

A Systematic Survey of Molecular Taggants, Distributed Ledger Technology, and Machine Learning for Verifiable Plastic Waste Traceability

Purva Shinde^{1*}, Adweta Dengale¹, Abhiraj Gatmane¹, Nayan Rathod¹, Umesh Pawar²

¹ Department of Computer Science and Engineering, Sandip University, Nashik, Maharashtra, India ² Department of Computer Science and Engineering, Sandip University, Nashik, Maharashtra, India

Abstract - Extended Producer Responsibility (EPR) is now the dominant regulatory instrument for plastic packaging in India, the European Union and elsewhere. It rests on an assumption that current systems cannot test: that a producer's declaration of recovered and recycled tonnage is true. Compliance evidence is aggregate, self-reported and detached from the material it describes, so the central claim of the regime is cheap to falsify and hard to disprove.

This paper is a systematic survey conducted under a PRISMA-style protocol. We screened 312 records from five databases and retained 45 sources published between 2023 and 2026 across four communities that seldom cite one another: molecular marking of polymers, machine-readable identifiers, distributed ledger technology (DLT) for provenance, and machine learning for waste classification, forecasting and anomaly detection. The included studies are quality-appraised, and the four bodies of work are compared against one another using a four-layer analytical framework and six capability criteria derived from the regulatory requirements.

Three principal findings emerge. Not one reviewed configuration satisfies all six criteria, and the sharpest — whether an identifier survives mechanical reprocessing — is failed categorically by every approach built on digital identifiers alone. Where one layer is weak another is characteristically strong, yet no reviewed study evaluates the layers in composition. The third finding carries furthest: because no shared dataset, protocol or benchmark permits reported accuracies and throughputs to be compared, the field's apparent maturity cannot be verified, and its limiting constraint proves evaluative rather than technical. We specify what an adequate benchmark requires, and derive a reference framework whose design decisions are each traced to supporting evidence. Implementation and empirical validation are outside the scope of this survey.

Keywords: plastic waste traceability; Extended Producer Responsibility; distributed ledger technology; molecular taggants; digital product passport; machine learning; circular economy; GS1 EPCIS

Abbreviations: CNN, convolutional neural network; DLT, distributed ledger technology; DPP, digital product passport; EPCIS, Electronic Product Code Information Services; EPR, Extended Producer Responsibility; ESPR, Ecodesign for Sustainable Products Regulation; IPFS, InterPlanetary File System; LSTM, long short-term memory; MSW, municipal solid waste; NFT, non-fungible token; NIR, near-infrared.

Terminology. *Molecular taggant* is the umbrella term for markers resident in the material itself; *DNA taggant* is the subclass encoding information in nucleotide sequence; *optical identifier* denotes printed machine-readable codes. These terms are not interchangeable. *Programme-scale* denotes deployment across a state-level or national EPR programme, as distinct from a single municipality.

1. INTRODUCTION

Global plastic waste generation exceeds 350 million tonnes annually, and little re-enters the material economy in usable form [1]. Policy has responded by shifting end-of-life cost onto the entities placing material on the market. The United Nations Environment Programme frames the transition as depending less on recycling capacity than on the accounting structures determining where material flows [2].

Two review streams reach that diagnosis from opposite directions. Where information-based plastic tracking is surveyed, the finding is that physical marking and digital record-keeping matured separately, so a digital record asserting a material property cannot be checked against the material [3]. Surveys of distributed ledger technology (DLT) in plastic waste management identify the complementary defect: immutability confers no accuracy on the claim a record contains [4]. Neither literature treats the other's failure mode as its own concern.

The international legal architecture has not resolved this. The Intergovernmental Negotiating Committee convened under UNEA Resolution 5/14 has met repeatedly without agreeing a binding text, and the reconvened fifth session in Geneva in August 2025 adjourned with the deadlock over production caps intact [5], [6]. Commentators attribute the collapse partly to a compressed timeline that forced negotiators to litigate implementation before agreeing scope [7], and argue increasingly that national instruments should not wait [8]. Implementation falls to domestic regulation.

In India that regulation is the Plastic Waste Management Rules and their amendments, which impose category-wise obligations on producers, importers and brand owners through a centralised EPR portal [9]. A producer discharges its obligation by acquiring certificates from registered processors. The quantity behind each certificate is attested, not demonstrated, and nothing lets a regulator, auditor or competitor trace a certificate

to the material it represents. The European instrument differs structurally — the ESPR mandates DPPs for regulated product groups [10] — but shares the verification problem.

1.1 Research gap

Recent reviews cover adjacent ground competently. DPP surveys map the regulatory and data-model landscape [11], bibliometric analyses chart growth after the European Green Deal [12], and structured reviews catalogue use cases by sector [13]. All three treat the digital record as the object of study. None treats the *binding between physical material and digital record* as the problem, and none brings the molecular marking literature into contact with the ledger literature. The gap addressed here is therefore not an absence of work on traceability, but an absence of work across the seam where traceability fails.

1.2 Contributions

We separate two categories. Items 1–4 are findings of the survey; item 5 is a design artefact derived from them and is not claimed as a survey result.

1. A systematic review of 45 sources (2023–2026) spanning four communities not previously reviewed together, under the protocol of Section 2 and appraised in Section 3.6.
2. A four-layer analytical framework and taxonomy organising the comparison in Section 3.
3. A comparative evaluation of six configurations against six capability criteria (Section 4.2), showing that none satisfies all six and isolating survivability through reprocessing as the neglected criterion.
4. Identification of the field’s limiting constraint as evaluative, with a specification of what an adequate benchmark requires (Section 5.3).
5. A reference framework instantiating the missing composition, each decision traced to supporting evidence (Section 4.3, Table 4).

Section 2 states the review protocol; Section 3 surveys the literature by layer and appraises its quality; Section 4 derives the evidence-derived reference framework; Section 5 discusses cross-cutting findings; Sections 6–8 cover advantages, social welfare and future work; Section 9 concludes.

