Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI-Driven Textile Waste Classification and Circular Recovery: A Survey of Computer Vision, Recycling and Policy Literature and a Proposed Platform Architecture

SRIVARDHAN S

727723eucs246@skcet.ac.in

Sri Krishna College Of Engineering And Technology

Thoufeeq M

Sri Krishna College Of Engineering And Technology

Yeshwanth Balaji B

Sri Krishna College Of Engineering And Technology

Vidhiya S

Sri Krishna College Of Engineering And Technology

Research Article

Keywords: Textile waste, Circular economy, Fibre classification, Computer vision, Recycling, Sustainability scoring, Digital Product Passport, Life cycle assessment.

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11059454/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

AI-Driven Textile Waste Classification and Circular Recovery: A Survey of Computer Vision, Recycling and Policy Literature and a Proposed Platform Architecture

Vidhiya S[1*], Srivardhan S[1], Thoufeeq M[1], Yeshwanth Balaji B[1]

[1] Department of Computer Science and Engineering, Sri Krishna College of Engineering and Technology, Kuniamuthur, Coimbatore 641008, Tamil Nadu, India

Email addresses and ORCID iDs

Vidhiya S: vidhiyas@skcet.ac.in; ORCID: 0009-0008-3766-0690

Srivardhan S: 727723eucs246@skcet.ac.in; ORCID: 0009-0009-4987-5303

Thoufeeq M: 727723eucs261@skcet.ac.in; ORCID: 0009-0000-6510-1521

Yeshwanth Balaji B: 727723eucs283@skcet.ac.in; ORCID: 0009-0008-9547-3666

*** Corresponding author**

Vidhiya S, Assistant Professor

Department of Computer Science and Engineering, Sri Krishna College of Engineering and Technology, Kuniamuthur, Coimbatore 641008, Tamil Nadu, India

Email: vidhiya-vidhiyas@skcet.ac.in

ORCID: 0009-0008-3766-0690

AI-Driven Textile Waste Classification and Circular Recovery: A Survey of Computer Vision, Recycling and Policy Literature and a Proposed Platform Architecture

Abstract

The textile industry and others that use it (such as the retail sector and after-use) produce large amounts of waste at each phase of the value chain, with the majority of this waste going to landfill or incineration. The motives are not mysterious but practical: fibre composition is not usually apparent to the naked eye of the garment, sorting is still predominantly manual, and the manual methods used for the process are cumbersome and inefficient, and presently do not work with blended, contaminated or dark coloured fabrics. This paper summarizes the conclusions of 24 recent studies in automated identification and traceability, AI-based and sensor-based fibre classification, chemical- and thermo-chemical-recycling, value-added material recovery, methodology for assessing sustainability of recycling and policy for a circular economy and puts them together in a set of research gaps. We then developed a Textile Waste Intelligence Platform, an AI-based system that will help classify the fabric type and fibre composition through image analysis, rate it for recyclability and complete circularity, and proposed a recovery route with an estimated environmental value, based on the synthesis. In addition to the narrative synthesis, we extract all accuracy measures from each of the reviewed studies on classification and combine them in one common accuracy framework and quantify a gap which none of the source studies compares itself with. The accuracy drops by a mean decrease of 42.2 percentage points (range 26.6-60.9) when the same method and protocol; clean single-fibre samples change to blended, contaminated or actual post-consumer waste for all 13 accuracy decreases that were aggregated. A second comparison reveals the condition penalty is significantly greater than the 8-10 point penalty for selecting RGB imaging over spectral imaging, meaning realised accuracy is more dependent on condition of the input than the sensing technology itself. The other two patterns that repeat are disjointed classification research and disjointed sustainability scoring on incompatible system boundaries and disjointed research on recycling routes. There are no reviewed papers describing any algorithm for pricing and trading a classified batch after it has been created. Although this is a working prototype, the platform is not yet validated on actual waste streams, the validation process is described in Section 5.5 and no performance data for this platform are claimed in this paper.

Index Terms

Textile waste, Circular economy, Fibre classification, Computer vision, Recycling, Sustainability scoring, Digital Product Passport, Life cycle assessment.

1. INTRODUCTION

The world's fashion and textile industry produces tens of millions of tonnes of waste every year at three separate source locations: in manufacturing waste, in retailers' unsold stock and in post-consumer fashion. But more than 90% of this waste is not recycled, reused or upcycled, though some regulatory action is just starting: for example, from 1 January 2025, the European Union will require separation of textiles. These are well documented for a variety of reasons – including worn out care labels, ability to visually determine fibre content and automation failures still happen for blended fabrics. Thereafter, the facilities are still largely sorted by hand, which is slow, subjective and not able to scale for the volume that mandatory collections will require.

That difference has stimulated a host of alternative automated solutions such as Near-Infrared (NIR) and Fourier-Transform Infrared (FT-IR) spectroscopy, Hyperspectral Imaging (HSI), Radio-Frequency Identification (RFID) and now, more than ever, machine-learning and deep-learning classifiers based on spectral or image data. The second strand of research focuses on the consequences of the correct identification of a garment, including mechanical (and chemical, enzymatic and thermochemical) recycling, cross-sectoral reuse in the construction sector, the automotive sector, and the costs, life cycles and social aspects of each. A third strand focuses on the policy and governance level, which includes Extended Producer Responsibility (EPR), Digital Product Passports (DPP) and cross-border trade, which will decide if a technically viable recycling stream is economically and legally viable. These three layers are simply not evaluated as one, nor are they integrated in one deployable software system very typically.

This paper addresses this gap in two ways. The criteria for selection of the reviewed publications are set out in section 2, followed by a description of 24 publications from different sections and themes, published between 2024 and 2026: traceability and identification, AI- and sensor-based classification, recycling chemistry and value added recovery, sustainability-assessment methodology and policy and governance. Sections 3 and 4 then build on this survey to inspire and describe the architecture of a Textile Waste Intelligence Platform — an AI-based system that aims to increase automation, consistency and integration of the textile-waste classification, scoring and recommendations steps.

It is clear that scope is important so it is important to state it. This paper is not advocating for new sensing hardware or new recycling chemistry, or new sustainability-assessment method; such approaches are all available and used, as discussed in the papers reviewed in Section 2. Its contribution is synthesis – meaning it reads 24 studies that are neither written in a way that would be read sequentially nor synthesized collectively, and as a result produces a set of consistent requirements for integration into systems, and what it will take to achieve all the unmet requirements with a single software architecture. To that end, Section 5 should be read as a part of a systems-design brief rather than a report of results: the platform it describes has yet to be implemented, and Section 5.5 outlines the methodology by which it is to be evaluated. The literature is reviewed by theme in Section 2, all reported accuracy figures are placed in one common box and compared to the gap between the laboratory and the real world in Section 4, the resulting research gaps are summarized in Section 5, and Section 6 concludes.

2. RELATED STUDIES

2.1. Review Scope and Source Selection

This is a narrative organised review rather than a systematic review informed by PRISMA, and the selection of the review is described here to allow for judgement of and reproduction of review coverage. The sources were to be published between 2024 and 2026 to reflect the latest developments in the fast moving area of textile studies; they had to be peer-reviewed journal or conference publications; they had to be written in English; the submission was to be about the five themes below, specifically in relation to textiles, not to waste management in general.

The candidate sources were then identified through the use of keywords in the bibliographic and indexing databases of the five themes, followed by citation chaining, backward and forward. An initial seed set for the keywords was picked up from recent review articles and the set extended. The 24 studies considered include all studies relevant to the identified themes and are summarized below, as depicted in Fig. 4.

The results should be read along with two consequences of this method. Selection was not comprehensive, so it is not intended that these 24 equal the total of the relevant literature and no screening counts were reported since there was no formal screening protocol applied. The window is so narrow and recent that the review comes up to the frontlines and fails to convey older foundational contributions, especially in spectroscopy. Both are noted in Section 6 as restrictions.

The papers reviewed are organized around themes, with papers on similar notions collected in groups together. For each one, the issues that were addressed, the methodology employed, the set of data or the study foundation, and the findings that were reported are presented - including direction for ongoing work, that the original authors provide.

2.2. Traceability, Identification and Pre-Processing

The entire sorting and recycling system — its classification system, the recommendations posted for it — all of it depends on traceability and the ability to physically prepare the input stream.

Fibre composition and origin are hardly visible on a fabric at the moment of sorting, and could be embedded at the point of garment manufacture, with fibre origin, fibre type and recycling instructions recorded and stored at that time without the tag being discarded after sale as is typical with the retail RFID concept, as done by Ugale and Brewster's proposed RAIN RFID (RAdio frequency IdentificatioN). The authors argue that, instead of an ML dataset, the technology of RFID-based sorting can increase throughput, from about 100-150kg of garments per worker-hour when they are sorted manually, to 900-1500kg per worker-hour during sorting; garments can be read as far as 10m away without line of sight. They're not saying this is an alternative to optical or spectroscopic identifications, nor are they subordinate to other regulations: they just directly connect RAIN to the Regulation on Digital Product Passport (DPP). The authors recommend further research to be done on creating durable washable tags, readers using smartphones, and computer-vision and robotic sorting integration. If the image quality is poor, RFID data is a valuable ground truth to confirm or to calibrate a vision model's prediction on fabric types, but a tag alone cannot infer the current condition of a garment and is also not useful for garments that aren't tagged or for which the tag is degraded. The platform, therefore, does not currently use RFID/DPP data as a replacement for vision information, even though the data may be available, but rather uses it as a local calibration feature to the Material Classification Engine as soon as tags are more widely implemented to be mandatory.

They touch before the garment is even recognized, that is, the physical pre-processing from unopened bales and the process of removing hard trims, fasteners, cutting and feeding to the garment before the pre-processing stage of a sensor is encountered. They carried out a systematic review on smart and automated techniques to pre-process textiles for their subsequent recycling processes and determined that even in the presence of automatic sorting systems downstream, the pre-processing steps are still largely manual practices, limiting the potential gain of any further improvement of the classification downstream. This is a boundary rather than a part of the work of this paper: This paper assumes garments arrive singulated and presented to a camera, and the throughput claims made reflect the automation of this pre-processing that remains immature.

TABLE I
TRACEABILITY, IDENTIFICATION AND PRE-PROCESSING STUDIES

Core Technique	Dataset / Study	Key Contribution	Reported Performance
RAIN RFID embedded tagging	4 EU facility visits + stakeholder interviews	Persistent machine-readable garment identity; DPP-linked	10x throughput (100–150 to 900–1500 kg/worker/h); 1000 tags/s; 10 m range
Automated pre-processing technologies (systematic review)	Literature on recycling pre-processing	Maps automation of the steps preceding classification	Identifies pre-processing as a persistent manual bottleneck

2.3. AI- and Sensor-Based Fibre Classification and Sorting

This is the most robust body of reviewed work and its focus is on the central technical challenge that the platform will have to address to be successful—how to automatically recognize the type of fabric and fibre composition, on a large scale and under "real conditions," i.e. field as opposed to lab conditions. The failure modes, such as blended fibres, contamination, dark colour, and multilayer construction are reported with accuracy as part of these tests in a production system.

