The Wear Index

An Open Specification for Objective, Machine-Produced Condition Grading of Used Consumer Electronics

Version 0.1 — Draft for Public Comment · July 2026

Khawar Iftikhar

Founder and Technical Principal, Dubai, United Arab Emirates

Correspondence: khawar.j@pm.me

STATUS OF THIS DOCUMENT. Version 0.1 is a draft published for public comment. It specifies a grading method, record format, conformance levels, and validation requirements; it is a statement of specification and design intent, not a report of results. No performance claims are made in this document, and none should be inferred from it.

The imaging-and-classification method on which this specification draws is under development and is the subject of a pending United States provisional patent application (filed May 27, 2026); it is described here at the level of design. As of this draft, no validated implementation of this specification exists (see Sections 9 and 10). Comments on every element of this draft are invited; see Section 11.

Abstract

Condition grading in the United States secondary market for consumer electronics is subjective: grades are assigned by eye, scales differ across marketplaces, and the same device can carry a different grade with every seller. This draft specification proposes the Wear Index, an objective, machine-produced, auditable condition score computed from images captured under controlled cross-polarized illumination and classified by a trained convolutional neural network. It defines a three-sub-score model with a 0–100 composite and fixed grade bands; informative mappings to existing market vocabularies, including the CTIA Wireless Device Grading lexicon; a Device Condition Certificate that binds every grade to content-addressed evidence and versioned models; three conformance levels for adopters of different capacity; and a validation methodology that any implementation — including the author's own — must publish before describing itself as validated. Version 0.1 is a draft for public comment and makes no performance claims.

Keywords: device grading; refurbished electronics; computer vision; polarized-light imaging; condition certification; circular economy; interoperability; electronic waste.

Suggested citation: *Iftikhar, K. (2026). The Wear Index: An Open Specification for Objective, Machine-Produced Condition Grading of Used Consumer Electronics (Version 0.1, draft for public comment).*

1. The Trust Problem in the Secondary Device Market

The United States secondary market in consumer electronics is large by any measure: industry estimates place the used and refurbished mobile-phone market at roughly 84.6 million units in 2024, projected to reach about 112 million units by 2030 [7]. Tens of millions of devices change hands every year, and every one of those transactions rests on a single load-bearing artifact: the condition grade. The grade sets the price, frames the buyer's expectation, and determines whether the transaction ends in satisfaction or in a return.

That artifact is produced by human judgment. A grader looks at a device and chooses a word. Two graders looking at the same device often choose different words. The words themselves are not shared across the market. One platform sells in “Excellent / Good / Fair,” another in “Premium / Excellent / Good / Acceptable,” and no mapping between them is defined by method — only by each seller's private habit. Some damage classes compound the problem because they are genuinely difficult for the eye: a hairline crack hidden under specular glare on glossy glass, or the early stage of a display separating from its frame, can pass visual inspection and surface only after sale.

The consequences are structural rather than incidental. Inconsistent grading drives returns, suppresses prices below what verified condition would support, and corrodes buyer trust in refurbished hardware as a category — across the whole market, not within any one company. The environmental stakes ride along with the commercial ones. The world generated a record 62 million tonnes of electronic waste in 2022, of which only 22.3 percent was documented as formally collected and recycled [5]; the United States has no federal standard requiring consumer-electronics recycling at all [6]. Every device kept in trusted reuse is a device diverted from that stream. Trust is the constraint, and the grade is where trust is made or lost.

What the market lacks is not another adjective. It is a number: a machine-produced measure of condition, computed by a defined method under defined capture conditions, bound to the evidence it was computed from, and checkable after the fact by anyone who cares to look. This document specifies that number and the record that carries it.

2. What Existing Standards Do — and the Gap They Leave

The secondary-electronics industry is not unregulated territory; it is served by several serious standards. Each does its job. None of them produces an objective measure of a device's cosmetic and structural condition, and it is worth being precise about that, because this specification is designed to sit beside them, not to displace them.

- CTIA Wireless Device Grading, Version 5.1 (June 2026) [1]. A constructive and recent step: a voluntary common vocabulary that participating sellers may apply to wireless devices. It standardizes the words used to describe a grade. It does not standardize the method of producing one — a human grader still exercises subjective judgment to decide what “Good” means — and it is limited to phones and tablets. This specification treats the CTIA lexicon as an interoperability target (Section 8): the Wear Index is designed to supply the objective method beneath vocabularies like CTIA's, not to compete with the vocabulary itself.
- R2v3 (SERI) [2] and e-Stewards (Basel Action Network) [3]. Facility certifications governing environmental handling, data security practice, and downstream accountability in electronics

recycling and reuse. They certify how a facility behaves; they do not define a machine-produced measure of an individual device's condition.