2. REVIEW METHODOLOGY

2.1 Databases and search strategy

Searches covered IEEE Xplore, ScienceDirect, SpringerLink, MDPI and Nature Portfolio; ACM Digital Library coverage came by reference chaining. Four Boolean strings were used, one per layer, each with a 2023–2026 filter:

- **S1 (physical):** (“molecular taggant” OR “DNA taggant” OR “chemical tracer” OR “fluorescent marker”) AND (plastic OR polymer) AND (traceability OR authentication OR sorting)
- **S2 (identifier):** (“digital product passport” OR “QR code” OR “digital watermark” OR EPCIS) AND (plastic OR packaging OR “circular economy”)
- **S3 (ledger):** (blockchain OR “distributed ledger” OR “Hyperledger Fabric” OR “non-fungible token”) AND (“supply chain” OR traceability OR “waste management”)
- **S4 (analytics):** (“deep learning” OR “machine learning” OR CNN OR LSTM OR “anomaly detection”) AND (“waste classification” OR “waste generation forecasting” OR “plastic sorting” OR “fraud detection”)

A fifth string covering EPR and treaty instruments was applied to policy sources.

2.2 Inclusion and exclusion criteria

Inclusion. (i) Published 2023–2026; (ii) peer-reviewed journal article, peer-reviewed conference paper, or official standards or regulatory instrument; (iii) reports a technical method, systematic review, or binding requirement bearing on one of the four layers; (iv) English language; (v) full text retrievable.

Exclusion. (i) Vendor white papers, patents and trade press; (ii) unreviewed preprints; (iii) studies where the waste application is incidental; (iv) architectural proposals with neither evaluation nor systematic review; (v) duplicates and extended versions of included conference papers.

Standards and government instruments are exempt from criterion (iii), cited as primary sources for regulatory fact rather than research contributions.

2.3 Selection process

Figure 1 records the selection flow. Searches returned 312 records; deduplication removed 58; screening on title and abstract removed 143, mostly studies where waste was an incidental application domain plus proposals lacking evaluation; full-text assessment of the remaining 111 removed 66, chiefly for failing criterion (iii). Forty-five sources were retained. That figure is an outcome, not a target: each retained study directly addressed one or more layers of the analytical framework, and no source meeting that condition was excluded on grounds of count.




Figure 1. PRISMA-style flow of study identification, screening, eligibility assessment and inclusion. Counts at each stage appear in parentheses; boxes on the right record the number excluded and the dominant reason. Of 312 records identified across five databases, 45 were retained for synthesis.

2.4 Quality appraisal criteria

Included studies were appraised on five attributes: study type; evidence base; validation method; principal strength; principal limitation. No composite score was computed, the corpus spanning experimental, review and design-science work whose evidence is not commensurable on one scale. Results appear in Section 3.6, once the studies are introduced.

Of the 45 retained sources, 32 (71%) appeared in Q1/Q2 journals, 8 (18%) in peer-reviewed conference proceedings, and 5 (11%) are standards or regulatory instruments. Most date from 2024–2026, reflecting both the recent acceleration of research and, for the ledger and passport work, the EU ESPR entering force in July 2024.

3. LITERATURE SURVEY

This section compares the four layers instead of describing studies serially, and each subsection closes with a comparative statement. Figure 2 gives the organising taxonomy.

3.1 Regulatory instruments as engineering requirements

Regulatory instruments determine what a traceability system must prove, so we treat them as requirements. The two dominant regimes differ decisively: the ESPR mandates DPPs at product-group granularity with specified information and operational requirements [10], while the Indian regime is certificate-based, aggregating at the tonne [9]. Requirements analyses find the binding constraints are not storage or identification but data ownership, granularity and update semantics — who writes which fields at which lifecycle stage [14].

This structural difference has a consequence easy to overlook. A certificate regime creates a fungible tradable instrument, and fungibility is exactly what makes double-claiming attractive; the passport model carries no such property. Systems designed against the European requirement will not transfer to the Indian one until they address an adversary the European design never faces.

Deployment context adds constraints. MSW management in India combines a large informal collection sector with uneven municipal capacity and thin digital infrastructure at the collection point [15], and critical reviews of circular economy frameworks warn that formal traceability can displace informal livelihoods when it works as a gatekeeping mechanism [16]. Any proposal must therefore be judged on whom it excludes, not only on what it detects.

3.2 Physical layer: molecular taggants

A molecular taggant belongs to the material, not the packaging. Three families are in commercial or near-commercial use: fluorescent tracer particles read by modified NIR sorters, taggant masterbatches compounded into the polymer, and elemental or chemical signatures [3]. DNA taggants form a fourth, distinguished by carrying information as well as a signature [17]. Table 1 compares them on the dimensions determining deployability.

Table 1. Comparison of molecular taggant families.

Family	Information capacity	Read instrument	Survives reprocessing	Principal limitation
Fluorescent tracers [3]	Class signature	Modified NIR sorter	Partial	No article identity; sorter retrofit
Taggant masterbatch [3]	Batch level	Dedicated reader	Yes	Alters material spec; supplier lock-in
Elemental signature [3]	Signature only	Spectroscopic	Yes	Costly instrument; low discrimination
DNA taggant [17]–[19]	Article level, encoded	PCR/sequencing; lateral-flow emerging	Yes, encapsulated	Assay cost, latency, lab dependence

The evaluation the traceability literature has never performed concerns that final column. Surveys of the DNA tagging market document commercial deployment in brand protection and fuel marking, silica encapsulation supplying the robustness bare oligonucleotides lack [17]. Recent work moves the read operation towards a field test using hybridization-encoded tags with lateral-flow readout [18]. The most careful assessment of the trade-off space finds no universal configuration: encapsulation maximises protection at the cost of accessibility, and supply chain tracking needs cheaper, faster read paths than high-security applications tolerate [19].