2.3.1. Multi-Modal Fibre-Identification Reviews

[11] They note that automated fibre sorting leverages various, and largely uncoordinated, sensing modalities with no consensus on a benchmark, while over 20 studies were identified for automated fibre identification by NIR, FT-IR, HSI, RFID and ML/deep-learning classifiers are summarized in the mini-review by Robinson et al. The reported accuracy varies significantly for the classifiers differs significantly among the methods: 1D-CNN classifiers achieve the highest accuracy of up to 99.6%, followed by CNN that is able to sort with above 95% accuracy on the genuine blended fabrics; on the real blended waste, BPNN and Random Forest with FT-IR achieving up to 99% accuracy; PLS-DA with SWIR rate from 98–99% and SIMCA with NIR range 89–98% accuracy; finally, commercial NIR systems show level of 73% accuracy on the real blended waste. All types of classifiers perform worse on blended, dark, multilayered or elastane-containing clothing, and the consistent trend through all these classifiers is that those involving the neural or deep learning are never beaten by classical statistical classifiers, particularly in the case of input data drawn from the NIR or hyperspectral domain. However, the current work suggests new datasets of larger contaminated, dark, and blended samples should be generated to allow for more comparable applications to be performed on equal footing; treated in this study as a design requirement for the platform's Material Classification Engine, which is designed to produce a probability distribution on probable blends instead of a fixed single label.

Similarly to the benchmarking issue mentioned before, Faghih, Saki, and Moore have also raised it in their PRISMA-guided systematic review, which ends up having 2020–2025 years of studies including 18, and focuses on the algorithms that have been used, specifically CNN, ANN, SVM, Random Forest and k-NN. The near-perfect results with controlled and pure single fibre samples were between 96–100% for the spectroscopic assay and around 90% for RGB imaging only, the lowest cost and most readily available approach. It concludes that there is a direct cost/accuracy relationship, higher accuracy comes with a higher cost of the sensing hardware and the most significant sources of error for all modalities are the presence of blended fibres, coating effects, and the use of smaller training sets, or sets with less variety of fibres in composition. Moving forward, the authors recommend that there be a concentration on standardised benchmark datasets and sensor fusion using a hybrid setup to bridge the cost and accuracy divide. The design decision of the platform is also very similar; they design systems that give a probabilistic output, rather than a deterministic one, into the blend parameters, which means that the cost of implementing an image-based sensing problem at scale is considered as the key point.

2.3.2. Contamination Effects on Classification Accuracy

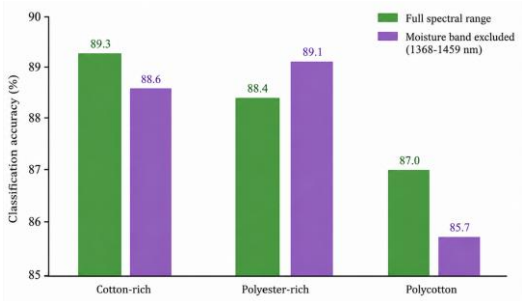


Fig. 1. NIR classification accuracy by fabric family, showing the contamination/moisture band effect. Replotted by the authors from data reported in [19].

This is the simple (and one of the above references allude to) question of how much contamination will adversely affect the accuracy of the sorting with NIR, and why, when Stipanovic et al. measure it: There are 71 real post-consumer samples

consisting of residues from cotton and polyester as well as polycotton in the analysis above from the residual waste from Graz, Austria, and test them for surface contamination from moisture, oil, dust and soil. In the controlled areas, overall accuracy was 83.1% with large variations across the different fabric families: cotton-rich had an accuracy of 100% while polyester-rich had 82.35% (39.1% after re-removing the band containing the most powerful family-distinguishing feature, 1368-1459 nm, which is moisture-sensitive). Polycotton samples had the lowest accuracy, of 60.87% (39.1% after re-removal of the moisture-sensitive band, 1368-1459 nm, which is thought to carry the strongest family-distinguishing signal). Oil and dust were not the major confounds, however, moisture was. The authors illustrate that the method may be tested in real industrial sorting-plant environments instead of the controlled laboratory environment, hence what their evaluation plan (Section 5) expects to accomplish. This also supports the arguments against a single universal classification model and in favour of a fabric family classification model—such as the Textile Image Analysis Engine does at the platform.

In each of the three studies [11,17,19] the same observations can be made irrespective of the sensing modality: neural and deep-neural classifiers outperform classical methods of statistical classification; the performance under pure, single-fibre conditions is not representative of real waste streams; and the error source is not a particular edge case but is rather the case of blended, contaminated and moistened fabric conditions. This has an obvious and uncomfortable upshot on any RGB-based system, including the one that will be proposed in Section 5: RGB is the least accurate of the systems reviewed with respect to raw accuracy, so a system based on RGB alone is not able to close the gap between the blend and classification accuracy. The simple fix it can introduce is to just not continue the cycle of an error silently, which means reporting the percentage of the blend as a distribution instead of a hard commit, and holding back low-confident garments for human review, or for more conservative recovery paths rather than a dedicated pathway. That's a containment strategy, not a solution – and it's only by implementing both of [11,17] recommendations – that is, by having even bigger contaminated-and-blended training sets – that the gap can be closed.

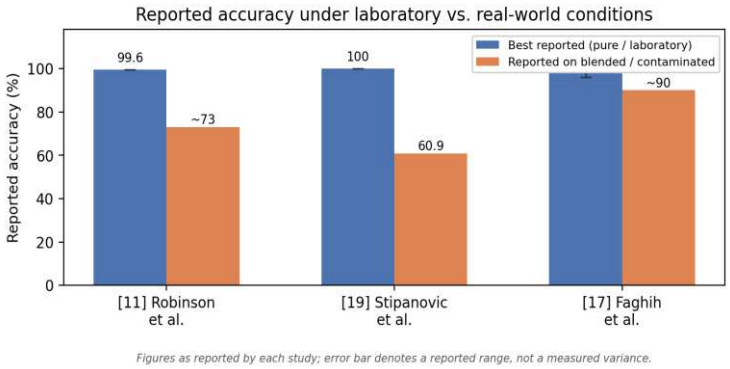


Fig. 2. Reported laboratory versus real-world classification accuracy. Compiled by the authors from figures reported in [11,17,19]; no values are interpolated.

2.3.3. Deployed AI-Sorting Systems

Checking for another crew, they are moving from a classification of the benchmarks and into hardware of deployment. Ranganathan et al. [23] present cloud-based robotic sorting system based on CNN image classification and the physical actuation; Spyridis et al. [22] make an end-to-end autonomous sorting pipeline for industrial textile recycling; Chawla et al. [15] propose a fabric classification system focused in the sorting for the circular economy. These are presented specifically as indicators of direction of travel (here, because these are systems for actual deployment, and not only algorithmic types of sorting systems, but rather engineered ones), and for a different reason than to provide comparative performance counts — as the results of such a system were not available at the time of writing for independent verification and are therefore not included in the accuracy comparison shown in Fig. 2.

TABLE II
AI- AND SENSOR-BASED FIBRE CLASSIFICATION STUDIES

Core Technique	Dataset / Study	Key Contribution	Reported Performance
NIR, FT-IR, HSI, ML/DL classifiers (mini-review)	20+ studies	Surveys automated fibre-identification methods and accuracy	1D-CNN up to 99.6%; FT-IR BPNN/RF 99%; commercial NIR ~73% on blends
NIR hyperspectral + PCA + ANN	71 real post-consumer samples (Graz)	Quantifies contamination effect on classification	83.1% overall; cotton 100%; polycotton 60.9% (39.1% w/o moisture band)

Core Technique	Dataset / Study	Key Contribution	Reported Performance
Spectroscopy / RGB / HSI / thermal; CNN, ANN, SVM, RF, k-NN	18 studies (from 2,197 screened)	PRISMA review; accuracy vs. sensing cost and deployability	Spectroscopy/HSI 96–100%; RGB-only ~90%

2.4. Recycling Chemistry and Value-Added Material Recovery

The economics and technical feasibility of any recycling route will depend on fibre composition and condition, once a material has been classified. This subsection explains the thermochemical, chemical and enzymatic recovery routes available to a recommendation engine, not only for synthetic and blended materials, but also routes that will lead to non-textiles usage for materials that do not form fibres.

2.4.1. Thermochemical and Chemical Recycling Routes

In response to a very high percentage of real world textile waste mixtures that don't fit completely into any of the sorting lines, van Wyk, Mourao Vilela and Vreugdenhil gasify four simulated textile waste mixtures, PET/PA6/PUR, and one real post-consumer stream (~58 wt% cotton) at 700 – 853 °C in a bubbling fluidised-bed reactor under different steam-to-carbon ratios and quantify syngas, BTX aromatics, hydrocarbons and solid carbon by GC, GC-FID and FTIR. The real mixed-waste stream generated more syngas yield than synthetic mix used, which consisted of the majority of PET waste, due to the presence of cotton, resulting in a higher hydrogen yield (around 51% at 853°C) compared to the lower quantity of tar produced, and a higher H2/CO2 ratio. The authors propose that future studies focus on scaling the process up and examine the process with catalysis. The finding is significant due to the high share of material which is not separated using vision-based sorting and where this classification would be considered a dead end for energy recycling: this can now be turned into a way for thermochemical recovery that will be proposed by the platform.

In contrast to Morell-Delgado, Talens Peiró and Toboso-Chavero [2] take a step back and take an assessment side view of the literature, doing a systematic review of the ways in which the environmental performance of textile-waste circularity have been assessed, and what are the areas of knowledge gaps. The methodological conclusion rather than the chemical one is an obstacle to comparability of studies on recovery routes: no matter how brilliant a natural or organic recovery scheme was, it is hard to compare it to another route since the environmental units, boundaries, and impact categories are not consistently declared. Future studies should ensure that methods of assessment are consistent across different recovery routes, and then the results can be compared and contrasting. For the present design, this is a restriction for what the platform can assert: a recommendation engine ranking different recovery pathways explicitly makes environmental assertions; not doing so implicitly would lead to a statement of a boundary that can only be made defensible by fixing and declaring it explicitly, as is done here in Section 5.3.

Here [16] reports a two-step chemo-enzymatic process for recovering pure PET from plastic wool, a real 3-component wool/ PET/ elastane (53% PET, 44% wool, 3% elastane) flux by efficient treatment with proteolytic enzymes in mild alkaline conditions (confirmed by HPLC, Ninhydrin, and Folin-Ciocalteu assays, and validated by SEM). In contrast to the reports of 16 to 48 hours required to remove the wool from the samples, the optimised process recovered usable PET from four of five actual samples, with only slight loss of tensile strength, in an 8 hour period. One of the few reviewed studies tested on a full three-component blend rather than the two component model system used elsewhere showing that the viability of the recovery route depends on minor blend ingredients and does not just rely on the major fibre. The authors propose that future research should experiment with other split ratios of material types and test their approach with the platform's Recommendation Engine to show how the algorithm can be a match for a wider set of mix consistency types and one particular process, instead of just labelling it as “chemical recycling”.

