, fabrics, wearers or buyers. GFT should keep research cases, local development cases and local operating records in separately labelled collections.

Launch is ready for the next stage when the team can produce a baseline report with stable definitions, traceable records and an explicit account of missing observations. The required coverage should follow the product envelope and the decision's consequences rather than an arbitrary universal sample size. If records are sparse, the report should say so and restrict the comparison. A weak baseline is a reason to repair measurement before claiming improvement, not a reason to let the assisted workflow define its own success retrospectively.

14.4

Scale should prepare a bounded partner evaluation

Scale should convert the research and procurement questions into a bounded evaluation of a complete handoff. The unit under test is not an image model or a software subscription. It is the path from an authorised brief through generation, deterministic checks, technician correction, material and fit review, receiving-CAD readback and release. The matched pilot design belongs to the preceding chapter; this chapter's concern is the organisational preparation that prevents the pilot becoming a demonstration detached from core work.

The product envelope should be chosen before the partner runs the attractive cases. GFT recommends representative approved blocks and fabrics, with explicit permission for every input. The evaluation set should include routine variants and held-out cases that challenge the intended use, such as an unfamiliar but permissible collar configuration or a material record with missing evidence. Held-out cases must not be used repeatedly to tune prompts and then presented as independent confirmation. Research datasets can inform case design, but their sampling assumptions remain distinct from local production conditions. [S30]

The manual and assisted paths should receive the same brief, material and acceptance requirements. Task order should be balanced where learning could affect the comparison. Every attempt, failed run, manual rescue and expert intervention belongs in the record. Excluding an output because it is obviously unusable would overstate the workflow's value if a technician would still have spent time reviewing it. The evaluation report should preserve per-case outcomes so that a strong average cannot hide a failure at a critical handoff.

Partner evidence should be inspected in the maker's intended environment. CLO documents several DXF formats and controls for units, grading and annotations, but documented options do not prove that a particular generated file will survive a particular receiving workflow. [S22] The partner should supply editable sample artefacts, the exact configuration used, model and integration identifiers, data-handling terms, support boundaries and a current commercial quotation. The technical team should reopen the files and compare dimensions and annotations rather than accept a demonstration video as readback evidence.

Hard release conditions should be fixed before results are reviewed. GFT proposes no unresolved critical geometry, rights, material or fit issue, all applicable required fields present, receiving-CAD readback complete and a named release authority. Buyer tolerances and acceptable residual risk must come from the accountable product process. NIST's framework supports contextual evaluation but supplies no apparel pass threshold. [S19] A partner's benchmark score therefore cannot substitute for the maker's decision rules.

The pilot must also protect core orders. Its review queue, sample capacity and technician time should be visible in the same management system as other development work. No unapproved pilot output should enter a production folder or cutting queue because a delivery deadline is pressing. If the assisted path is unavailable, the team should retain the last authorised specification and continue through the manual route. This keeps the experiment reversible and prevents artificial urgency from becoming a reason to bypass a gate.

Costs should follow the same boundary. Software, integration, training, partner support, review, correction, physical samples and repeat failures should be charged to the evaluation record. A partner's unpaid demonstration assistance is still effort and should be noted even if it does not create an invoice. The resulting evidence can support stop, repair or bounded expansion. It cannot support a claim that additional garment families or suppliers are qualified when they were outside the tested envelope.

14.5

Optimisation should expand only within the tested envelope

Optimisation should treat qualification as a scoped authorisation, not a general statement that generative design works. The decision record should identify the product families, block revisions, size range, material classes, artwork types, model and integration versions, export route, permission classes and human checks covered by the evidence. It should attach complete gate results, downside economics, unresolved limitations and the release roles responsible for the authorised use.

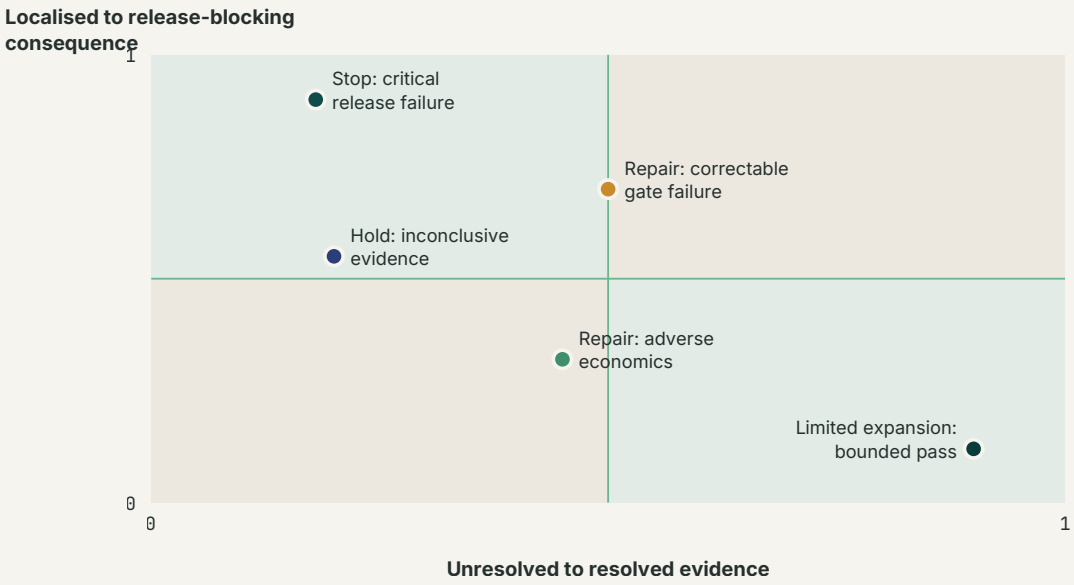
Expansion can then occur one boundary at a time. A maker might retain the same block and material while increasing the number of authorised styles, or add a closely related material after reopening material and fit review. It should not infer that a passed shirt variant qualifies a different garment category, a new body-size range or a new supplier. The change record should state which dimension moved and which evidence remains applicable. This is more deliberate than repeating a full pilot for every minor edit, but safer than treating the original approval as permanent and universal.