Several practical limitations follow. **Read economics dominate:** the binding cost is assay, not tagging, incurred per verification rather than per article. **Latency precludes checkpoint use:** the fastest published readouts exceed the dwell time of a custody handover. **Regulatory acceptance is unestablished:** not one source in the corpus shows a molecular assay accepted as compliance evidence by any regulator. So a DNA taggant is an adjudication technology, suiting contested identity rather than routine operation; architectures proposing molecular verification at every transfer misread the economics in [19].

3.3 Identifier layer: optical codes and passports

Optical identifiers are the pragmatic counterpart: cheap, readable on commodity hardware, readable by non-specialists. Their weakness is categorical — a printed code belongs to the surface and does not survive shredding, washing or extrusion [3]. That asymmetry against Section 3.2 is the structural fact this survey turns on.

Interoperability here rests on the GS1 EPCIS standard, whose 2.0 revision defines object, aggregation, transaction and transformation event types permitting exchange without bilateral integration [20]; recent work extends it towards knowledge-graph representation [21]. Against that substrate, the three most substantial DLT-backed passport contributions each solve something different. A two-year design-science study produced an end-to-end DPP prototype for textiles, articulating principles for binding upstream transformation to downstream record [22]. Multi-blockchain implementations with decentralised identifier providers address identity portability across consortia [23]. An NFT-enabled passport for construction waste demonstrates cross-jurisdictional trading, the closest analogue to EPR certificate trading [24].

These solve identity, portability and transferability respectively; none addresses persistence. The passport literature has optimised every property of the identifier bar the one that fails first.

3.4 Ledger layer

DLT is the most mature of the four literatures, and the most prone to overclaiming. The question is not immutability, which holds by construction, but whether throughput and operational complexity suit programme-scale volumes.

Comparative analysis of Ethereum and Hyperledger Fabric establishes the order-of-magnitude gap between public and permissioned architectures under equivalent workloads [25]. A cold-chain implementation reports stable operation under sustained load, query transactions substantially outperforming invoke transactions [26]; that asymmetry matters, because verification queries dominate once a traceability ledger is populated. Benchmarking on heterogeneous hardware — the realistic municipal condition — finds throughput sensitive to peer specification and endorsement policy in ways idealised benchmarks conceal [27]. A food-processing framework demonstrates the multi-organisation governance EPR requires [28], and a survey of Fabric business cases catalogues recurring deployment patterns and failure modes [29].

Reading these together exposes two things the individual papers do not. The reported figures are not comparable: hardware, endorsement policy, topology and workload mix all differ, and no study reports the others' configuration. The second point matters more — this literature evaluates ledger performance in isolation from the capture layer feeding it, so the end-to-end latency an operator would experience goes unreported throughout.

Storage is the standard constraint. On-chain storage grows without bound and every full node bears the cost, so evidence artefacts sit off-chain with only a content identifier committed. Work on collaborative on-chain/off-chain storage using IPFS quantifies the reduction in ledger growth

[30]. Its integrity guarantee is routinely overstated: committing a content identifier proves the artefact unaltered since commitment, but says nothing about authenticity at creation.

3.5 Analytical layer

Classification. Hyperspectral sensing with machine learning has been evaluated across eight polymer types including bioplastics, under artificial and natural illumination — illumination proving a material determinant of accuracy [31] — and applied to mixed municipal streams for stream-level characterisation [32]. On commodity hardware, MobileNetV2 shows lightweight architectures suited to edge inference remain competitive [33], and hybrid CNN ensembles report higher accuracy on curated datasets [34]. Detection-and-classification pipelines have been assessed for sorting-line deployment [35], and shape-based detection proposed for chemically similar polymers such as PET and PET-G that defeat NIR discrimination [36]. A recent review situates these against the computational bottleneck of industrial deployment [37], and comparative k-fold evaluation gives the most defensible relative ranking available [38].

Grouped rather than listed, these studies divide by sensing modality far more than by algorithm — and modality carries the deployment constraint. Spectroscopic approaches achieve material discrimination unavailable to RGB, at instrumentation and throughput costs that shut them out of the collection point. That trade-off, not accuracy, is the operative comparison. No individual study makes it.

Forecasting. Multi-model approaches to MSW generation and recycling demand project capacity over decadal horizons [39]. LSTM models with downstream risk assessment report mean absolute percentage errors in low single digits against multi-decade series [40]; transformer and multi-layer perceptron approaches are competitive [41]. A systematic review notes the field's dependence on socioeconomic covariates that are themselves forecast [42], and bidirectional LSTM ensembles address regime change [43].

Set side by side, these studies share a resolution mismatch none of them frames as a limitation. All forecast *generation*, at municipal or regional aggregation and monthly or annual resolution. Forecasting *accumulation* at grid-cell and weekly resolution — where collection decisions are made — is a different problem, and no reviewed source establishes the transfer.

Anomaly detection. Labelled instances of fraudulent compliance claims do not exist in quantity, so applicable methods are unsupervised or semi-supervised. Semi-supervised supply chain fraud detection addresses this label-scarcity regime by iteratively expanding a small labelled set [44]. A critical analysis combining review with empirical validation finds supervised ensembles dominate where labels exist, isolation-based methods where they do not [45].

Set reported sensitivity against unreported precision and the difficulty is plain. In an unlabelled compliance setting the fraud base rate is unknown, so the flag-to-violation ratio cannot be estimated in advance — and a false positive costs the audited party, not the system operator. No reviewed source reports a calibration procedure for this asymmetry.

3.6 Quality appraisal of the reviewed evidence

Table 2 applies the criteria of Section 2.4 to a representative subset, selected to span the four layers and every evidence type in the corpus — primary experimental work, critical review, design science, comparative benchmarking and method development — so the appraisal reflects the range of warrant on offer rather than the strongest instances of any one kind.

Table 2. Quality appraisal of representative included studies.