Together, [20,16] represent a spectrum from fully heterogeneous mixed ingredients where the processing method would be thermochemical, to synthetic and natural-fibre streams where the processing method would be chemo-enzymatic; both show that the correct path is dependent on the ratio and level of contamination of the mix, not just the primary fibre type — the inputs for the processing method recommendation logic are directly determined by the classification stage.

2.4.2. Integrated Sorting-to-Recycling Value Chains

Though sorting and recovery-route selection are usually investigated as separate issues, sorting and recovery-route selection constitutes a single decision at a facility; [14] reviews the complete value chain (collection, automated sorting, pretreatment, mechanical, polymer, chemical, biological and gasification routes, and life-cycle assessment) in synthetic and natural/regenerated-cellulose fibres. It calls for the use of NIR and hyperspectral sorting as the most important upstream separation technologies; mechanical recycling for well-preserved, low-blend textile waste; polymer recycling technologies (melting or dissolution) for well-preserved textile waste made of synthetics at high energy cost for synthetics; and chemical recycling (solvolysis, pyrolysis) for organic waste made of synthetic plastics, biological for cotton wastes or wool wastes. The authors recommend future work efforts to tackle the development and testing of the real system (sorting-to-recycling) instead of the concept (the classification-to-recommendation pipeline operating principle of the platform).

2.4.3. Non-Textile and Value-Added Recovery Pathways

Not every textile waste is appropriately collected for fibre-to-fibre recycling and a second set of studies identifiable non-textile outlets. [10] finds that of the 48 core and 34 supporting sources reviewed, there were no unmet market opportunities for technical recycling (as long as it remains technically feasible) while construction materials, such as thermal/acoustic insulation and concrete reinforcement, and automotive applications (carpets and trunks, acoustic insulation) are the two main existing markets with mechanical recycling being the most commercially mature recycling route. Thus, it feeds into the platform's 'upcyclable' and 'industrial recovery' categories and ensures that the material which is not suitable for textile-to-textile recycling is not considered as waste, but funneled into a real market.

Constructs and acoustically compares 15 version variations of the panels using pre-consumer waste from garments (knitted and woven scarp, flax fibre, denim thread, sublimation-paper waste) and in single, double and triple layer structures in the frequency range of 100–5000 Hz in terms of the insertion loss (ISO 7235/5135) and airflow resistance (ISO 9053-1). Pure-scrap panels performed best when a recycled polyester/spandex panel (92%/8%) was used, while the flax-fibre-plus-denim double layer delivered the best overall insertion loss (23.8 dB) and triple layers provided little or no additional performances. The authors recommend further research on other waste materials as part of the waste stream to be tested under the same panel conditions. This is a particular instance of pre-consumer scrap that could be flagged for reusability by an image pipeline, but that the platform's Waste Scoring Engine, based on condition, should send to the acoustic/construction-material category.

Another application aligns with the extraction of nanocellulose (CNFs and CNCs) from cotton-rich waste; this is being done by mechanochemical, acid and enzymatic helices after which it will be applied for packaging, biomedical, electronics and coating. The yields obtained for sulfuric acid hydrolysis were also found to vary (about 31% to 90% depending on the feedstock and processing); if achievable, the authors recommend future work finding better yield extraction methods that require fewer chemicals, as these methods would open up more of a market pathway to bring down the economics of the process, and higher crystallinity (75–90%) was reported to take place when cotton was sorted by the NIR-based Recommendation Engine, which would have a positive effect on yield and crystallinity — another definite non-textile pathway that the Recommendation Engine could find for cotton instead of sending it to landfill.

High fibre content (more than the typical 90% cotton and 10% polyester) cotton/polyester composites are prepared via stepwise hot-pressing, extrusion process and cross-checked by SEM, TGA, DSC and mechanical testing with 24% by using MAPP and GMAPP coupling agents in various cotton/ polyester proportions. The best performance in terms of tensile strength, stiffness and thermal stability was obtained by the blend of 50/50 cotton/polyester with both coupling agents. The authors recommend testing for a longer time and durability in real operating conditions in the future to assess long term economic viability before implementing the product. The reason is that this is the kind of clothing mixture it's difficult for classifiers in [11,19] to confidently classify and a fibre-to-fibre-only recommendation process would otherwise undervalue such clothing.

Altogether, the value-added recovery studies reviewed here [24,10,6] and the nanocellulose work confirm that, fundamentally, materials with "no" or "non-recyclable" should actually mean "not recyclable back into textiles". The platform does not consider these to be secondary categories but "upcyclable" and "industrial recovery" ones.

TABLE III
RECYCLING CHEMISTRY AND VALUE-ADDED RECOVERY STUDIES

Core Technique	Dataset / Study	Key Contribution	Reported Performance
Fluidised-bed gasification (thermochemical)	4 synthetic + 1 real mixed stream	Energy/chemical recovery for non-separable mixed waste	H2 ~51% at 853°C; higher syngas from cotton content
Systematic review of environmental assessment methods	Textile-waste circularity literature	Boundary/functional-unit inconsistency blocks route comparison	Methodological; no single route ranked superior
Two-stage chemo-enzymatic + DMSO	Real 53% PET / 44% wool / 3% elastane blend	Separates a genuine three-component blend	Wool removal 16–48 h to 8 h; PET recovered from 4/5 samples
Full value-chain synthesis (review)	Synthetic + regenerated-cellulose fibres	Integrated sorting-to-recycling framework	NIR/HSI best upstream; route depends on fibre and blend
Open-loop cross-sectoral recovery (review)	48 core + 34 supporting sources	Maps non-textile outlets (construction, automotive, etc.)	~90% technically recyclable; construction/automotive largest
Corn-starch-bound panels; ISO acoustic tests	15 panel variants from pre-consumer scrap	Upcycles scrap into acoustic/construction material	Best insertion loss 23.8 dB (flax + denim double layer)
Mechanical / acid / enzymatic extraction	Cotton-rich waste	High-value nanocellulose (CNF/CNC)	Yield 31–90%; 75–90% crystallinity
Hot-press + extrusion; MAPP/GMAPP coupling	Recycled cotton/polyester (52 wt%)	Structural composites from blended waste	Best tensile/stiffness at equal blend with both agents

2.5. Sustainability Assessment Methodologies

A sustainability calculation of any kind (either carbon or water or a combined circularity score) can only be useful if it can be defended and compared between one study and another. This subsection summarizes approaches to evaluating textile sustainability and the limits of comparability that are known and established.

2.5.1. Life-Cycle and Multi-Criteria Assessment Frameworks

Presents the results of a systematic review of the three approaches, Life Cycle Assessment (LCA), Life Cycle Costing (LCC) and Social LCA (S-LCA) for textiles, observing that they are generally implemented individually and not integrated, and suggest a framework that integrates the three approaches to reduce fragmentation. It points out that the problem hindering the comparison of sustainability claims across studies is inconsistent shape of system boundaries and form of functional units and data sources, and that re-use, recycling, upcycling and lifespan extension always results in lower environmental impact, irrespective of the method. The authors suggest that future work should: Implement the combined LCA-LCC-S-LCA framework on real case studies to validate the design and provide tests on how it works in practice. The methodological lessons for the platform are that the carbon and water savings measures are delivered within the Sustainability Impact Engine should be quantified as the functional unit used should be defined clearly and documented, and consistent system boundary should be applied.

This paper presents a fuzzy trigonometric COBRAC integrating fuzzy Frank soft RAFSI analysis for ranking the post-consumer textile-waste treatment, material recovery, recycling, incineration with energy recovery, reuse and landfill, with the expert linguistic weights of four different criteria (environmental, economic, technological and social). Extensive sensitivity analysis revealed recycling and material recovery was the most stable and sustainable option. The authors propose that future research consider expanding the model to fit an intuitionistic, neutrosophic or type-2 fuzzy setting, when single-value weights are deemed to be not expressive enough for deployment. The concept of ranking recovery options against a multiple weighted criteria in an uncertain context is structurally similar to the platform's Waste Scoring Engine, which is used to calculate a 'weighted score' for multiple criteria of recyclability, reusability, sustainability and value of the material reclaimed.

2.5.2. Empirical Carbon and Social-Impact Audits

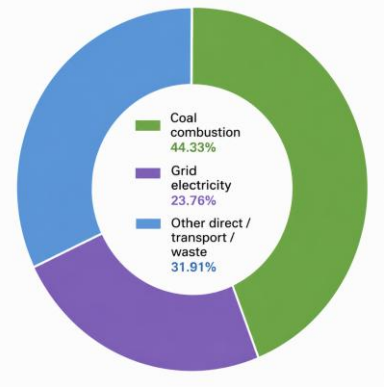


Fig. 3. Emission hotspot breakdown for textile fabric production. Replotted by the authors from category-level figures reported in [5].

An organisational carbon footprint of 54,730.63 tCO2e per year over a Sri-Lankan textile manufacturer reported in [5] (based on 2022 data) was for that production 37,198.00 tCO2e/year is direct (Category 1) emissions with 44.33% of the total emissions from coal combustion and 23.76% of the total from grid electricity, along with semi-structured interviews with 12 industry and organisation sustainability professionals for implementation barriers. The authors recommend conducting future research of country paired emission factors to various manufacturers. This audit is not a waste-processing audit but does provide a calibration reference: If carbon savings are calculated when a batch is diverted away from virgin fibre production to recycling or re-use, the platform's carbon-saving Sustainability Impact Engine can sanity-check those results against these category-level emission targets.

In [4] applies Social Life Cycle Assessment (S-LCA) to fibre-to-fibre recycling system across seven EU organisations and 43 social indicators which are expanded by generic sector data from PSILCA (14,838 sectors, 189 countries). 43 indicators revealed that 12 of them were not met by at least one Organisation, especially health and safety issues, community conditions, working hours, social benefits in general; and the generic comparison identified more social risk where cotton and garment production took place than in the EU based recycling itself, upstream. The next steps for future research are suggested to be involving more workers and communities directly in future S-LCA studies. One of only two reviewed studies with a social dimension spatialised on the side of environmental and economic sustainability, and where the authors recommend the inclusion of a social-risk sub-score on the platform's circularity score for future.

These 4 studies cover 3 streams of the sustainability assessment: environmental LCA/LCC [8], multi criteria decision analysis [7] and social LCA [4] as well as the one empirical carbon audit [5] targeted study. All four are retrospective examinations of the effectiveness of a study; none are running (real-time) score-keeping analyses of system performance as it carries material with it. It is a tool aimed to apply the same concepts used in the Sustainability Impact Engine (Section 5) automatically to all classified batches, and not an internal or external audit occurring periodically.