A failed critical test holds expansion even if other measures improve. NIST's generative-AI profile supports empirical capability checks and lifecycle monitoring rather than extrapolation from narrow assessments. [S20] GFT should therefore preserve both aggregate results and gate-level outcomes. The release reviewer needs to see unresolved cases, missing records and reviewer disagreement. A high mean score across incompatible checks would conceal precisely the failures that the release process exists to catch.

Exhibit 57, p. 217 summarises the proposed decision logic. Its coordinates arrange decision types for reading. They are not probabilities, risk scores or evidence that a pilot has occurred.

EXHIBIT 57 / ILLUSTRATIVE DECISION POSITIONS, NOT SCORES

A stop decision is a valid pilot outcome



ILLUSTRATIVE [S19] [S20] GFT proposed decision map informed by NIST contextual evaluation and lifecycle-monitoring guidance. Horizontal placement moves from unresolved towards resolved evidence. Vertical placement moves from localised towards release-blocking consequence. Coordinates are arbitrary layout positions. Actual decisions require the defined gates, evidence and accountable review.

The decision categories have different next actions. Stop preserves the evidence and prevents further use of the failed route. Hold identifies the missing observation or unresolved requirement and keeps the current manual process. Repair assigns a correction and a retest to the affected gate. Limited expansion authorises only the passed envelope and retains any extra review conditions. Adverse economics may justify stopping even when technical gates pass, because production readiness is necessary but not sufficient for adoption.

Model, supplier and integration updates require an impact review rather than automatic acceptance. A model change may reopen geometry and regression cases. A changed material may reopen fit and simulation checks. A changed export library may reopen CAD readback even if the generated design is unchanged. [S20] [S22] The last approved artefacts should remain available where permitted, with a manual route if the previous service version cannot be restored.

A stop decision protects the programme's learning value. It prevents sunk software or partner costs from becoming an argument for release. The final record should say what was tested, what failed, what remains unknown and whether another use case is worth evaluating. Optimisation then means improving the measured system, including deciding not to use a component, rather than maximising the quantity of generated work.

14.6

GFT should measure approved outcomes and correction effort

The measurement system should answer whether the workflow produces usable specifications at an acceptable total effort. It should not reward the number of prompts, generated images or exported documents. GFT proposes a compact KPI contract in which every ratio has a defined numerator, denominator and owner, and every period record retains its product, material and revision context. Thresholds should be selected by the accountable team before evaluation, with critical gates reported separately from advisory indicators.

Exhibit 58, p. 218 defines the proposed measures. The roles are responsibilities rather than assumed GFT appointments. The table deliberately avoids a composite readiness score: hours, percentages, counts and release failures do not become comparable merely because they appear on one dashboard.

EXHIBIT 58 / ILLUSTRATIVE MEASUREMENT CONTRACT

Every proposed KPI needs a numerator, denominator and owner

Measure	Numerator or event	Denominator or basis	Decision owner
Approved-tech-pack lead time	Release minus authorised-brief time	Accepted tech packs	Product development
Required-field completeness	Complete applicable required fields	All applicable required fields	Specification owner
CAD readback pass rate	Passing receiving-system readbacks	Attempted readbacks	Pattern technician
Fit approval rate	Samples approved to stated conditions	Evaluated physical samples	Fit owner
Correction effort	Review and correction hours	Accepted styles	Workflow owner
Permission completeness	Assets with complete required permission	Assets requiring permission	Rights reviewer
Physical sampling intensity	Completed physical sample rounds	Accepted styles	Sample owner
Incremental cost	Actual assistance and changed costs	Accepted styles	Management owner

ILLUSTRATIVE [S19] [S20] [S22] GFT proposed measures. Applicable fields, acceptance states and thresholds must be fixed by the accountable team before evaluation. Lead time uses recorded event timestamps. Missing attempts and missing records are reported separately, not counted as passes. No measured GFT values are supplied.

Lead time should be decomposed before it is interpreted. A reduction in technician touch time is useful, but it is not equivalent to a reduction in calendar time if buyer approval or material availability dominates the queue. Required-field completeness should count only fields applicable to the product family, with not-applicable reasons

recorded. CAD pass rate concerns completed receiving-system inspections, while the number not attempted remains a coverage gap. Fit approval concerns evaluated physical samples and must not be inferred from a rendered image.

Correction effort needs an equally strict boundary. It includes the technician's diagnosis, editing, validation, repeat export and expert help for every candidate, including rejected work. Dividing only by accepted styles shows the effort required to obtain the useful outcome. Reporting effort per attempted style alongside it can help explain workload, but substituting that denominator would make rejected work disappear from the unit cost of acceptance. Specification escapes after release should remain separate and carry their severity and containment record.