Ref.	Study type	Evidence base	Validation	Strength	Limitation
[17]	Market/technology review	Published and commercial sources	Narrative synthesis	Breadth across deployed applications	No primary experiment
[18]	Primary experimental	Laboratory-synthesised tag library	Laboratory readout trials	Field-deployable read path	Not tested in a waste stream
[19]	Critical review	Cross-technology literature	Comparative synthesis	Explicit trade-off analysis	No quantitative benchmark
[22]	Design science	Textile supply chain instantiation	Two-year prototype demonstration	End-to-end coverage	Single sector; no external validation
[25]	Comparative benchmark	Synthetic transaction workload	Throughput and latency measurement	Direct platform comparison	Configuration-specific results
[31]	Primary experimental	Eight polymer types, two lighting conditions	Accuracy on own held-out data	Quantifies illumination effect	Laboratory samples; no shared benchmark
[33]	Primary experimental	Public waste-image dataset	Hold-out accuracy	Lightweight, edge-deployable	Curated imagery
[40]	Experimental modelling	Multi-decade municipal series	Error vs. observed	Low reported error	Single city; annual resolution
[44]	Method development	Supply chain transaction data	Semi-supervised evaluation	Addresses label scarcity directly	Not a waste-domain dataset

The appraisal exposes a pattern worth naming. Validation is internal in every primary study — each reports against its own data, none against a shared reference. And the strongest evidence sits in the layers of weakest deployment relevance: laboratory-validated molecular work and benchmark-validated ledger work both stand further from the collection point than the conditions either would face.




Figure 2. Taxonomy of plastic waste traceability enablers derived from the 45 reviewed sources, organised into the four layers used throughout Section 3. Each layer answers a distinct question: the physical layer what can be asserted about the material, the identifier layer who may assert it, the ledger layer whether the assertion can later be revised, and the analytical layer whether it is plausible. Bracketed examples name technologies appearing in the corpus.

4. EVIDENCE-DERIVED REFERENCE FRAMEWORK

4.1 Cross-layer synthesis

Table 3 consolidates Section 3.

Table 3. Synthesis across the four layers.

Layer	Dominant methods	Strength	Characteristic failure	Compensated by
Physical [3], [17]–[19]	Tracers; masterbatch; DNA taggants	Persists through reprocessing	Read cost and latency preclude routine use	Identifier layer
Identifier [3], [20]–[24]	Optical codes; DPPs; EPCIS vocabulary	Zero-instrument, non-specialist read	Destroyed by reprocessing	Physical layer
Ledger [25]–[30]	Permissioned ledgers; token registries; IPFS	Prevents retroactive revision	Cannot assess truth at time of writing	Analytical layer
Analytical [44], [31]–[45]	CNN; LSTM; unsupervised detection	Assesses plausibility of claims	Requires a trustworthy record to assess	Ledger layer

Look down the last column and the regularity is hard to miss: in every row the failure is compensated by a layer the study in question does not employ. We term this the *composition gap*. It differs from the limitations documented within each literature, and no improvement to any single component will close it.

4.2 Comparative capability assessment

Figure 3 evaluates six configurations against six criteria derived from Sections 3.1–3.5.

Selected enablers (status quo)	 Item-level granularity	 Sentiment transparency	 Reduces unnecessary recalls	 Verifiable by auditors	 Low margins / cost savings	 Usable without specialists
 Serialized items (status quo)	none	none	none	none	strong	strong
 Centralized EPC global data repository	weak	none	weak	weak	strong	strong
 Optical code + centralized database	strong	none	weak	weak	strong	strong
 Optical code + permissioned ledger	strong	none	strong	strong	partial	strong
 Molecular taggant (DNA)	partial	strong	partial	partial	none	none
 Taggant + code + ledger + analytics	strong	strong	strong	strong	partial	partial

 strong High capability / benefit
  partial Moderate capability / benefit
  none / weak Low or no capability / benefit

Figure 3. Capability comparison of six traceability configurations against six criteria. The criteria derive from the regulatory requirements of Section 3.1 (item-level granularity, external verifiability) and the technical constraints of Sections 3.2–3.5 (survivability, resistance to unilateral revision, marginal cost, specialist independence). Ratings were assigned through qualitative synthesis of evidence from the reviewed literature rather than quantitative scoring, and are not measurements. The discriminating column is *survives reprocessing*, where every configuration built on optical identifiers alone fails categorically.

The comparison is informative in three ways. Aggregate self-declaration fails every evidentiary criterion yet scores well on cost and usability — which explains its persistence rather better than accounts resting on weak enforcement. Adding a permissioned ledger to an optical-code system improves revision-resistance and external verifiability, but changes nothing about survivability, a property fixed at the physical layer. And no configuration reaches the highest rating on all six; the composite comes closest only by accepting moderate cost and specialist-independence.

4.3 A reference framework and its evidential basis

Table 4 states each design decision alongside the evidence requiring it, answering the concern that architectural proposals in this field are asserted rather than derived.

Table 4. Design decisions and their evidential basis.

Design decision	Evidence	Reasoning
Assay as exception path	[19]	Assay cost and latency govern taggant economics
Optical identifier for routine events	[3], [20]	Only option readable without instrumentation
Permissioned ledger for custody	[25]–[29]	Throughput and role-based write authority
Public token registry for compliance	[9], [24]	Irreversible retirement removes fungibility
Content-addressed off-chain evidence	[30]	On-chain artefact storage is unbounded
Lightweight CNN at capture	[33], [37]	Spectroscopic accuracy unavailable in the field
Anomaly score to human review	[44], [45]	Precision unmeasurable without labels
Commodity smartphone capture	[15], [16]	Specialised terminals exclude informal collectors

Figure 4 presents the resulting layered structure.




Figure 4. Evidence-derived reference framework comprising six operational layers mapped onto the four analytical layers introduced in Section 3, with capture and reporting separated out. L1–L2 are physical and operational; L3–L4 provide the tamper-evident record; L5 interpretation and L6 the regulatory interface. Cross-cutting concerns at the foot — credentialing, key custody, location privacy, event vocabulary — belong to no single layer and are the most common source of integration failure.