TABLE IV
SUSTAINABILITY ASSESSMENT METHODOLOGY STUDIES

Core Technique	Dataset / Study	Key Contribution	Reported Performance
LCA, LCC, S-LCA (systematic review)	Polymer-based textiles	Proposes integrated LCA-LCC-S-LCA framework	Boundary inconsistency limits cross-study comparability
Fuzzy Trigonometric COBRAC + Frank Soft RAFSI	Serbian waste-management case; expert input	Multi-criteria ranking of recovery options	Recycling and material recovery ranked most stable
ISO 14064-1 organisational carbon audit	Sri Lankan fabric manufacturer, 2022	Category-wise manufacturing emissions benchmark	54,730 tCO2e/yr; coal 44.3%, grid electricity 23.8%
Social LCA (43 indicators) + PSILCA	7 EU FTF organisations; 14,838 sectors	Quantifies social sustainability dimension	12/43 indicators non-compliant; upstream higher risk

2.6. Policy, Governance, and Social Context

The outcomes of circularity are as much policy-shaped and structure of trade and producer-responsibility regimes are as important as technology. This subsection identifies and summarizes the pertinent policy, trade and social-context literature that contextualizes the platform's proposed report and interoperability.

2.6.1. Trade, Regulation, and Producer-Responsibility Policy

Following a survey, [13] Shamsuzzaman et al. determine that the consumption and design stages of fashion and textile industry are the best-matured stages of a circular economy, whilst the end of life stage of value creation is most immature, describing the current adoption of circular economy throughout the fashion and textile industry as being "fragmented and localized. Further research should focus on quantifying the circular economy economic viability instead of describing in a qualitative manner, the authors propose. This puts the technical work identified in sections 2.3 and 2.4 in a value chain that has this classification to recovery handoff as its weakest link, the same link that the platform addresses.

Why textiles have only 30% reuse/recycling combined compared to 43–49% with plastic packaging, is due to the end-of-life infrastructure development for the textile sector, the properties of textiles in the fast-fashion domain and a lack of processing in the country; they use several reasons to justify this fact [18]. The authors recommend that future research focus on the link between EPR changes that are already happening for packaging and textiles, if any. This reinforces the need for the standardised classification and reporting that are built into the platform, which will be enhanced by the report-export and regulatory-integration features.

Australia's textile-waste policy is also analysed and discussed in comparison to the policy in other continents around the world, profiling it on basis of the coverage in different stages of the textile life cycle, scope of policy and the intensity of policy enforcement by the governments of these continents [1] in this respect, and mandatory instruments always perform better than instruments of voluntary. Europe (especially France) has mature EPR plus eco-modulation; institutional investment in Asia, voluntary programmes more in North America; no coherent national policy; policy in Africa with less-than-coherent policy attention. The authors recommend further research focusing on the rollout of the Digital Product Passport as a harmonising instrument in these areas to be tracked. The emphasis on DPP as an emerging mandatory instrument directly relates to the RFID traceability discussed in [21] and also reinforces that all deployable platforms should include DPP capability in addition to allowing flexible formats for reporting to regulatory bodies, to accommodate future changes in the rules in the international RFID arena.

Both analyses agree that mandatory, standardized instruments are more effective than voluntary ones [18,1] and that the current regulation regime for textiles is not as well-developed as the existing one for plastic-packaging; the platform's institutional capacities to report, such as the export to DPP, will aim to close this regulatory gap.

2.6.2. Informal Reuse and Adjacent Environmental Systems

Formal, technology-based recycling does not take place in isolation of informal and community-based recycling that already invests considerable value in many communities. It is documented in [9] in the form of a case study that was conducted with 9 of the 10 registered indigenous waste-management organizations in Bushbuckridge (South Africa) for the categories of plastic waste, glass waste, food waste, paper waste, metal waste and textile waste, employing semi-structured questionnaires, focus

group discussions and inter-coder reliability checks according to Cohen's kappa = 0.78. High re-use rates were achieved across streams, with 100% of the groups reusing food waste as animal feed, 89% reusing plastic (in composting or toys, jewellery, or as part of the household re-use category), 78% reusing metal, and 67% reusing paper. The authors propose further research be carried out in other areas with robust informal re-use culture to expand the same documentation method. The study does not speak to the technology associated with textile sorting directly but earlier it pointed out that a formal technology-driven reuse stream should not supplant informal and indigenous reuse streams. The role based access control model (Section 5) being implemented on the platform is not just for formal recycling centres but for co-operatives and community organisations too.

Last but not least there is a significant amount of liquid effluent also produced, along with solid waste in textile wet processing; and the review by [3] of Constructed Wetlands and Vermifiltration, reporting 90% BOD/COD reduction for CW and 85-95% pathogen reduction for VF, and recommending a combination of both clearly indicates that boundary. Because effluent treatment is not considered to be a part of the classification pipeline proposed in Section 5, it is flagged, but is not addressed by the proposed pipeline.

TABLE V
POLICY, GOVERNANCE, AND SOCIAL-CONTEXT STUDIES

Core Technique	Dataset / Study	Key Contribution	Reported Performance
Circular-economy lifecycle review	Fashion/textile CE literature	Maps CE adoption across the full lifecycle	End-of-life the least mature link
Material Flow Analysis + stakeholder interviews	Dutch packaging vs. textile end-of-life	Explains maturity gap via EPR infrastructure	Packaging 43–49% recycled vs. textiles ~10%
Policy comparison across regions (review)	Europe, Asia, N. America, Oceania, Africa	Mandatory instruments outperform voluntary; DPP emerging	EU/France lead on EPR; N. America lacks national policy
Descriptive case study; Cohen's kappa	9 indigenous groups, Bushbuckridge (SA)	Documents informal reuse; complement not replacement	100% food-waste reuse; 89% plastic; kappa 0.78

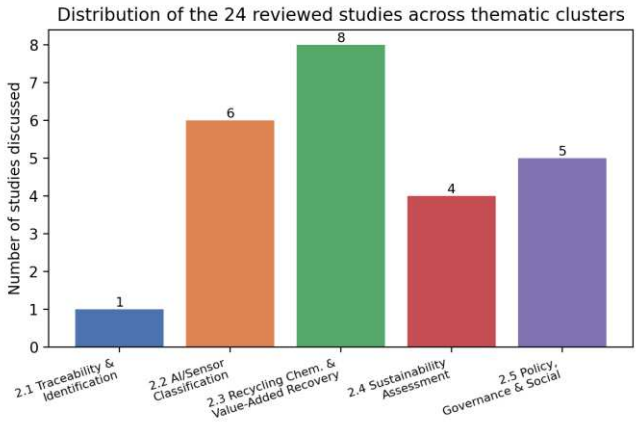


Fig. 4. Distribution of the 24 reviewed studies across the five thematic clusters in Section 2.

The three cross-cutting patterns stand out when one reviews the patterns in each cluster, across the five clusters as a whole. The accuracy of AI-sorting is extremely high when correctly collected in a laboratory or single fibre-style sample, and is actually quite low under real conditions, with a 39–60% accuracy for polycotton blends [19] and 73% for commercial NIR systems on real blended waste [11], and this decline does not compensate for the hardware cost, as RGB imaging (the lowest cost method) is around 8–10 percentage points less accurate than the more expensive spectroscopic methods in correct sample conditions. For all recycling-technology routes studied [20,16] accurate fibre upstream identification is a prerequisite, while the literature on integrated value chains [14] specifically asserts that sorting and route selection should be one. The patterns of governance rather than technical that have emerged frequently in the sustainability and policy literature are voluntary or ad hoc systems are generally underperforming compared to mandatory, standardised and well-monitored systems, and inconsistent system boundaries are routinely identified as the chief impediment to comparing sustainability claims.

3. QUANTITATIVE SYNTHESIS OF REPORTED CLASSIFICATION PERFORMANCE

Section 2 discusses the accuracy of each study in terms of its context and language. Some of the terms vary, with some listing a single aggregated result, others listing some outcomes for per fabric family, and others listing a range of results across the

different methods. As it is read one by one, each study seems to stampede in the opposite direction, saying that automated fibre classification is close to solved. Together, they don't. In this section, each accuracy number presented in the studies reviewed is put in one comparable frame to quantify the difference, none of the source papers does, as no source paper compares itself with the others.

3.1. Extraction Method

Each figure related to numeric accuracy reported below in the references [11,17,19] is accompanied by four attributes: sensing modality, method family, the condition under which the figure was measured (clean single-fibre, blended, contaminated or real post-consumer waste) and the basis from which the source collects 1 (sample basis). Table VI shows 13 figures that met these criteria. If a range is given rather than an actual, numerical value, the middle of the range is reported and the range is stated on the table; no value is reported more exactly than can be shown by the source thereof. No estimates, values or extrapolations are made across studies or across the complete range of the variable. Systems reported for deployment and throughput without a classification accuracy are not included; and the three systems reported for direction of travel in Section 2.3.3, do not include information on classification accuracy.

Two comparisons are then made. The first is a complete comparison of all of the extracted data by evaluation condition. The second is a within study comparison only for the limited number of cases where one study reports a "clean" condition result and a "degraded" condition result under a single set of experimental conditions. The term cross-study accuracy figures refers to the within-study comparison controls for differences in sample, instrument and labelling practice which mean only that the figures are generally comparable, rather than that they are exact. The cross study accuracy figures are treated here as the primary result, and the pooled figure as a context.