Every KPI also needs a change log. If the product mix, buyer conditions or acceptance definition changes, the team should mark a break in the series rather than compare unlike periods. A revised denominator should be applied prospectively, with the earlier definition retained. Where a baseline is restated to make a paired comparison possible, the report should show both the original and restated basis. This prevents an apparent improvement from arising through reclassification rather than better work.

The review should connect each measure to a decision and an anti-gaming check. A high readback pass rate is weak evidence if difficult exports are never attempted. A low correction-hours figure is not an improvement if unresolved candidates are removed from the record. A short lead time is not beneficial if required physical checks are skipped. The owner should therefore review outcome, coverage and critical failures together, then record the action, responsible role and evidence required at the next review.

WORKED EXAMPLE: DIFFERENT DENOMINATORS PREVENT A MISLEADING READINESS AVERAGE

An illustrative evaluation batch contains 12 candidate styles, of which 9 reach accepted-tech-pack status. The team records 108 correction hours across all candidates. Correction effort per accepted style is 108 / 9 = 12 hours. Effort per candidate is 108 / 12 = 9 hours, but that second figure must not replace the first in the acceptance-cost comparison.

Within the same illustrative constructed batch, 8 of 10 attempted receiving-CAD readbacks pass, giving 8 / 10 = 80.0 percent. Eight of 11 evaluated physical samples receive fit approval, giving 8 / 11 = 72.7 percent, rounded to a decimal place. The two unattempted readbacks and one unevaluated sample are coverage gaps, not passes. Two styles also retain unresolved permission records, so the proposed critical permission gate remains closed.

These figures demonstrate only the arithmetic and reporting boundary. They are hypothetical, are not a recommended sample size and do not describe a GFT pilot. Averaging the percentages would not resolve the missing coverage or the permission hold. The decision record should show each denominator and keep expansion on hold until the critical condition is addressed.

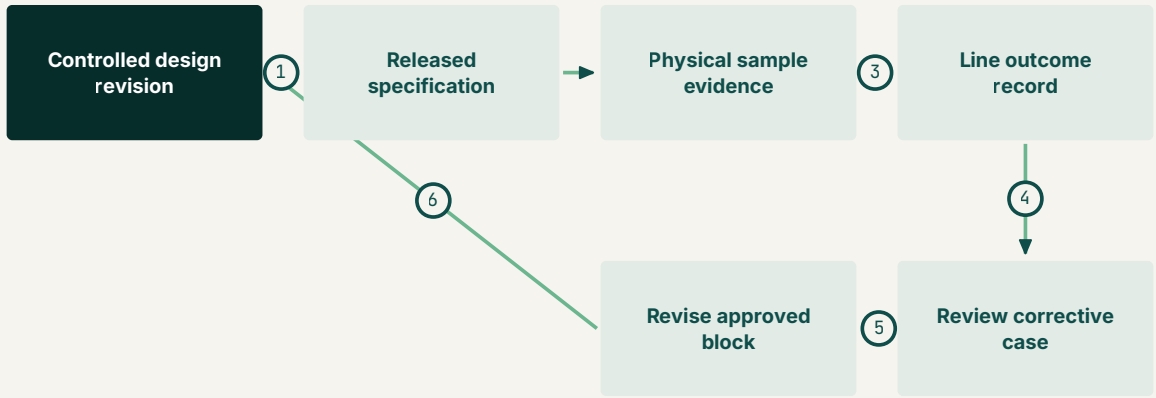
The approved style should then feed the same learning spine as later physical outcomes. TailorCoPilot's TailorTrace research records discrete pattern states and explicit transformation operations, illustrating how corrections can become inspectable

learning records rather than undocumented file overwrites. [S15] GFT's proposed loop extends that principle across specification release, physical sampling and line feedback. It does not claim that the research system has been deployed or that a production outcome has improved.

Exhibit 59, p. 220 shows the required return path. A line outcome or sample failure should not directly rewrite an approved block. It should create a corrective case, receive technical review and produce a new controlled revision only after the relevant gates pass.

EXHIBIT 59 / ILLUSTRATIVE RECORD AND DECISION SEQUENCE

The approved style should feed the same measurement spine



1. Gates passed
3. Authorised use
4. Exception or learning
5. Retest and approve
6. New controlled revision

ILLUSTRATIVE [S15] [S19] [S22] GFT proposed learning loop. Research on version-controlled pattern states informs the record mechanism; NIST informs evaluation; CAD documentation informs exchange checks. Physical and line outcomes are proposals for local capture, not observed GFT results. Corrective action requires review before a new block revision is approved.

At each pilot and operational review, the team should examine results by product family and failure category, not just the portfolio total. Missingness, model version, material revision and manual overrides should travel with the KPI extract. Proposed tools such as `kpi-cascade`, `sampling-cost` and `risk-matrix` can organise these records, but the tool names are not evidence that the measures have been collected or that a benefit has occurred.