At L1 each article carries an optical identifier and, for regulated categories, a taggant compounded into the polymer; the identifier encodes article, manufacturer, declared resin, batch and a salted hash of the taggant sequence — never the sequence itself. At L2 an operator reads the identifier on commodity hardware, the device recording image, position, time and credential. At L3 the event is committed as an EPCIS 2.0 event [20], [21], with a token operation where compliance status changes. At L4 image and position records go to content-addressed storage, only the identifier on-chain [30]. At L5 three models operate: resin inference [33], [34], accumulation forecasting [40], [43], and plausibility scoring [44], [45]. At L6 evidence is rendered in the regulator’s format, every figure resolvable to constituent events.

Operationally the framework executes in six sequential stages: article identification, capture, ledger registration, evidence storage, analytical assessment, and regulatory reporting. Molecular verification is invoked only as an exception path when identity is disputed, per the read economics of [19]; anomaly detection triggers human review only for flagged claims, per the precision limitation of [44], [45]. Ordering matters within that sequence: classification precedes ledger commitment so the inferred resin class enters the immutable record, evidence pinning precedes it so the content identifier is available to the transaction, and scoring follows commitment, a claim being well-formed only once its constituent events exist.

4.4 Scope and status of the framework

Table 4 specifies a framework, not a validated system. Its warrant is derivational: every decision follows from evidence in the surveyed corpus, the reasoning stated so both can be contested — a weaker warrant than empirical validation, and no substitute for it. Validation is deliberately outside this survey’s scope, for the reason the survey itself identifies: without shared benchmarks, an isolated evaluation of one system would reproduce the incommensurability of Section 5.3 instead of remedying it.

Three limitations are structural. Taggant synthesis and assay require a molecular marking partner and lie outside any software specification. Accumulation forecasting faces a cold start, since the event stream it depends on does not exist before deployment. The anomaly component cannot be validated without labelled fraud instances, so no detection rate is claimable.

5. RESULTS AND DISCUSSION

This section reports three results emerging from the corpus as a whole, not visible in any individual study.

5.1 Reported ledger throughput does not scale to programme demand

Table 5 collects the reported throughput.

Table 5. Transaction throughput reported in the reviewed ledger literature.

Study	Platform and configuration	Reported throughput (tx/s)
[25]	Ethereum, public network	~30
[27]	Hyperledger Fabric, heterogeneous hardware	~300
[26]	Hyperledger Fabric, cold-chain (query transactions)	~442
[30]	Hyperledger Fabric with IPFS off-chain storage	~1500

Note: Reported transaction throughput values originate from independent studies conducted under different hardware configurations, endorsement policies, and workloads; they should therefore be interpreted as indicative rather than directly comparable. Set against demand, they reassure at municipal scale and less so beyond it. Assuming six checkpoints per article across a year — an arithmetic projection, not a forecast — a single ward generates well under one transaction per second, while a state-level programme handling a billion articles annually approaches several hundred. Although these values appear sufficient under average operating conditions, real deployments experience peak workloads that may substantially reduce effective throughput. Hence the case for hierarchical or partitioned deployment over a monolithic ledger — an argument no individual study can make, each evaluating one network alone.

5.2 Sensing modality, not architecture, determines deployability

Table 6 collects the reported accuracy.

Table 6. Classification accuracy reported in the reviewed sorting studies.

Study	Method	Sensing modality	Reported accuracy range (%)
[33]	MobileNetV2, 6-class recyclables	RGB	88–96
[34]	Hybrid and enhanced CNN ensemble	RGB	90–97
[35]	Detection-and-classification pipeline	RGB	84–93
[36]	Shape-based detection, PET vs. PET-G	RGB	91–99
[31]	Classical ML, eight polymer types	Hyperspectral	88–98
[32]	Machine learning, mixed MSW stream	Hyperspectral	85–96

Values span the configurations each study reports on its own held-out data. Although generally high, they are not comparable between rows: datasets, acquisition conditions and evaluation protocols differ substantially, and no common test set exists. What the table does show is that accuracy varies less by architecture than by modality — and modality carries the deployment constraint: spectroscopic methods achieve discrimination unavailable to RGB, at an instrumentation cost and data volume that exclude the collection point [37]. A designer choosing on reported accuracy alone would select on the least informative dimension available.

5.3 The evaluation gap and what an adequate benchmark requires

This survey’s sharpest result is negative, and Tables 5 and 6 are its evidence. No two classification studies share an evaluation set; no two ledger studies share a configuration. Section 3.6 shows the same from the appraisal side. Anomaly detection is worse, absent ground truth making detection performance unmeasurable so studies fall back on proxies [44], [45]. Forecasting is healthier — municipal series are publicly reported, errors computed against observation [40]–[42] — but the resolution mismatch of Section 3.5 persists.

An adequate benchmark would require four components. A **common dataset** captured at operational collection points rather than in the laboratory, spanning the resin classes the regulation distinguishes, with documented acquisition conditions and class balance. A **standard protocol** fixing the train/validation/test partition, prohibiting test-set reuse for model selection, and specifying how out-of-distribution material is handled. **Shared metrics** matched to the operational cost structure — for compliance, precision and false-positive burden alongside accuracy, since a false flag costs the audited party. And a **reproducible reference implementation** with released code, fixed seeds and declared hardware. Addressing this evaluation gap would substantially improve the reproducibility and comparability of future plastic waste traceability research.

None of these exists. The field therefore cannot answer the question a designer most needs answered — which method to select — and the maturity signalled by high reported accuracies is partly an artefact of evaluation on favourable data. Benchmark construction is consequently the field’s most urgent methodological need, and a more valuable near-term contribution than any incremental architectural proposal, including the framework of Section 4.3.

5.4 Threats to validity

Several caveats bear on how far these findings travel. The corpus covers only English-language sources in the five databases of Section 2.1, under-representing regional venues in the jurisdictions carrying the highest waste burdens. ACM coverage came by chaining rather than direct query, so the computing-for-sustainability community may be thinner here than it should be. The four-layer framework is our imposition on a literature that does not organise itself this way; a different decomposition would foreground different gaps. The ratings in Figure 3 and the appraisal in Table 2 are also qualitative judgements over heterogeneous evidence, and readers weighting the criteria differently would reach different conclusions.