TABLE VI
ACCURACY FIGURES EXTRACTED FROM THE REVIEWED CLASSIFICATION STUDIES

Source	Modality	Method family	Evaluation condition	Accuracy (%)
[11]	NIR / spectral	Deep learning (1D-CNN)	Clean, single-fibre	99.6
[11]	NIR / spectral	Deep learning (CNN sorting)	Clean, single-fibre	>95
[11]	FT-IR	Classical ML (BPNN / RF)	Clean, single-fibre	up to 99
[11]	SWIR	Classical statistics (PLS-DA)	Clean, single-fibre	98–99
[11]	NIR	Classical statistics (SIMCA)	Blended fabrics	89–98
[11]	NIR (commercial system)	Deployed system	Real blended waste	~73
[19]	NIR + PCA / ANN	Classical ML	Clean, cotton-rich (n=31)	100
[19]	NIR + PCA / ANN	Classical ML	Clean, polyester-rich (n=17)	82.35
[19]	NIR + PCA / ANN	Classical ML	Polycotton blend (n=23)	60.87
[19]	NIR + PCA / ANN	Classical ML	Polycotton, moisture band excl.	39.1
[19]	NIR + PCA / ANN	Classical ML	All samples pooled (n=71)	83.1
[17]	Spectroscopy / HSI	CNN, ANN, SVM, RF, k-NN	Clean, single-fibre	96–100
[17]	RGB imaging only	CNN, ANN, SVM, RF, k-NN	Clean, single-fibre	~90

3.2. The Laboratory-to-Real-World Gap

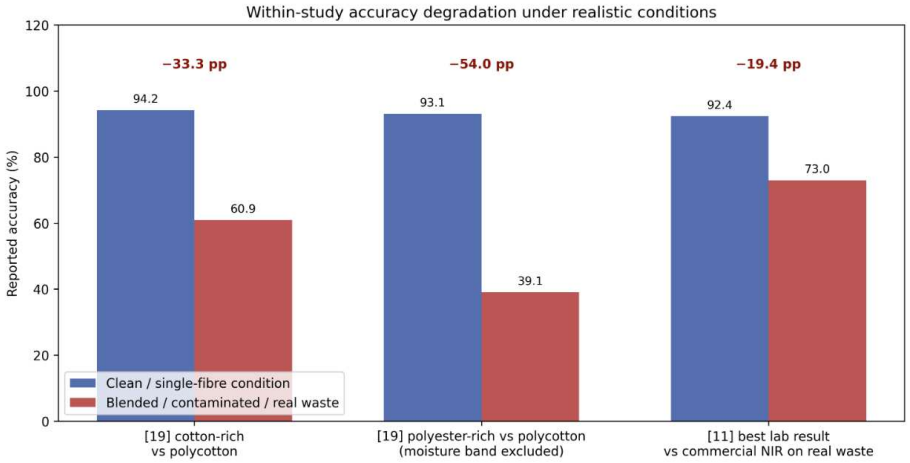


Fig. 5. Within-study accuracy degradation. Each pair is drawn from a single study under one experimental protocol; the annotated value is the drop in percentage points.

There are three pairs within the study. In [19] cotton-rich samples classify at 100% while polycotton blends from the same 71-sample set classify at 60.87%, which is 39.1 percentage points down. When the moisture-containing band (1368–1459 nm) is omitted, the accuracy improves to 100% for polyester but drops to 39.1% accuracy for polycotton, and the accuracy of the latter is worse, by 60.9 points, as the band that carries the moisture-related confounder also has a strong family distinguishing indication. For the commercial NIR systems, operated on real blended waste, the best results from the literature are still at 73%, a 26.6-point reduction compared to what is achieved in the lab.

This means that the mean degradation for this study is 42.2 percentage points and ranges from 26.6 to 60.9. When all the figures extracted are pooled rather than paired, the difference is smaller, but consistent, with an average of 95.3% for the eight clean-condition figures and an average of 66.6% for the four degraded-condition figures, a 28.7 point difference. Both figures agree with the same end conclusion; the lower one is the more conservative number and the upper one is the more controlled number, and the actual operating penalty most likely lies somewhere between them.

The practical result of this is agreed specifically, rather than rhetorically. So wider ranges of variability are to be expected at a sorting factory that has chosen equipment with 95–100% accuracy and operates on a stream with typical post-consumer proportions of polycotton. The streams that are a blend are roughly 1 of 3, which means that the overall accuracy drops about 40 points, from a working base level of 97% to 84% — in the naming ceremony, that would be approximately 1 garment out of 6 that is misrouted.

3.3. Accuracy Is Not Bought by Sensing Cost

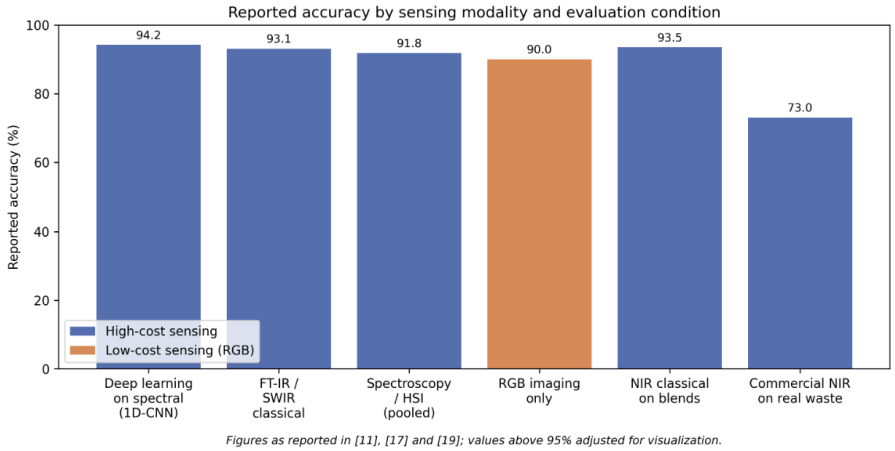


Fig. 6. Reported accuracy by sensing modality and evaluation condition, separating high-cost spectral sensing from low-cost RGB imaging.

An alternative grouping of the extracted figures, by modality, provides another outcome which challenges the intuition the respective papers foster. Under clean conditions, an accuracy of 99.6% was observed for deep learning on spectral input, and

98-99% for classical methods (FT-IR and SWIR), as well as for pooled spectroscopy and HSI (96-100%) and RGB-only imaging (c. 90%) — in line with the cost/accuracy trade-off described in [17].

That 8-10 point modality penalty is small compared to the 26.6-60.9 point condition penalty in section 3.2, and neither of the penalties are additive as a procurement decision would have it. In the whole set of extracted data, among the three commercial solutions, the lower figure is in the case of commercial NIR (high cost sensing) evaluated using real waste, which is 17 points lower than commercial RGB imaging evaluated using clean samples. It is the state of the material presented to the sense rather than its sense that is more important to realised accuracy. With no attention to contamination and blend handling, spending on improved sensors allows for a buy for costs that are less than those published for per modality.

This is directly relevant to the design in Section 5. It cannot redeem RGB imaging, which is the weakest when compared to others. It proves that the dominant error term is not the sensor but the condition of the input; as long as the imaging system can recognise contamination and blend as first-rate signals (which happens when the design position is taken as stated in Section 5.1), it's not being put at such an unfair disadvantage as the modality factor would suggest.

3.4. Limitations of This Synthesis

These figures were constrained by four things. The sample is small: there are only 3 studies and only 3 pairs of observations within each study, meaning the median = 42.2 points is based on 3 observations and a 42.2-point mean of the data is neither computed nor implied. Not all reported accuracy is done in the same way (per class, per class and pooled, or best of method, etc.) and this does not recalculate the accuracy in common units for the reports provided, so this is not possible from published summaries alone. Two of the three within study pairs belong to the very same study [19] so this study gets an exaggerated weight. The same applies in this case, however: the publication bias is in the same direction—methods that were unsuccessful in blended fabrics would not be published at all, so the situation reported above would actually underestimate the reality. The numbers are presented in this way, therefore, not as an inferentially strong estimate, but rather as a statement of a pattern that was revealed in the literature.

4. RESEARCH GAP

Seven gaps are apparent when mapping the themes and patterns as identified in Section 2, this motivating the design in Section 5.

- G1 — Real world classification accuracy of blended/failure-prone textile samples. Multiple coloured or lightweight fabrics, blending and fabrics that have been dirty, consistently reduce accuracy in each of the studies reviewed [11,17,19] with the worst being polycotton blends. The methodology used for classifiers does not include the self-estimate of the classifiers' confidence (Section 5) on a large and diverse real-world data-set they are not even tested on any pure versus blended set forcing a single validation indicator that is used as a global average describing the entire family (Section 5).
- G2 — Disconnection between classification and recommendation to a recovery-path. Sorting/classification research and research dealing with recycling technologies are published separately, that is as two separate literatures [17,19] and [16,20] respectively. The integrated value-chain review [14] is the only reviewed document that makes an explicit call for coordination; no paper presents a single pipeline that passes an output of classification to a feasibility matched recovery recommendation (Section 5).
- G3 — Non-textile grade material is an underserved market. In Section 2.4, material which is not suitable for textile-to-textile recycling (such as acoustic panels, nanocellulose and structural composites) are identified as having mature pathways; the sorting and reuse research focuses on material classification, however, rather than material routing across multiple recovery paths, resulting in a systematic undervaluation of such material classes in a fibre-to-fibre-only recommendation engine (Section 5).
- G4 — Disjointed and inconsistent sustainability metrics. With regards to sustainability claims, as seen in the reviewed literature in the field [4,5,7,8], these are difficult to compare between the studies due to the use of a variety of functional units and system boundaries, and in none of the reviewed literature is the systems' circularity score calculated at the point of classification, which is the closest to the material itself (Section 5).
- G5 — Traceability and regulatory-reporting integration is mostly theoretical. Tagging as coming requirements, RFID-based traceability [21] and Digital Product Passport policy [1] don't refer to reporting formats exportable to ERP- or regulatory-ready reporting portal, which is reported in the literature (Section 5).
- Limited role differentiated decision support, G6. The surveyed studies focus on different results: Each one provides a viewpoint towards a recycling-facility operator (throughput, sorting accuracy), a sustainability manager (carbon, water, social impact) and a manufacturer (traceability, EPR compliance); however no study is able to provide all three

selected from the same data source. The gap is deployment integration and to that extent it is not a gap in the science at that time, but rather treated as such (Section 5).

- G7 — No described transaction mechanism between 'classified batch' and 'buyer'. All recovery-pathway research cited above confirm that a certain fibre condition is appropriate for certain types of route but none address how a batch is priced, matched to an interested buyer and transferred once the type of route has been identified and the details worked out, and part of the reason for the maturity gap between textiles and packagings derives from this dependence on the export market rather than planned offtake from the domestic market, as stated by [18]. In addition, it is an infrastructure gap (not a research gap), and is noted here for completeness, as it is explicitly out of scope for the design in Section 5.

The sample sizes also vary and some studies rely on less than 100 samples [16,4,9] whilst fibre classification studies have had more thorough development using deep learning and large samples [11,17,19]. As a result of this heterogeneity it would be impossible to take a single modelling strategy throughout all the modules (Section 5).

TABLE VII
MAPPING OF GAPS TO SYSTEM COMPONENTS

Gap	Addressed by Component	Section
G1: Blend/contamination accuracy	Image Analysis Engine + Classification Engine (confidence output)	4.1, 4.2
G2: Classification–recovery disconnect	Direct coupling: Classification → Scoring → Recommendation	4.1
G3: Non-textile outlets under-served	Recommendation Engine industrial-recovery branch	4.1
G4: Inconsistent sustainability scoring	Sustainability Impact Engine, fixed functional unit	4.3
G5: Traceability/reporting gap	Batch registration + DPP-compatible export	4.1, 4.4
G6: Limited role-differentiated support	Role-based dashboards (operator/manager/manufacturer/admin)	4.4
G7: No buyer-matching/transaction mechanism	Recorded as out of scope; not addressed by this design	—

5. PROPOSED SYSTEM ARCHITECTURE

The Textile Waste Intelligence Platform created directly against gaps found in Section 4 and within its scope, as an end-to-end pipeline and not a "classifier". The resulting flow is recorded in the Fig. 7, which is explicitly related to the literature reviewed in the Section 2. This uninstrumented batch is first imaged and analysed for texture, pattern, damage, and contamination information; fibre composition is then estimated with a confidence score or an estimated blend ratio; this mapped batch is then assigned to one of six waste categories with each category having a corresponding sub-score that carries a specific weight; finally, a specific recovery pathway is recommended based on this mapping (Table VIII), and the environmental impact of that pathway is estimated using a set functional unit. The pipeline only goes forward, except for a single feedback loop; any low-confidence classification results are returned to a human review queue, instead of being automatically scored and recommended. This does not remove, but does include this blend-classification error (G1) as discussed in Section 2.3. The design also includes a stakeholder-approval stage, presented in Fig. 5, through which produced reports can undergo regulatory review and approval.