14.7

The evidence agenda must remain open

Conditional adoption leaves a concrete research agenda. Before a partner pilot, the maker still needs actual product and material records, authorised blocks, buyer requirements, current vendor terms and quotes, and evidence that the intended export route can be read back. Before bounded expansion, it needs physical validation, per-case pilot outcomes, complete correction and cost records, and a decision on every critical gate. During authorised use, it needs repeatability checks after relevant model, supplier and integration changes. [S19] [S20]

GFT recommends maintaining an open-evidence register beside the roadmap. Each entry should state the question, required artefact, accountable owner, decision blocked and current status. A question should close only when the named evidence has been inspected. A vendor document can close a question about documented functionality while leaving local interoperability unresolved. [S22] A synthetic research dataset can inform test construction while leaving local wearer and material representativeness unresolved. [S30]

The same discipline applies to value claims. No guaranteed saving follows from a smaller number of sample rounds unless acceptance is maintained and all changed costs are counted. No automatic legal clearance follows from a generation log or a software permission. No demonstrated environmental benefit follows from digital sampling alone; the NIST generative-AI profile identifies environmental impacts as a risk, so resource use and any avoided physical flows need their own boundaries and records. [S20] These remain evaluation questions, not reasons to reject learning.

GFT's recommendation is therefore to proceed with records and operational measurement, then qualify a bounded partnered generative-specification pilot when the prerequisites exist. Expand only the envelope that passes, retain manual production authority and stop where critical evidence fails. There is no field-proven GFT generative-design result in this research base, no guaranteed economic saving, no automatic legal clearance and no demonstrated environmental benefit. This chapter is not investment, legal or tax advice.

Scale the evidence of readiness, not the volume of generation.

CHAPTER SOURCE NOTES

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

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

S19 National Institute of Standards and Technology. 2023-01-26.

S09 Authors via arXiv; CVPR 2025. 2024-12-11.

S22 CLO Virtual Fashion. 2026-07-15.

S30 Authors via arXiv; ECCV 2024. 2024-05-27.

S20 National Institute of Standards and Technology. 2024-07-26.

S15 Authors via arXiv; UIST 2026 listed. 2026-08-26.

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

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.

R02 Inherited evidence R02 GRADE C

Fibre2Fashion; Business in Cameroon. Publication date unrecorded. fibre2fashion.com; 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.

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

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.

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

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.

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

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.

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

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.

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

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.

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

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.

c01-S01 Transforming Fashion with AI: A Comparative Study of Large Language Models in Apparel Design GRADE B

Authors via 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-S02 GarmentCode: Programming Parametric Sewing Patterns GRADE B

Authors via arXiv; ACM Transactions on Graphics. 2023-06-06. https://arxiv.org/abs/2306.03642

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S09 Design2GarmentCode: Turning Design Concepts to Tangible Garments Through Program Synthesis GRADE B

Authors via arXiv; CVPR 2025. 2024-12-11. https://arxiv.org/abs/2412.08603

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S10 **ChatGarment: Garment Estimation, Generation and Editing via Large Language Models** GRADE B

Authors via arXiv; CVPR 2025. 2024-12-23. https://arxiv.org/abs/2412.17811

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S12 **AutoSew: A Geometric Approach to Stitching Prediction with Graph Neural Networks** GRADE B

Authors via arXiv; WACV 2026. 2026-02-25. https://arxiv.org/abs/2602.22052

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S13 **FIT: A Large-Scale Dataset for Fit-Aware Virtual Try-On** GRADE B

Authors via arXiv; SIGGRAPH 2026. 2026-04-09. https://arxiv.org/abs/2604.08526

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S17 **EasyFashion: A Human-AI Co-Creation System for Personalized Fashion Design and Sewing Pattern Generation** GRADE B

Authors via arXiv. 2026-09-16. https://arxiv.org/abs/2609.18483

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S19 **Artificial Intelligence Risk Management Framework (AI RMF 1.0)** GRADE A

National Institute of Standards and Technology. 2023-01-26. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S20 **Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile** GRADE A

National Institute of Standards and Technology. 2024-07-26. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S22 **2D Pattern (DXF) Import/Export** GRADE C

CLO Virtual Fashion. 2026-07-15. https://support.clo3d.com/hc/en-us/articles/115000493067-2D-Pattern-DXF-Import-Export

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S26 **IMAGGarment: Fine-Grained Garment Generation for Controllable Fashion Design** GRADE B

Authors via arXiv. 2025-04-17. https://arxiv.org/abs/2504.13176

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S27 **EASE: Parametric garment design with explicit and local ease control** GRADE B

Authors via arXiv; Computers & Graphics. 2026-06-28. https://arxiv.org/abs/2606.29419

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S30 **GarmentCodeData: A Dataset of 3D Made-to-Measure Garments With Sewing Patterns** GRADE B

Authors via arXiv; ECCV 2024. 2024-05-27. https://arxiv.org/html/2405.17609v3

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S0101 **GarmentDiffusion: 3D Garment Sewing Pattern Generation with Multimodal Diffusion Transformers** GRADE B

Authors via arXiv. 2025-04-30. https://arxiv.org/abs/2504.21476

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S0102 **GarmageNet: A Multimodal Generative Framework for Sewing Pattern Design and Generic Garment Modeling** GRADE B

Authors via arXiv. 2025-04-02. https://arxiv.org/abs/2504.01483

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S02 Same source record as c01-S02.

c02-S03 **Towards Garment Sewing Pattern Reconstruction from a Single Image** GRADE B

Authors via arXiv; ACM Transactions on Graphics. 2023-11-07. https://arxiv.org/abs/2311.04218

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S04 **DressCode: Autoregressively Sewing and Generating Garments from Text Guidance** GRADE B

Authors via arXiv. 2024-01-29. https://arxiv.org/abs/2401.16465

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S05 **Garment3DGen: 3D Garment Stylization and Texture Generation** GRADE B

Authors via arXiv; 3DV 2025. 2024-03-27. https://arxiv.org/abs/2403.18816

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S07 **HAIGEN: Towards Human-AI Collaboration for Facilitating Creativity and Style Generation in Fashion Design** GRADE B

Authors via arXiv. 2024-08-01. https://arxiv.org/abs/2408.00855

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S08 AIpparel: A Multimodal Foundation Model for Digital Garments GRADE B

Authors via arXiv. 2024-12-05. https://arxiv.org/abs/2412.03937

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S09 Same source record as c01-S09.