6. ADVANTAGES OF THE EVIDENCE-DERIVED FRAMEWORK

These are architectural properties entailed by the framework of Section 4.3, each with its enabling condition — not empirical claims.

Item-level evidence. Every custody transfer yields a discrete event bound to an article identifier, so a compliance figure aggregates independently inspectable records — conditional on checkpoint coverage, since a chain with gaps yields aggregations no more trustworthy than the declarations they replace.

Structural prevention of double-claiming. Compliance status held as a token whose retirement is irreversible causes a second claim to fail at submission, addressing the fungibility problem of Section 3.1 [9], [24]. Provided that the article entered the registry at production.

Verification without trusting the operator. A public registry and content-addressed evidence let an outsider confirm that a stated figure follows from committed events and that the referenced evidence is unaltered [30]. Accuracy at the time of writing remains unestablished — the residual trust assumption, and the reason the physical layer exists.

An appeal mechanism for reprocessed material. Once material is reprocessed the optical identifier is gone, and a material-resident taggant is the only reviewed technology with anything to offer [17]–[19].

Inherited interoperability. The EPCIS 2.0 vocabulary [20] permits integration without bilateral mapping and makes the event stream consumable by standard-conformant tools [21].

Bounded participation cost and auditable output. A commodity smartphone as capture device keeps informal collectors inside the system — decisive for equitable implementation [15], [16] — while committing the classifier’s inference alongside the event gives later composition disputes a tamper-evident record naming the responsible model version.

7. SOCIAL WELFARE IMPLICATIONS OF THE EVIDENCE-DERIVED FRAMEWORK

The welfare argument concerns who bears the cost of an unverifiable system. Each point ties to reviewed evidence; none projects outcomes.

Environmental integrity of compliance. Where obligations are dischargeable on paper, the marginal tonne a producer is nominally responsible for may never be recovered [9]. Verifiability shifts the cheapest compliance route from paperwork to operations. The size of that shift is unknowable in an unverifiable system — which is itself the argument for building the measurement.

Occupational recognition for informal collectors. Much of the plastic recovered in India passes through informal hands whose contribution is invisible in official statistics [15]. A credentialed collection event produces a durable record of individual contribution as a by-product — which depends on inclusive credentialing, making the smartphone constraint of Table 4 a welfare requirement rather than a convenience [16].

Public accountability. A public registry lets civil society, journalists and researchers audit corporate claims without institutional cooperation — among the few instruments available while the international process stalls [5]–[8].

Service efficiency and exposure reduction. Accumulation forecasting, subject to the resolution caveat of Section 3.5, would let collection resources be positioned prospectively, with gains accruing disproportionately to under-served areas — the same settlements where uncollected plastic drives drain blockage, localised flooding and open burning [15].

Traceability also generates location traces of individuals engaged in collection work, and such traces are sensitive. The welfare case is contingent on privacy design that no reviewed source provides, and we treat it as open.

8. FUTURE ENHANCEMENT

This survey proposes neither an implementation nor a validation study. Instead, it identifies several research needs based on the reviewed literature. These include the development of a shared classification benchmark under real deployment conditions, field-deployable molecular read paths, and end-to-end latency evaluation from molecular readout to ledger commitment. The review also highlights the need for operational-scale accumulation forecasting, privacy-preserving analytics using federated learning, and reliable calibration methods for unlabelled anomaly detection. Finally, it suggests extending the proposed approaches to other waste streams, including electronic waste, batteries, and textiles, which face similar verification challenges.

9. CONCLUSION

EPR requires producers to prove something that current systems cannot check. The principal obstacle is the absence of composition among existing approaches, not the absence of any particular technology.

Under a PRISMA-style protocol we screened 312 records and reviewed 45 sources drawn from four literatures that do not cite one another. No reviewed configuration satisfies all six capability criteria derived from the regulatory requirements, and the discriminating criterion — survival through mechanical reprocessing — is failed categorically by every identifier-only approach. Each layer’s characteristic failure is compensated by a layer that studies in that layer do not employ. The limiting constraint, though, is evaluative: appraisal confirms that validation is internal in every primary study, so reported accuracies and throughputs are incommensurable, and claims about the maturity of AI-enabled waste traceability describe particular datasets instead of the world. We specify the four components an adequate benchmark requires.

Separately, we derive a reference framework instantiating the missing composition, each decision traced to its evidential basis. That framework is a specification offered as such; its validation, and the benchmark work that would make validation meaningful, remain open.

ACKNOWLEDGEMENTS

The authors thank the Department of Computer Science and Engineering, Sandip University, Nashik, for institutional support.

Conflicts of interest: The authors declare no conflict of interest.

Data availability: No new data were generated or analysed. The search strings and selection criteria of Section 2 permit replication of the corpus.

Funding: This research received no external funding.