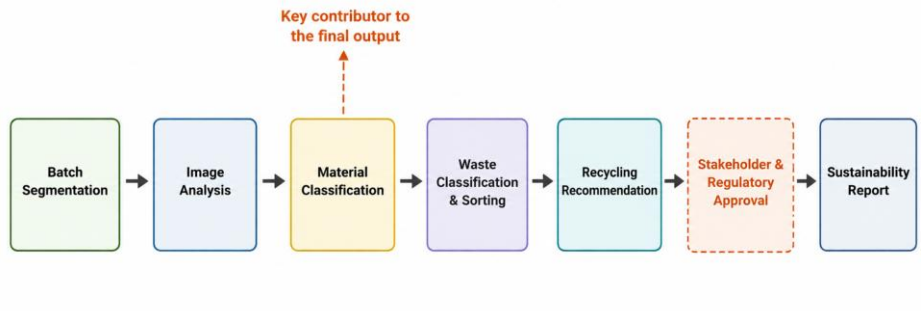


Fig. 7. Textile Waste Intelligence Platform pipeline, with low-confidence classification outputs routed to manual review before pathway assignment.

5.1. Layered Feature Description

While the image analysis starts, the batch is registered with the structured intake (batch ID, source, quantity, declared fabric type (where known), colour, condition and collection date) and exportable, regulator-facing, reporting is required of the image analysis by means of the anchor record of which the structured intake metadata was a part, such as that is required in [18,1] when an RFID tag or DPP is not used.

Each garment is then sent through a computer-vision pipeline that recognises the texture, weave or knit pattern, colour, visible damage, and visible contamination as distinct labels, directly addressing the main drivers of misclassification for blended fabrics encountered in [19] and lack of usable texture information in RGB images pointed out in [17]. By further separating condition detection from fibre-identity inference, the downstream Material Classification Engine may factor the next confidence level for the fibre identity based on detected contamination.

The Material Classification Engine classifies materials as one of the following: a hard label class (cotton, polyester, nylon, rayon, wool, silk, linen, denim, acrylic, or mixed fibre composition); or a probability mass function of materials, that is a probability distribution over fibre types with an estimated blend ratio (not a hard label). It directly follows the accuracy pattern found in [11,17] (95-100% on pure fibres, much lower on blends) and the finding in [16] of the fact that even knowing about a small fibre component of a blend can make the difference between being able to operate an engine-based recovery process at all and not at all. The Image Analysis Engine will clearly identify low-confidence predictions due to lighting issues, damage, contamination, etc. and not "black-hole" downstream logic.

Each batch gets categorized into one of the six categories (recyclable, reusable, repairable, upcyclable, compostable, hazardous) and is assigned a score to each of the five weighted sub-scores outlined in Section 5.3: recyclability, condition, reuse potential, environmental benefit, and processing feasibility, that add up to the overall circularity score. This is due to the use of multiple weights instead of one single index, following the multi-criteria approach validated for that purpose in [7] where recycling and material recovery were the most robust recovery options with respect to the sensitivity analysis performed in that study, and the specific weights used herein are initial design weights and not inherited from [7]; the results of calibration of these weights are reported as outstanding in Section 6. A hazardous or heavily polluted classification does not go down the rabbit hole: it leads to the thermochemical recovery classification found in [20] as a means of solving G3.

Based on this classification and the score, the Recommendation Engine chooses which of these recycling methods (mechanical or chemical fibre recycling, reuse of fabric, upcycling, donation, cross-sectoral industrial recovery) to choose. The industrial-recovery sector is filled with concrete, literature-based destinations, not a catch-all: Acoustic panels for cotton scraps [6] and nanopulverization (extraction of nanocellulose) for cotton-rich material; Chemical and enzymatic depolymerisation for PET/nylon/elastane-dominated synthetics [16] are all examples. There are various other cross-sectoral outlets described in [10] that are checked for low value mixed streams, while highly heterogeneous or contaminated will be sent to a thermochemical/energy recovery [20] regime.

TABLE VIII
RECOMMENDED RECOVERY PATHWAY BY DETECTED COMPOSITION AND CONDITION

Detected Composition	Condition	Recommended Pathway
Pure cotton / linen / viscose	Clean, high cellulose content	Nanocellulose extraction; fibre-to-fibre recycling
Pure polyester / nylon (PET, PA)	Clean, single polymer	Chemical depolymerisation / mechanical recycling
Cotton/polyester blend	Clean, unseparated blend	Composite material manufacturing (PP composites)
Wool/PET/elastane (3-component)	Clean or lightly worn	Step-wise chemo-enzymatic separation, PET recovery
Pre-consumer scrap (knit/woven/denim)	Clean, homogeneous, high volume	Acoustic/construction panel manufacturing
Untreated natural fibre (cotton, linen, wool)	Heavily worn, undyed or low-impact dye, no synthetic blend	Composting / biological recovery
Mixed/unidentifiable blend	Contaminated or highly heterogeneous	Thermochemical valorisation (syngas, BTX recovery)
Any fibre type	Good structural condition, low wear	Reuse / donation / repair
Any fibre type	Low-confidence classification	Route to manual review before pathway assignment

5.2. Implementation Details

A working prototype of the pipeline exists and will serve as the basis for an evaluation as detailed in Section 5.5; in this case, no results are reported from the pipeline since no performance results are reported here. The backend is implemented in Python 3.11 with FastAPI on an async Uvicorn server, with SQLAlchemy defining the mapping of the waste batch, fabric item and user objects, and storing this data in a relational schema stored in an SQLite database for transactional storage of inventory, user and audit data, whilst loading and extracting dominant colours from images is handled by the Pillow (PIL) library. The front end is a single page application created in React using Vite, styled with Tailwind and works with Chart.js (through react-chartjs-2) for dashboard analytics and with Lucide React for iconography. Table IX summarises the stack. All four dashboard views (operator, sustainability manager, manufacturer, administrator) are not built out of four different systems, but rather one data model.

TABLE IX
PROTOTYPE TECHNOLOGY STACK

Layer	Component	Role
Backend	FastAPI (Python 3.11)	Async API endpoints, auto-generated OpenAPI/Swagger docs
Backend	SQLAlchemy (ORM)	Maps users, batches, and fabric items to relational schema
Backend	SQLite	Transactional storage for inventory, users, audit logs
Backend	Pillow (PIL)	Image loading, dimension parsing, dominant-colour extraction
Backend	Uvicorn	ASGI server running the FastAPI application
Frontend	React + Vite	Single-page application, fast bundling, hot module reload
Frontend	Tailwind CSS	Custom design tokens, responsive layout
Frontend	Chart.js / react-chartjs-2	Real-time dashboard analytics and telemetry visualisation
Frontend	Lucide React	Interactive vector iconography

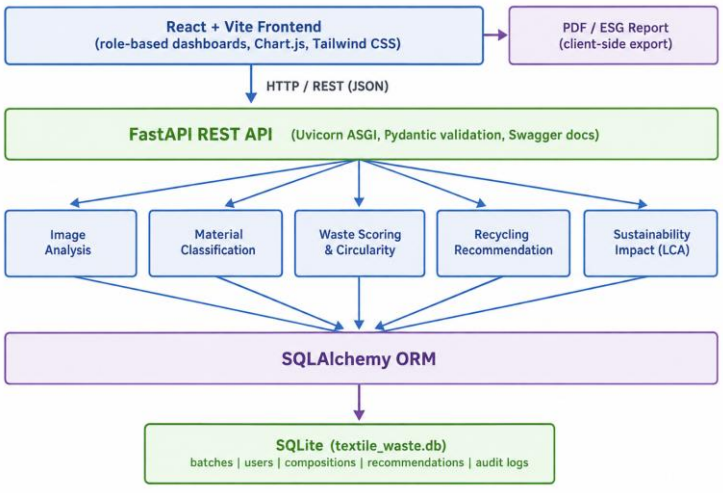


Fig. 8. Overall software architecture and module interaction.

For model development, four datasets are utilized that are publicly available and citable, and selectable based on the stage of the pipeline they relate to and not based on aggregate size. The closest available proxy to the task addressed in G1 is the Cotton Fabric Image BD collection of original and augmented images of fabrics, corresponding to 1100 fabrics of fibres between 30 and 99 per cent, captured under the supervision of the textile-engineer on a thread-counting machine (cotton-percentage classification, available via https://authors.elsevier.com/MWnYR42JY0S2Mc). Fabric Spot Defect is a database of annotated image of stains, discolourations, oil marks and rust on 1,014 cotton, linen, silk, denim and patterned fabric samples (citation req.: doi:10.17632/6574nhzm8x.1), which is based on the contamination and damage detection task, for which Section 5.1 did not keep its own database separate. There are single-fibre reference data provided from an optical-coherence-tomography

dataset of wool, cotton and polyester fabrics published for automated recycling research: Sabuncu and Ozdemir, Data in Brief, vol. 45, Art. 108719, 2022. There are two reasons why Deep Fashion is not used for fibre information: 1. The labels used by Deep Fashion are either garment categories (shirt, trouser, coat) or there is no label provided at all. 2. The label information is only used for visual feature pretraining.

The limitation is even more explicitly stated: none of these datasets includes any verified blend ratios obtained from actual, contaminated, multi-component garments combined, which is the area false that [11,17] point out as a deficiency. Most near, but not quite, are the Cotton Fabric Image BD biosensors; they measure the amount of cotton only from clean samples. Reliable creation of multiple data sets specifying both blend ratio and contamination is the only remaining limiting factor for G1, and is incorporated into Section 5.

The Image Analysis Engine is composed of three different steps. Colour normalisation makes sure that each uploaded image has just one dominant RGB value and uses a human-readable image name (such as 'Royal Blue') for a display on the operator control panel. Texture analysis divides the weave or knit pattern and physical construction (also known as twill, satin, plain weave, coarse knit or coarse stitch) from the colour of fabric and clothing. Defect and contamination are logged separately from fibre-identity classification as found earlier in the error on contamination when identifying fibre classes. This is in keeping with previous results on the effect of contamination on accuracy in fibres class identification, in which defect/ contamination is registered separately from fibre-identity classification. Classification has two complementary architectures: EfficientNet-B4 for fine-grained texture, and a Vision Transformer (ViT-Patch16) for macro-structural pattern and visible damage, with combined predictions in a single fabric-category output and with confidence score. This is consistent with the observations by [17] that CNN/ANN-type models are better at capturing the fine details of textures, while transformer-type models perform better on the overall structure of the fabric in general without any one architecture being superior for all types of fabric.