---

c02-S10 Same source record as c01-S10.

---

c02-S11 Image2Garment: Simulation-ready Garment Generation from a Single Image GRADE B

Authors via arXiv. 2026-01-14. https://arxiv.org/abs/2601.09658

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S13 Same source record as c01-S13.

---

c02-S14 PatternGSL: A Structured Specification Language for Template-Free and Simulation-Ready 3D Garments GRADE B

Authors via arXiv. 2026-06-23. https://arxiv.org/abs/2606.24564

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S16 GarmentWeaver: Schema-Aware Structured Synthesis for Multimodal Sewing Patterns GRADE B

Authors via arXiv. 2026-08-31. https://arxiv.org/abs/2608.30550

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S18 Inverse Garment Design with Visual Agents GRADE B

Authors via arXiv. 2026-10-11. https://arxiv.org/abs/2610.15582

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S20 Same source record as c01-S20.

---

c02-S22 Same source record as c01-S22.

---

c02-S27 Same source record as c01-S27.

---

c02-S31 Adobe Firefly FAQ GRADE C

Adobe. 2026-05-13. https://helpx.adobe.com/firefly/web/get-started/learn-the-basics/adobe-firefly-faq.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S01 Same source record as c01-S01.

c03-S02 Same source record as c01-S02.

c03-S04 Same source record as c02-S04.

c03-S07 Same source record as c02-S07.

c03-S09 Same source record as c01-S09.

c03-S10 Same source record as c01-S10.

c03-S13 Same source record as c01-S13.

c03-S19 Same source record as c01-S19.

c03-S20 Same source record as c01-S20.

c03-S21 Copyright Office Releases Part 2 of Artificial Intelligence Report GRADE A

U.S. Copyright Office. 2025-01-29. https://www.copyright.gov/newsnet/2025/1060.html

In the United States copyright context, human expressive control is distinguished from merely supplying text prompts to a generative model.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S23 Indigenous Peoples and Fashion GRADE A

World Intellectual Property Organization. 2026. https://www.wipo.int/en/web/traditional-knowledge/fashion

Fashion references to traditional cultural expressions require context-sensitive, respectful engagement with custodian communities rather than informal copying.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S24 Traditional Cultural Expressions GRADE A

World Intellectual Property Organization. 2026. https://www.wipo.int/en/web/traditional-knowledge/traditional-cultural-expressions/index

Traditional cultural expressions embody community identity and know-how; legal and ethical protection depends on jurisdiction and context.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S26 Same source record as c01-S26.

c03-S0301 GarmentSketch: Large-scale Sketch-to-Fashion Benchmark GRADE B

Authors via arXiv. 2026-06-12. https://arxiv.org/abs/2606.14025

A benchmark of 26,249 fashion sketches across 21 garment categories paired with MLLM captions refined by humans demonstrates both utility and current limitations in sketch-guided generation.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S0302 **Fine-Grained Controllable Apparel Showcase Image Generation via Garment-Centric Outpainting** GRADE B

Authors via arXiv. 2025-03-03. https://arxiv.org/abs/2503.01294

Garment-centric outpainting preserves an existing segmented garment image while generating showcase poses and model appearance around it, illustrating that marketing showcase synthesis and garment construction design are distinct tasks.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S02 Same source record as c01-S02.

c04-S03 Same source record as c02-S03.

c04-S04 Same source record as c02-S04.

c04-S08 Same source record as c02-S08.

c04-S09 Same source record as c01-S09.

c04-S10 Same source record as c01-S10.

c04-S12 Same source record as c01-S12.

c04-S14 Same source record as c02-S14.

c04-S15 **TailorCoPilot: Enabling Agentic Pattern Making with Version-Controlled State Tracking** GRADE B

Authors via arXiv; UIST 2026 listed. 2026-08-26. https://arxiv.org/abs/2608.25462

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S16 Same source record as c02-S16.

c04-S17 Same source record as c01-S17.

c04-S18 Same source record as c02-S18.

c04-S19 Same source record as c01-S19.

c04-S20 Same source record as c01-S20.

c04-S22 Same source record as c01-S22.

c04-S27 Same source record as c01-S27.

c04-S30 Same source record as c01-S30.

c04-S0401 **Generating Datasets of 3D Garments with Sewing Patterns** GRADE B

Authors via arXiv; NeurIPS 2021 Datasets and Benchmarks. 2021-09-12. https://arxiv.org/abs/2109.05633

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0402 **NeuralTailor: Reconstructing Sewing Pattern Structures from 3D Point Clouds of Garments** GRADE B

Authors via arXiv; ACM Transactions on Graphics. 2022-01-31. https://arxiv.org/abs/2201.13063

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0403 **ISP: Multi-Layered Garment Draping with Implicit Sewing Patterns** GRADE B