REFERENCES

- [1]. B. Gazeau, A. Zaman, R. Minunno, and F. Shaikh, "Developing Traceability Systems for Effective Circular Economy of Plastic: A Systematic Review and Meta-Analysis," *Sustainability*, vol. 16, no. 22, pp. 9973–9973, Nov. 2024, doi: 10.3390/su16229973.
- [2]. U. N. Environment, "Turning off the Tap: How the world can end plastic pollution and create a circular economy," May 2023. Accessed: July 23, 2026. [Online]. Available: <https://www.unep.org/resources/turning-off-tap-end-plastic-pollution-create-circular-economy>
- [3]. T. Rumetshofer and J. Fischer, "Information-Based Plastic Material Tracking for Circular Economy—A Review," *Polymers*, vol. 15, no. 7, pp. 1623–1623, Mar. 2023, doi: 10.3390/polym15071623.
- [4]. K. Bułkowska, M. Zielińska, and M. Bułkowski, "Blockchain-Based Management of Recyclable Plastic Waste," *Energies*, vol. 17, no. 12, pp. 2937–2937, June 2024, doi: 10.3390/en17122937.
- [5]. S. Pottinger *et al.*, "Pathways to Reduce Global Plastic Waste Mismanagement and Greenhouse Gas Emissions by 2050," *Science*, vol. 386, no. 6726, Nov. 2024, doi: 10.1126/science.adr3837.
- [6]. M. Spring *et al.*, "Effective progress and implementation of the INC-5 plastics treaty through scientific guidance," *Nature Sustainability*, vol. 8, no. 7, pp. 728–730, May 2025, doi: 10.1038/s41893-025-01574-0.
- [7]. F. Shah and W. Wu, "Overcoming roadblocks to a global plastic pact," *Science*, vol. 388, no. 6744, pp. 259–259, Apr. 2025, doi: 10.1126/science.adv0976.
- [8]. S. Winton and A. March, "Beyond the Global Plastics Treaty," *Nature Reviews Earth & Environment*, Dec. 2025, doi: 10.1038/s43017-025-00752-0.
- [9]. Ministry of Environment, Forest and Climate Change, Government of India, *Plastic Waste Management (Amendment) Rules*, G.S.R. 133(E), Feb. 2022; as amended by G.S.R. 201(E), Mar. 2024.
- [10]. "Regulation - EU - 2024/1781 - EN - EUR-Lex," 2024. Accessed: July 23, 2026. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2024/1781/oj/eng>
- [11]. F. Psarommatis and G. May, "A Comprehensive Review of Digital Product Passports in Sustainable Manufacturing: Current State, Challenges, and Future Directions," *Circular Economy and Sustainability*, vol. 6, no. 3, Apr. 2026, doi: 10.1007/s43615-026-00927-x.
- [12]. N. Karabulut, Gamzegül Çalıkoğlu, and Zeki Atıl Bulut, "Discussing and Reviewing the Digital Product Passport: An Up-to-Date Bibliometric Analysis," *Systems*, vol. 13, no. 11, pp. 930–930, Oct. 2025, doi: 10.3390/systems13110930.
- [13]. Zhang and S. Seuring, "Digital product passport for sustainable and circular supply chain management: a structured review of use cases," *International Journal of Logistics Research and Applications*, vol. 27, no. 12, pp. 1–28, July 2024, doi: 10.1080/13675567.2024.2374256.
- [14]. M. Jansen, T. Meisen, C. Plociennik, H. Berg, A. Pomp, and W. Windholz, "Stop Guessing in the Dark: Identified Requirements for Digital Product Passport Systems," *Systems*, vol. 11, no. 3, pp. 123–123, Feb. 2023, doi: 10.3390/systems11030123.
- [15]. Anitha Rengerla and Parthiban Angamuthu, "Evolving dynamics of municipal solid waste management in India: a comprehensive review of status, challenges, opportunities and policy innovations," *Waste Disposal & Sustainable Energy*, vol. 7, no. 1, pp. 127–150, Mar. 2025, doi: 10.1007/s42768-024-00220-3.
- [16]. F. B. Awino and S. E. Apitz, "Solid Waste Management in the Context of the Waste Hierarchy and Circular Economy Frameworks: An International Critical Review," *Integrated Environmental Assessment and Management*, vol. 20, no. 1, Apr. 2023, doi: 10.1002/ieam.4774.
- [17]. Kuzdraliński, M. Miśkiewicz, H. Szerba, W. Mazurczyk, J. Nivala, and B. Książkowski, "Unlocking the potential of DNA-based tagging: current market solutions and expanding horizons," *Nature Communications*, vol. 14, no. 1, p. 6052, Sept. 2023, doi: 10.1038/s41467-023-41728-2.
- [18]. J. Li *et al.*, "Hybridization-encoded DNA tags with paper-based readout for anti-forgery raw material tracking," *Nature Communications*, vol. 16, no. 1, July 2025, doi: 10.1038/s41467-025-60282-7.
- [19]. Kuzdraliński *et al.*, "Advancements in <sc>DNA</sc> Tagging and Storage: Techniques, Applications, and Future Implications," *Wiley Interdisciplinary Reviews Computational Molecular Science*, vol. 15, no. 4, July 2025, doi: 10.1002/wcms.70040.
- [20]. "EPCIS & CBV | GS1." Accessed: July 23, 2026. [Online]. Available: <https://www.gs1.org/standards/epcis>
- [21]. Y. K. Tolcha, G. Park and D. Kim, "Building Digital Product Passports: EPCIS 2.0 Meets Knowledge Graphs," *2025 IEEE International Conference on RFID (RFID)*, Atlanta, GA, USA, 2025, pp. 1-6, doi: 10.1109/RFID64926.2025.11015538.
- [22]. H. Abreu, V. Pereira, and J. Barata, "Blockchain-based digital product passport: design principles and demonstration," *International Journal of Production Research*, pp. 1–20, Feb. 2025, doi: 10.1080/00207543.2025.2464161.
- [23]. Mihai Hulea, R. Miron, and V. Muresan, "Digital Product Passport Implementation Based on Multi-Blockchain Approach with Decentralized Identifier Provider," *Applied Sciences*, vol. 14, no. 11, pp. 4874–4874, June 2024, doi: 10.3390/app14114874.
- [24]. L. Wu, W. Lu, Z. Peng, and C. Webster, "A blockchain non-fungible token-enabled 'passport' for construction waste material cross-jurisdictional trading," *Automation in Construction*, vol. 149, p. 104783, May 2023, doi: 10.1016/j.autcon.2023.104783.
- [25]. Yusuf Ucbas, Amna Eleyan, M. Hammoudeh, and Manar Alohaly, "Performance and Scalability Analysis of Ethereum and Hyperledger Fabric," *IEEE Access*, vol. 11, pp. 67156–67167, Jan. 2023, doi: 10.1109/access.2023.3291618.
- [26]. E. Duman and Ebru Aydoğan, "Enhancing Traceability and Reliability in Cold Chain Logistics Through Hyperledger Fabric and IoT," *Applied Sciences*, vol. 15, no. 22, pp. 12149–12149, Nov. 2025, doi: 10.3390/app152212149.
- [27]. M. M. Zulkarnain, Nabilah Ramli, and A. N. Nordin, "Performance Benchmarking of Hyperledger Fabric on Heterogeneous Hardware for IoT Applications," *IIUM Engineering Journal*, vol. 26, no. 3, pp. 156–170, Sept. 2025, doi: 10.31436/iiumej.v26i3.3610.
- [28]. M. Rahaman, F. Tabassum, V. Arya, and R. Bansal, "Secure and Sustainable Food Processing Supply Chain Framework Based on Hyperledger Fabric Technology," *Cyber Security and Applications*, pp. 100045–100045, Mar. 2024, doi: 10.1016/j.csa.2024.100045.
- [29]. O. Pandian and Y.-T. Song, "A Survey Of Hyperledger Fabric For Business Use Cases," *2025 IEEE/ACIS 23rd International Conference on Software Engineering Research, Management and Applications (SERA)*, pp. 80–88, May 2025, doi: 10.1109/sera65747.2025.11154535.
- [30]. T. Zhu *et al.*, "Research on an On-Chain and Off-Chain Collaborative Storage Method Based on Blockchain and IPFS," *Future Internet*, vol. 18, no. 2, p. 92, Feb. 2026, doi: 10.3390/fi18020092.
- [31]. M. Moroni, M. Balsi, and Soufyane Bouchelaghem, "Plastics detection and sorting using hyperspectral sensing and machine learning algorithms," *Waste Management*, vol. 203, pp. 114854–114854, May 2025, doi: 10.1016/j.wasman.2025.114854.
- [32]. J. Tao *et al.*, "Combination of hyperspectral imaging and machine learning models for fast characterization and classification of municipal solid waste," *Resources, Conservation and Recycling*, vol. 188, p. 106731, Jan. 2023, doi: 10.1016/j.resconrec.2022.106731.
- [33]. L. Yong, Le Ma, D. Sun, and L. Du, "Application of MobileNetV2 to waste classification," vol. 18, no. 3, pp. e0282336–e0282336, Mar. 2023, doi: 10.1371/journal.pone.0282336.