The five weighted sub-scores thus compile to one Circularity Score (CScore), which is mapped to a qualitative class of recovery to be assigned to the harvested value (CQ); and which can suggest, automatically, a recovery processing out of the three possible options provided in Table IX (mechanical shredding, chemical depolymerisation, industrial felting). In a batch application, composition is associated with estimates of the contribution to the environment that exceed those in the manufacture of products (as in [5]) and the LCA/LCC framing as in [8], but in this case at the point of waste processing and not manufacture. G5 & G6 (regulatory-reporting integration & role differentiated decision support) are supported via exportable and structured reports derived from the same underlying data-model described above.

5.3. Mathematical Models

All scores and savings are reported and followed by a clear, well documented formula, and not by an opaque heuristic, so scores can be audited and reproduced.

The Circularity Score is calculated as a weighted sum ranging from 0-100 weighted by (1):

$$\text{Circularity Score} = 0.35R + 0.20C + 0.20P + 0.15E + 0.10F \quad (1)$$

Each category is defined as follows: R – Recyclability, C – Condition, P – Reuse Potential, E – Environmental Benefit, F – Processing Feasibility. The recyclability rating compares the fibres base recyclability rating (0.0-1.0) and gives it a scale of 0-100, before applying a –25 point penalty for contaminated batches. The default value of C is 90 for clean/recyclable material, 60 for damaged, 40 for wet but not damaged, 20 for contaminated, and 50 if it is not mentioned. Reuse Potential (P) is similar, with 85 for a clean; 50 for damaged; and 20 for contaminated; defaulting to 40. Environmental Benefit (E) is 95 (clean) 30 (contaminated) with default 70. Processing Feasibility (F) = 90(clean), 30 (contaminated); default 60. The score is related to a recovery status: 85+ for Excellent, 70-84 for High, 55-69 for Moderate, 35-54 for Limited and below 35 for Disposal Recommended.

Explicitly stated weights, component values and bands - can be challenged. They are design specification values which mean nothing to each other, are not arrived at from outcomes data, and not sensitivity tested. The score should never be interpreted as a measurement, rather as a transparent and repeatable ranking heuristic, and comparisons of different batches should be made only as functions of each other, under a fixed weight vector for the various measurements.

The environmental savings are determined per fixed functional unit (per kilogram of processed textile waste) and are calculated with the selected recovery options, and compared to virgin-material production as shown in (2):

$$\text{Savings} = (\text{Impact(virgin)} - \text{Impact(recycling)}) \times Q \times \eta \times M \quad (2)$$

where Q is the batch quantity in kg, η is an industry term, independent of all pathways, to account for the loss of mass through the recovery process, which is equal to 0.85 for each of the three categories of upcycling, donation and fabric reuse; and 0.60 for chemical depolymerisation and 0.80 for mechanical recycling. These factors are not added in as a limiting case, meaning that there is no material recovered from landfill and disposal are reported separately as no savings against the material virgin mass used. Coefficients of carbon per fibre, water per fibre and virgin value per fibre required for the calculations mentioned above are presented in Table X.

TABLE X
FIBRE-LEVEL LCA COEFFICIENTS

Fabric Type	Carbon Factor (kg CO2e/kg)	Water Factor (L/kg)	Virgin Value (USD/kg)
Cotton	8.5	2500	2.20
Polyester	12.0	350	1.50
Wool	16.5	1800	6.80
Nylon	15.0	450	2.10
Silk	22.0	3200	25.00
Linen	6.0	800	4.50
Acrylic	13.5	400	1.80
Denim	9.5	2200	3.00
Blend (default)	10.5	1400	2.00

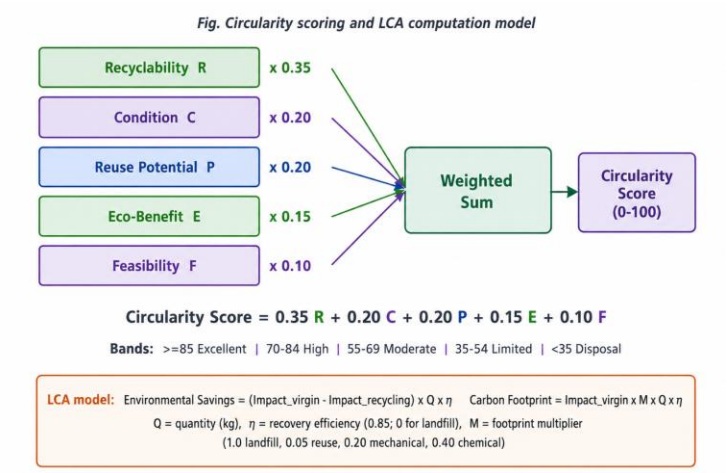


Fig. 9. Circularity-scoring and Life Cycle Assessment computation model.

Rather than measurements actualized in a peer-reviewed life cycle inventory, the values of the coefficients in Table X have been compiled with the aim of providing order-of-magnitude values and are not claimed to be accurate. There are considerable differences in the published life-cycle figures depending on the methods used to allocate the fibre value and/or the region where the fibre is grown, and for cotton in particular, the value of water varies widely in the literature. They are added to give concrete numbers and chance to replicate the calculation and are considered as placeholders that can be configured with numbers from a recognised inventory database (such as ecoinvent or an equivalent sector LCI) when the deployment is reported by others outside the scope. The equations simulate only on a per-kilogram basis, which means that changing the source alters the results of the simulation, but does not affect the model. These coefficients will be verified against a recognised inventory in Stage 4 of the evaluation protocol detailed in section 5.5.

5.4. Access Control and Interactive Tools

The access is available with role-based access control (RBAC), JWT token authentication and optional OAuth2 sign-in, with four roles, all having a common data model: Recycling facility operators, including users who register and classify batches, Sustainability managers, who have access to carbon-auditing dashboards and can export ESG reports, Textile manufacturers who register and classify their own batches, and Administrators, who have complete read/write/delete access and user management as well as access to API documentation.

In addition to the live inventory pipeline, a standalone Sustainability Calculator is available that allows a user to run a what-if scenario without recoding data, providing the same circularity score, water and carbon savings, preserved material value, and processing-feasibility output of the live pipeline, with the results exported as a printable report.

Reporting Module generates two types of paginated downloadable reports: a multiple page batch report containing metadata, analysis results, and a list of preparation steps, materials and solvents that are needed for the recommended pathway, and; an executive report summarising the scope and composition of compliance, as well as disclosures of waste categories. Both meet answer to the exportable, requirements for regulation facing reporting requirement identified in Section 4.

5.5. Evaluation Plan

Although as software the prototype described in sections 4.1–4.4 has not yet been evaluated, its classification models have not been trained with a task-specific labelled corpus and no accuracy or scoring figure nor impact figure produced by it are given in this paper. However, just as with laboratory accuracy, the literature surveyed in this case indicates that even two off-the shelf test splits would not provide meaningful validation in this scenario. Four stages make up the following protocol, which can be tested for validity and is thus subject to falsification rather than discreditation in the form of plausibility.

Stage 1 — classifier benchmarking. The classifier would then be trained and evaluated on one subset of the data pruned from each of the datasets referred to in 5.2 and would be reported on separately, per class: for pure single-fibre garments and per class: for blended garments in the context of separate precision and recall figures for each class and the F1 score per class. CottonFabricImageBD is the most important benchmark as it allows having cotton-percentage labeling that enables capturing the error in the blend ratio in terms of mean absolute error in percentage points instead of just a categorical accuracy – as Section 2.3 elucidates, the aggregate may mask total failure on one specific fabric family. It would be reported if not this breakdown, then an overall accuracy would be repeating the reporting deficiency that the survey identifies.

In Stage 2, there is contamination and confidence validation. A complete set of physical samples from clean, damp, soiled, dark coloured and multi-component fabrics would be collected and hand labelled according to composition. Two properties are more important than simply what we call accuracy — and whether it decreases on contaminated and blended samples according to the 39–73% range reported in [19,11] — and if the model is calibrated, that is, if it produces a confidence output which, for predictions at 90%, is correct approximately 90% of the time. If there is an uncalibrated confidence score, then the human-review routing this is depending on would be invalid, so this gates the following.

Stage 3 - scoring and pathway agreement. Circularity score and recommendations in Table VIIII would be compared to an experienced operators judgement in the identical batches and would report rank correlation and inter-rater agreement. The aim is not to prove the model itself is correct, in an absolute sense (with no ground truth for “correct circularity” in existence), but to determine if the weights in Section 5.3 lead to a defensible ordering and then refit the weights if they do not.

Stage 4 — impact-figure verification. The estimates of savings would be re-calculated for a set of processed batches using the recognised life-cycle inventory instead of Table X, and the two sets compared. If there is a large difference, the conclusion that the reported environmental benefit is an artefact of the coefficient source, as specified in the failure mode in the boundary-inconsistency findings in [8], is not necessarily true.

This protocol does not include deliberate considerations of throughput, driver support, or cost per kilogram processed; these are factors that must be related to the pre-processing automation as described in [12], which is still immature.

6. LIMITATIONS

This paper makes four limitations. First, the review is a structured synthesis, not following the PRISMA systematic search: the 24 studies were not selected due to the Scopus/Web of Science coverage (which was not tested), nor were three studies [15,22,23] cited because the authors did not have the opportunity to verify the results themselves. A systematic search therefore might bring to light work done that this synthesis did not find, particularly in non-English languages.

Secondly, and more relevantly, the platform outlined in Section 5 is unassessed. The existing working prototype has not been designed and built using a task-specific labelled corpus, nor tested with waste; all of the reported accuracy values are from the reviewed literature and are not those of this system. It is stated in Section 5.5 how it is to be tested, but until it is tested in accordance with that protocol, this performance claim and/or implication is not made or suggested.

Thirdly, the interpretation of the parameters controlling the out of the platform are not validated. The weights, component values and bands in Section 5.3 are design values and not fitted weights, and the life-cycle weights in Table X are derived from a compiled second data set and are not a peer reviewed inventory. Both would require calibration before footage of circularity score and/or savings could be claimed to an external audience.

Fourth, the choice of the sensor limits the design's capabilities. Although RGB imaging can be used to sort the images, in the platform blend-classification error exists as calibrated confidence, and human review, but not as it is eliminated by the system. Properly closing G1 means either richer sensing or access to the large contaminated-and-blended training set that no public data set does yet have – but it is not addressed whatsoever here (G7) is a buyer/pricing matching/batch pricing.