Authors via arXiv. 2023-05-23. https://arxiv.org/abs/2305.14100

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S0404 **Learning Sewing Patterns via Latent Flow Matching of Implicit Fields** GRADE B

Authors via arXiv. 2026-01-25. https://arxiv.org/abs/2601.17740

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S06 **ContourCraft: Learning to Resolve Intersections in Neural Multi-Garment Simulations** GRADE B

Authors via arXiv; SIGGRAPH 2024. 2024-05-15. https://arxiv.org/abs/2405.09522

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S12 Same source record as c01-S12.

c05-S14 Same source record as c02-S14.

c05-S18 Same source record as c02-S18.

c05-S19 Same source record as c01-S19.

c05-S20 Same source record as c01-S20.

c06-S01 Same source record as c01-S01.

c06-S02 Same source record as c01-S02.

c06-S05 Same source record as c02-S05.

c06-S06 Same source record as c05-S06.

c06-S09 Same source record as c01-S09.

c06-S10 Same source record as c01-S10.

c06-S11 Same source record as c02-S11.

c06-S12 Same source record as c01-S12.

c06-S14 Same source record as c02-S14.

c06-S18 Same source record as c02-S18.

c06-S19 Same source record as c01-S19.

c06-S20 Same source record as c01-S20.

c06-S22 Same source record as c01-S22.

c06-S27 Same source record as c01-S27.

c06-S28 OmniFabric: Coherent UV Space Texture Synthesis for 3D Garment Reconstruction GRADE B

Authors via arXiv; SIGGRAPH Asia 2026 acceptance listed. 2026-09-24. https://arxiv.org/abs/2609.30234

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S29 Product Standard GRADE B

GHG Protocol, World Resources Institute and World Business Council for Sustainable Development. 2026. https://ghgprotocol.org/product-standard

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S30 Same source record as c01-S30.

c06-S32 Creating 3D Fabric GRADE C

CLO Virtual Fashion, CONNECT Support. 2025-04-17. https://support-connect.clo-set.com/hc/en-us/articles/45304313148313-Creating-3D-Fabric

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0601 How Will It Drape Like? Capturing Fabric Mechanics from Depth Images GRADE B

Authors via arXiv. 2023-04-13. https://arxiv.org/abs/2304.06704

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S0602 **Learnable Persistent Wrinkle Formation in Cloth Simulation** GRADE B

Authors via arXiv. 2026-09-12. https://arxiv.org/abs/2609.13707

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S13 Same source record as c01-S13.

c07-S19 Same source record as c01-S19.

c07-S20 Same source record as c01-S20.

c07-S22 Same source record as c01-S22.

c07-S27 Same source record as c01-S27.

c07-S30 Same source record as c01-S30.

c08-S01 Same source record as c01-S01.

c08-S02 Same source record as c01-S02.

c08-S04 Same source record as c02-S04.

c08-S09 Same source record as c01-S09.

c08-S12 Same source record as c01-S12.

c08-S14 Same source record as c02-S14.

c08-S15 Same source record as c04-S15.

c08-S19 Same source record as c01-S19.

c08-S20 Same source record as c01-S20.

c08-S21 Same source record as c03-S21.

c08-S22 Same source record as c01-S22.

c08-S23 Same source record as c03-S23.

c08-S24 Same source record as c03-S24.

c08-S25 **Sustainable and Circular Textiles Strategy** GRADE A

European Commission, Directorate-General for Environment. 2022-03-30. https://environment.ec.europa.eu/strategy/textiles-strategy_en

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c08-S26 Same source record as c01-S26.

---

c08-S28 Same source record as c06-S28.

---

c08-S31 Same source record as c02-S31.

---

c08-S32 Same source record as c06-S32.

---

c08-S0801 **TexTile: A Differentiable Metric for Texture Tileability** GRADE B

Authors via arXiv. 2024-03-19. https://arxiv.org/abs/2403.12961

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c08-S0802 **FabricDiffusion: High-Fidelity Texture Transfer for 3D Garments Generation from In-The-Wild Clothing Images** GRADE B

Authors via arXiv. 2024-10-02. https://arxiv.org/abs/2410.01801

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c08-S0803 **Structured Pattern Expansion with Diffusion Models** GRADE B

Authors via arXiv. 2024-11-12. https://arxiv.org/abs/2411.08930

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c08-S0804 **Content Credentials: C2PA Technical Specification 2.1** GRADE B

Coalition for Content Provenance and Authenticity. 2024-09. https://c2pa.org/specifications/specifications/2.1/specs/C2PA_Specification.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c09-S10 Same source record as c01-S10.

---

c09-S15 Same source record as c04-S15.

---

c09-S19 Same source record as c01-S19.