- [34]. S. Shetty, S. Kallianpur, R. Fernandes, A. P. Rodrigues, and V. Padmanabha, "ECO-HYBRID: Sustainable Waste Classification Using Transfer Learning with Hybrid and Enhanced CNN Models," *Sustainability*, 2025, doi: 10.3390/su17198761.
- [35]. Faizul Rakib Sayem *et al.*, "Enhancing waste sorting and recycling efficiency: robust deep learning-based approach for classification and detection," *Neural Computing and Applications*, Dec. 2024, doi: 10.1007/s00521-024-10855-2.
- [36]. J. Choi, B. Lim, and Y. Yoo, "Advancing Plastic Waste Classification and Recycling Efficiency: Integrating Image Sensors and Deep Learning Algorithms," *Applied Sciences*, vol. 13, no. 18, p. 10224, Jan. 2023, doi: 10.3390/app131810224.
- [37]. F. Anchieta e Silva, A. de S. Cartaxo, A. D. de S. R. Esmeraldo, E. M. Senra, and J. C. Pinto, "Advancements in Plastic Waste Sorting: A Review of Techniques and Applications," *Processes*, vol. 14, no. 7, p. 1144, Apr. 2026, doi: 10.3390/pr14071144.
- [38]. Y. Sahu and N. K. Sahu, "Smart Waste Management: Deep Learning-Based Garbage Classification for Sustainability," *Proceedings of the ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies*, pp. 672–676, July 2025, doi: 10.1145/3715335.3736312.
- [39]. P. Xu and H. Zheng, "A multi-AI approach to predicting municipal solid waste generation and recycling demand in Hong Kong," *Resources, Conservation and Recycling*, vol. 225, p. 108590, Jan. 2026, doi: 10.1016/j.resconrec.2025.108590.
- [40]. Liu, J. Chen, F. Yang, J. Li, and B. Yan, "Prediction of municipal solid waste generation and environmental risk assessment of heavy metal based on long short term memory," *Environmental Earth Sciences*, vol. 83, no. 13, June 2024, doi: 10.1007/s12665-024-11702-2.
- [41]. N. Kumar, P. Rajeswari, D. Jeya Priya, and M. Uma Maguesvari, "Forecasting municipal solid waste generation using advanced transformer and multi-layer perceptron techniques," *Clean Technologies and Environmental Policy*, Dec. 2024, doi: 10.1007/s10098-024-03091-8.
- [42]. Z. X. Hoy, Z. X. Phuang, A. A. Farooque, Y. Van Fan, and K. S. Woon, "Municipal solid waste management for low-carbon transition: A systematic review of artificial neural network applications for trend prediction," *Environmental Pollution*, vol. 344, p. 123386, Mar. 2024, doi: 10.1016/j.envpol.2024.123386.
- [43]. M. S. Jassim, Gulnur Coskun, N. Sultana, and S.M. Zakir Hossain, "Forecasting domestic waste generation during successive COVID-19 lockdowns by Bidirectional LSTM super learner neural network," *Applied Soft Computing*, vol. 133, pp. 109908–109908, Dec. 2022, doi: 10.1016/j.asoc.2022.109908.
- [44]. L.-T. Zhao, D.-S. Wang, F.-Y. Liang, and J. Chen, "A recommendation system for effective learning strategies: An integrated approach using context-dependent DEA," *Expert Systems with Applications*, vol. 211, p. 118535, Jan. 2023, doi: 10.1016/j.eswa.2022.118535.
- [45]. Fouzia Juman and M. Raza, "Machine Learning for Anomaly Detection in Blockchain: A Critical Analysis, Empirical Validation, and Future Outlook," *Computers*, vol. 14, no. 7, pp. 247–247, June 2025, doi: 10.3390/computers14070247.