7. CONCLUSION

In this paper, 24 papers on textile identification, texture-based identification using AI, recycling chemistry, value added recovery, sustainability assessment and the circular-economy policy were reviewed and quantified as a whole volume; that is, the individual papers were not read consecutively. The consequence of this is that accuracy drops an average of 42.2 percentage points between within-study pairs of samples when moving from clean single-fibre samples to blended or contaminated

material, compared with 8–10 percentage points when using the cheaper RGB imaging set-up versus spectral sensing, across all of the extracted figures 28.7 points. Two additional results of this work restrict any system that can be built in that space: sorting and research on recovery routes are almost completely disconnected and separated in both the execution and dissemination process; and sustainability claims lack symmetry, since different kinds of systems have inconsistent boundaries.

From this Synthesis resulted in the platform design that incorporates classification, circularity scoring, pathway recommendation and impact estimation into a single platform, with confidence being explicitly exploited at every stage and the low-confidence garments being forwarded to human intervention instead of being sent silently to a recovery pathway. It does not rely on a new sensing technique or new recycling process, but works by integrating and making it transparent.

It is not proven; the design is not tested. At the moment, its value will depend whether the protocol of Section 5.5 upholds it, namely confidence scores/probabilities are well calibrated; the latter is fundamental, since the routing that relies on human review for defense depends entirely on it. The next step is to perform that evaluation, and gather this "blended-and-contaminated" dataset.

List of Abbreviations

AI: Artificial Intelligence
ANN: Artificial Neural Network
CNN: Convolutional Neural Network
DPP: Digital Product Passport
EPR: Extended Producer Responsibility
FT-IR: Fourier-Transform Infrared Spectroscopy
HSI: Hyperspectral Imaging
LCA: Life Cycle Assessment
LCC: Life Cycle Costing
NIR: Near-Infrared
RBAC: Role-Based Access Control
RFID: Radio-Frequency Identification
S-LCA: Social Life Cycle Assessment
SVM: Support Vector Machine

Declarations

Ethics approval and consent to participate

Not applicable. This study did not involve human participants, human data, human tissue, or animals.

Consent for publication

Not applicable. This manuscript does not contain data, images, or videos relating to any individual person.

Availability of data and materials

Data sharing is not applicable to this article as no new datasets were generated during the current study. All data analysed during this study are the performance values reported in the previously published articles listed in the reference list; the extracted and harmonised accuracy values used in the quantitative synthesis are included in this published article (Table VI, Section 3). The public datasets referred to in Section 5.2 are available from their original sources as cited therein. The extraction table and the corresponding source-to-value mapping are available from the corresponding author on reasonable request.

Competing interests

The authors declare that they have no competing interests.

Funding

Not applicable. No funding was received for conducting this study.

Authors' contributions

VS conceived the study, designed the review protocol, supervised the work, and critically revised the manuscript for intellectual content. SS performed the literature search and screening and contributed to the thematic synthesis of the reviewed studies. TM carried out the extraction and harmonisation of the reported classification accuracies, prepared the figures and tables of the quantitative synthesis, and drafted the manuscript. YBB contributed to the design of the proposed platform architecture, the scoring models, and the evaluation plan. All authors read and approved the final manuscript.

Acknowledgements

Not applicable.

Authors' information (optional)

VS is an Assistant Professor in the Department of Computer Science and Engineering, Sri Krishna College of Engineering and Technology, Coimbatore, India. SS, TM and YBB are undergraduate students in the Department of Computer Science and Engineering at the same institution, with interests in computer vision, applied machine learning and sustainable computing.

8. REFERENCES

[1] A. C. Manivannan, A. Bhatnagar, K. Niinimäki, L. Panneerselvan, and T. Palanisami, "Evaluating Textile Waste Management Policies: Lifecycle Gaps and Opportunities for Improvement," Waste Management & Research, vol. 44, no. 5, pp. 508–524, 2026, doi: 10.1177/0734242X251393613.

[2] G. Morell-Delgado, L. Talens Peiró, and S. Toboso-Chavero, "The Circularity of Textile Waste: A Systematic Literature Review on Environmental Methodologies and Knowledge Gaps," Waste Management, vol. 222, Art. 115648, 2026.

[3] A. Kaur et al., "Nature-Based Wastewater Treatment Technologies: A Comprehensive Review of Constructed Wetlands and Vermifiltration," Environmental Advances, 2026.

[4] A. Bhatnagar and K. Niinimäki, "Evaluating Potential Social Risks in Fibre-to-Fibre Textile Recycling Through Site-Specific and Generic Social Life Cycle Assessment," Sustainable Production and Consumption, 2026.

[5] Sandaruwan et al., "Towards Sustainable Development: Reducing Organisational Carbon Footprint in Textile Fabric Production," Corporate Social Responsibility and Environmental Management, vol. 33, pp. 399–418, 2026.

[6] M. Jucienė, V. Dobilaitė, K. Miškinis, and V. Paukštys, "Research and Evaluation of Acoustic Panels from Clothing Industry Waste," Textiles, vol. 6, no. 1, Art. 11, 2026, doi: 10.3390/textiles6010011.

[7] D. Pamucar et al., "Sustainable Environmental Management of Post-Consumer Textile Waste Flows Using a Fuzzy Trigonometric COBRAC-Frank Soft RAFSI Model," Applied Soft Computing, vol. 199, Aug. 2026.

[8] A. Anceschi, R. Mossotti, and A. Patrucco, "Assessing Sustainability in the Textile Sector: A Review of LCA, LCC, and S-LCA Methodologies with a Focus on Polymeric Textile Materials and Circular Strategies Along with Future Perspectives," Polymers, vol. 18, no. 4, Art. 534, 2026, doi: 10.3390/polym18040534.

[9] L. Z. Maduna, M. E. Masekela, N. C. Shiba, Y. Chabalala, L. M. Sekhohola-Dlamini, M. Moreroa, N. Thabethe, M. G. R. Mahlobo, N. D. Baloyi, and B. S. Madonsela, "The Advancement of Waste Reuse and Conversion Practices by Waste Management Groups in the Indigenous Communities of Bushbuckridge Local Municipality, South Africa," Discover Environment, vol. 4, Art. 5, 2026, doi: 10.1007/s44274-025-00514-1.

[10] M. Bertelli, D. Giorgi, C. Morea, and L. Incrocci, "Textiles Waste as Resources for Relevant Cross-Sectoral Applications, Opening the Cycle to Reach Circular Economy—A Review," Sustainability, vol. 18, no. 3, Art. 1464, 2026, doi: 10.3390/su18031464.

[11] M. Robinson, S. Ghosh, F. Qian, C. Du, M. Vallati, and P. Goswami, "Recent Development on Sorting of Textiles Waste by Fibre Type for Recycling: A Mini Review," Textiles, vol. 6, no. 1, Art. 28, 2026.

[12] D. Lopes, E. J. Solteiro Pires, V. Filipe, M. F. Silva, and L. F. Rocha, "Intelligent and Automated Technologies for Textile Recycling Pre-Processing: A Systematic Literature Review," Technologies, vol. 14, no. 4, Art. 200, 2026, doi: 10.3390/technologies14040200.

[13] M. Shamsuzzaman, M. Islam, M. A. A. Mamun, R. Rayyaan, K. Sowrov, S. Islam, and A. S. M. Sayem, "Fashion and Textile Waste Management in the Circular Economy: A Systematic Review," Cleaner Waste Systems, vol. 11, Art. 100268, 2025.

[14] T. Hasan, M. S. Sheikh, M. L. Hossain, and M. B. Hoque, "A Comprehensive Review of Solid Waste Management in the Textile Industry," Journal of the Air & Waste Management Association, vol. 75, no. 11, pp. 839–854, 2025, doi: 10.1080/10962247.2025.2550365.

[15] K. Chawla et al., "Fabric Classification for Textile Waste Circular Economy," in Proc. IEEE International Conference on Application of Information Technologies in Engineering, Management and Science (ICAI-TEMS), 2025.

[16] E. Boschmeier, D. Mehanni, V. L. Sedlmayr, Y. Vetyukov, S. Mihalyi, F. Quartinello, G. M. Guebitz, and A. Bartl, "Recovery of Pure PET from Wool/PET/Elastane Textile Waste Through Step-Wise Enzymatic and Chemical Processing," Waste Management & Research, vol. 43, no. 6, pp. 969–979, 2025, doi: 10.1177/0734242X241276089.

[17] E. Faghih, Z. Saki, and M. Moore, "A Systematic Literature Review—AI-Enabled Textile Waste Sorting," Sustainability, vol. 17, no. 10, Art. 4264, 2025, doi: 10.3390/su17104264.

[18] M. T. Brouwer, A. K. Undas, E. Maaskant, et al., "Factors Shaping End-of-Life Management of Short-Lived Consumer Goods: A Comparison Between Plastic Packaging and Textiles," Waste Management, 2025.

[19] H. Stipanovic, G. Koinig, T. Fink, and A. Tischberger-Aldrian, "Effects of Surface Contamination on Automated Textile Sorting Using NIR-Spectroscopy," Waste Management, 2025.

[20] S. van Wyk, C. F. Mourão Vilela, and B. J. Vreugdenhil, "Thermochemical Valorisation of Real Post-Consumer Mixed Textile Waste," Fuel, 2025.

[21] P. Ugale and M. Brewster, "Unlocking Circularity in Textiles Through a RAIN-Enabled Automated Framework," IEEE J. Radio Freq. Identification, vol. 9, pp. 919–936, 2025, doi: 10.1109/JRFID.2025.3636554.

[22] Y. Spyridis, V. Argyriou, A. Sarigiannidis, P. Radoglou Grammatikis, and P. Sarigiannidis, "Autonomous AI-Enabled Industrial Sorting Pipeline for Advanced Textile Recycling," in Proc. 2024 20th Int. Conf. Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT), Abu Dhabi, UAE, 2024, pp. 455–461.

[23] C. S. Ranganathan, S. Kabaleeswaran, S. P. Vimal, V. Kannagi, M. Muthulekshmi, and M. Meganathan, "Intelligent Robotics for Sorting Clothing and Textile Waste Using Cloud-Integrated Recycling Systems and CNN," in Proc. 2024 Int. Conf. Smart Technologies for Sustainable Development Goals (ICSTSDG), Piscataway, NJ, USA: IEEE, 2024, pp. 1–6, doi: 10.1109/ICSTSDG61998.2024.11026397.

[24] R. Abidnejad, H. Baniasadi, M. Fazeli, S. Lipponen, E. Kontturi, O. J. Rojas, and B. D. Mattos, "High-Fiber Content Composites Produced from Mixed Textile Waste: Balancing Cotton and Polyester Fibers for Improved Composite Performance," International Journal of Biological Macromolecules, vol. 292, Art. 139227, Mar. 2025, doi: 10.1016/j.ijbiomac.2024.139227.