---

c09-S20 Same source record as c01-S20.

c09-S22 Same source record as c01-S22.

c09-S32 Same source record as c06-S32.

c10-S01 Same source record as c01-S01.

c10-S02 Same source record as c01-S02.

c10-S07 Same source record as c02-S07.

c10-S08 Same source record as c02-S08.

c10-S09 Same source record as c01-S09.

c10-S10 Same source record as c01-S10.

c10-S11 Same source record as c02-S11.

c10-S12 Same source record as c01-S12.

c10-S13 Same source record as c01-S13.

c10-S14 Same source record as c02-S14.

c10-S15 Same source record as c04-S15.

c10-S16 Same source record as c02-S16.

c10-S18 Same source record as c02-S18.

c10-S19 Same source record as c01-S19.

c10-S20 Same source record as c01-S20.

c10-S21 Same source record as c03-S21.

c10-S22 Same source record as c01-S22.

c10-S23 Same source record as c03-S23.

c10-S26 Same source record as c01-S26.

c10-S27 Same source record as c01-S27.

c10-S30 Same source record as c01-S30.

c10-S32 Same source record as c06-S32.

c10-S1001 The Metacognitive Demands and Opportunities of Generative AI GRADE B

Authors via arXiv; ACM CHI 2024. 2023-12-18. https://arxiv.org/abs/2312.10893

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1002 Explanations Can Reduce Overreliance on AI Systems During Decision-Making GRADE B

Authors via arXiv; Proceedings of the ACM on Human-Computer Interaction. 2022-12-13. https://arxiv.org/abs/2212.06823

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1003 Same source record as c04-S0402.

c10-S1004 Aligning Constraint Generation with Design Intent in Parametric CAD GRADE B

Authors via arXiv. 2025-04-17. https://arxiv.org/abs/2504.13178

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1005 TailorCoPilot: Enabling Agentic Pattern Making with Version-Controlled State Tracking (v3 HTML full text) GRADE B

Authors via arXiv HTML; UIST 2026. 2026-08-30. https://arxiv.org/html/2608.25462v3

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S07 Same source record as c02-S07.

c11-S09 Same source record as c01-S09.

c11-S10 Same source record as c01-S10.

c11-S12 Same source record as c01-S12.

c11-S15 Same source record as c04-S15.

c11-S19 Same source record as c01-S19.

c11-S20 Same source record as c01-S20.

c11-S21 Same source record as c03-S21.

c11-S22 Same source record as c01-S22.

c11-S23 Indigenous Peoples and Fashion GRADE A

World Intellectual Property Organization. 2023-11-02. https://www.wipo.int/en/web/traditional-knowledge/fashion

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S24 Traditional Cultural Expressions GRADE A

World Intellectual Property Organization. n.d.. https://www.wipo.int/en/web/traditional-knowledge/traditional-cultural-expressions/index

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S30 Same source record as c01-S30.

c11-S31 Same source record as c02-S31.

c11-S32 Same source record as c06-S32.

c11-S1101 Copyright Office Launches New Artificial Intelligence Initiative GRADE A

U.S. Copyright Office. 2023-03-16. https://www.copyright.gov/newsnet/2023/1004.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S1102 Government launches public consultation on enhancing Copyright Ordinance regarding protection for artificial intelligence technology development GRADE A

Commerce and Economic Development Bureau, Hong Kong Special Administrative Region Government. 2024-07-08. https://www.cedb.gov.hk/en/news/press_release/2024/pr08072024a.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S1103 Content Credentials: C2PA Technical Specification 2.1 GRADE B

Coalition for Content Provenance and Authenticity. 2024-09. https://spec.c2pa.org/specifications/specifications/2.1/specs/C2PA_Specification.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S15 Same source record as c04-S15.

c12-S19 Same source record as c01-S19.

c12-S20 Same source record as c01-S20.

c12-S22 Same source record as c01-S22.

c12-S25 Same source record as c08-S25.

c12-S29 Same source record as c06-S29.

c13-S09 Same source record as c01-S09.

c13-S15 Same source record as c04-S15.

c13-S19 Same source record as c01-S19.

c13-S20 Same source record as c01-S20.

c13-S22 Same source record as c01-S22.

c13-S30 Same source record as c01-S30.

c14-S09 Same source record as c01-S09.

c14-S15 Same source record as c04-S15.

c14-S19 Same source record as c01-S19.

c14-S20 Same source record as c01-S20.

c14-S22 Same source record as c01-S22.

c14-S30 Same source record as c01-S30.

Stable definitions make comparisons useful.

Accepted style

A design whose specified review and release conditions have been met; its exact acceptance state must be recorded.

Agentic refinement

An iterative process in which software selects actions, inspects results and proposes further changes under defined controls.

AAMA DXF

An apparel convention for exchanging pattern data through DXF; exact options and receiving-system behaviour require verification.

ASTM DXF

An apparel pattern-exchange convention supported by some CAD tools; format support alone does not establish interoperability.

Avatar

A digital representation of a body used for visualisation or simulation.

Base block

A reference pattern used as the starting point for related designs and controlled variants.

Bill of materials

A structured list of the material and trim inputs specified for a product, with identifiers and quantity bases.

Body measurement

A dimension recorded on a wearer or body representation using a declared method, posture and unit.

CAD round trip

Export and receiving-system reopen/readback that checks whether required geometry and annotations survive exchange.

Collision

Contact or intersection between simulated garment elements or between garment and body geometry.

Construction specification

Instructions describing how components are assembled, including seam, stitch and reinforcement requirements.

Cut line

The boundary used for cutting a pattern piece; it can differ from the stitching line by the seam allowance.

Design intent

The intended visual and functional properties that technical development must preserve or explicitly revise.

Domain-specific language

A programming language designed to express the concepts and operations of a particular task domain.

Drape

The shape and deformation of fabric under its loading, constraints and material properties.

DXF

A CAD exchange format whose supported content and interpretation depend on conventions, software and settings.

Ease

The allowance of garment material relative to the body, specified for appearance, comfort or movement.

Evidence envelope

The proposed collection of source, version, permission, validation and approval records accompanying an artefact.

F1-score

The harmonic mean of precision and recall for a stated task and unit of evaluation; not a whole-product acceptance measure by default.

Fit approval

A recorded decision that a sample meets the agreed fit conditions for specified wearers, sizes and activities.

Grading

The controlled development of related pattern sizes from a base design using defined rules and checks.

Grain direction

The intended alignment of a pattern piece with the fabric structure, recorded for cutting and behaviour.

Held-out case

A test example excluded from the relevant development or adjustment process to assess transfer beyond tuned examples.

Interoperability

The demonstrated ability of systems to exchange and use the information required for the intended workflow.

JSON

A structured text format for data; valid syntax does not ensure that its values describe a valid garment.

Material preset

A collection of simulation material inputs whose measurement or estimation basis should be declared.

Mood board

A collection of visual references expressing creative direction rather than a complete manufacturing instruction.

Notch

A reference mark used to help locate, align or assemble pattern pieces.

Panel

An individual garment pattern component with geometry and relationships to other pieces.

Parametric pattern

A pattern constructed or modified through defined parameters and relationships.

Physical sample round

A distinct physical prototyping and review iteration, with its cost and outcome recorded.

Point of measure

A defined garment dimension with specified landmarks, method and units.

Provenance

The recorded origin and transformation history of data or an artefact; it does not itself confer legal permission.

Release authority

The accountable role authorised to approve a defined specification state for its intended downstream use.

Repeat tile

An artwork unit intended to repeat across a surface at a specified scale and orientation.

Schema

A formal description of required data structures, fields, types and relationships.

Seam allowance

Material outside an intended stitching line retained for assembly or finishing.

Seam correspondence

The relationship identifying which pattern boundaries are to be joined, including order and direction.

Simulation-ready

Prepared for a specified simulation workflow; it does not automatically mean production-ready.

Synthetic dataset

Data generated by a computational process rather than directly collected as the corresponding real-world observations.

Tech pack

A controlled set of product specification information used for technical review, sampling and production communication.

Traditional cultural expression

A cultural expression associated with community heritage and identity whose use requires attention to relevant rights and context.

UV map

A mapping between a surface and texture coordinates; it should not be confused with a physical print specification.

Virtual try-on

A method that visualises a garment on a person or avatar; fit accuracy depends on the method and its inputs.

Whole-garment assembly success

The proportion of evaluated digital garments assembled without the stated errors under a specified benchmark protocol.

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.

AAMA DXF 21 25 64 74 122

Accepted style 24 25 58 119 148 183 184 187 189 190 191 213 219

Agentic refinement 239

ASTM DXF 64 74 75 122

Avatar 21 87 95 99 100 102 103 108 110 111 116 117 118 119 125

Base block 45 47 54 68 73 75 95 103 148 162 173 183

Bill of materials 67 105 131 132 140 143 144 183

Body measurement 100 111 118

CAD round trip 22 138

Collision 70 85 87 100 101

Construction specification 82 151 170

Cut line 127

Design intent 38 111 114 122 138 152 162 163 176 213

Domain-specific language 31 61 63

Drape 18 25 53 70 92 98 101 102 103 111 145 156 158 196

DXF 21 22 23 24 25 28 39 65 69 73 74 75 101 104 115 122 125 135 141 144 152 156 158 161 162 164 165 174 176 177 178 179 202 204 213 215

Ease 17 18 21 23 25 28 32 36 37 45 66 75 82 86 102 103 104 108 109 111 112 113 114 116 151 152 153 155 158 162 164 165

Evidence envelope 144 168 172

F1-score 21 80 176

Fit approval 25 37 108 114 219

Grading 17 21 23 24 28 39 68 75 95 110 111 113 114 115 118 119 122 123 125 127 131 138 141 142 144 147 162 164 173 177 183 197 201 202 204 205 213 215

Grain direction 240

Held-out case 198

Interoperability 39 141 202 221

JSON 17 34 35 41 54 57 60 63 64 65 68 72 73 74 75 138 140 152 171

Material preset 95 116 205

Mood board 47 133

Notch 15 21 61 75 78 79 81 82 114 115 125 143 144 151 156 158 177 213

Panel 15 17 18 21 23 33 38 45 49 50 57 60 61 63 65 66 67 68 75 78 79 81 85 88 92 95 98 99 100 101 102 111 114 116 117 122 125 126 127 141 142 143 144 148 151 153 154 158 161 162 165 170 183 196 201 205 212

Parametric pattern 15 17 25 27 34 54 176 196 212

Physical sample round 240

Point of measure 18 112 113 114 115

Provenance 18 23 51 52 95 97 122 131 132 133 135 138 161 162 163 171 173 175 179 197

Release authority 18 138 141 152 162 176 184 196 201 203 210 213 215

Repeat tile 123 126 127 170

Schema 31 32 38 39 42 61 62 64 65 68 69 73 96 110 139 140 202

Seam allowance 18 74 83 87 104 112 125 126 127 151 202

Seam correspondence 25 78 83 101

Simulation-ready 27 37 61 92

Synthetic dataset 60

Tech pack 24 33 53 105 132 135 138 139 140 144 145 147 168 171 179 183 199

Traditional cultural expression 55 131 162

UV map 241

Virtual try-on 21 23 36 37 53 108 152 158

Whole-garment assembly success 241

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 wp10. 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.