- NIST Special Publication 800-88, Revision 1 [4]. The recognized standard for media sanitization. It answers the data-security question definitively — and this specification's certificate format references it directly (Section 6) — but it says nothing about physical condition.
- Marketplace grade scales. The de facto standards buyers actually encounter are the incompatible letter-and-adjective scales of individual platforms, each applied by eye, none defined by published method.

The gap, stated exactly: no existing standard defines an objective, machine-produced, auditable measure of a used device's cosmetic and structural condition that holds across sellers. That measure is the missing layer of trust infrastructure in a national-scale market, and it is the only layer this specification attempts to define.

3. Design Principles

Six principles govern every normative choice in this document. They are stated here so that public comment can challenge the foundations, not only the details.

1. Machine-produced objectivity. The grade is computed, not opined. Under this specification, a conformant grade is the output of a defined imaging procedure and a versioned classification model — never of unassisted human judgment. Humans retain disposition authority over machine-raised anomaly flags, and their disposition is recorded, but the score itself is produced by method.
2. Evidence-bound auditability. A score that cannot be checked is an opinion with extra steps. Every conformant grade binds, by content-addressed reference, to the exact images it was computed from, the model version that computed it, and the specification version that governed it — so that any party can later re-examine the evidence and reproduce the assessment.
3. Openness. The specification is published openly, versioned publicly, and adoptable by any party without permission or fee. A grading standard creates value only when others adopt it; this one is designed from the outset to be adopted rather than owned.
4. Interoperability, not replacement. The market already speaks in vocabularies — CTIA's lexicon, marketplace scales, institutional procurement language. The Wear Index maps to them (Sections 5.2 and 8) instead of asking anyone to abandon them. The score supplies the method beneath the words.
5. Proportional adoption. A small independent refurbisher and a national processor should both be able to participate honestly. Three conformance levels (Section 7) define appropriate depth for each, so that the standard's reach is set by the market's, not by the capacity of its largest members.
6. Honest status. The specification and its implementations must state plainly what is validated and what is not. Version 0.1 makes no performance claims; Section 9 defines what an implementation must publish before calling itself validated; Section 10 states current limitations without decoration. Versioning makes every claim checkable against the document that governed it.

4. Reference Architecture

The specification assumes a five-component grading station. The bracketed element numbers [100]–[140] follow the numbering used in the author's pending U.S. provisional patent application, to keep the public specification and the underlying filing cross-referenceable; the architecture is presented here as design intent, at the level of what each component is for, not as a report of an operating system. Figure 1 shows the components and how a device's record moves through them.

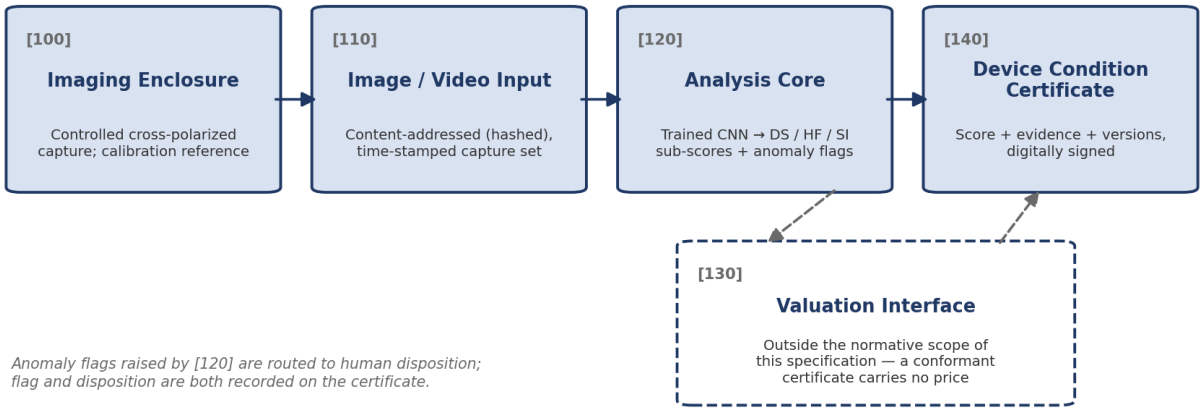

Figure 1. Reference architecture, concept level. Solid elements are within the normative scope of this specification; the valuation interface [130] is shown dashed because it is deliberately outside it (Section 4.4).

4.1 Imaging Enclosure [100]

A controlled lighting-and-camera enclosure in which the device is placed for capture. Its defining feature is cross-polarized illumination: polarizing filters on the light sources and on the camera, oriented to suppress the specular glare that ordinarily bounces off glossy glass and painted housings. Removing that glare is what exposes the fine surface detail — hairline cracks, early delamination indicators, subtle deformation — that defeats unaided visual inspection. The enclosure holds lighting geometry, background, and camera position constant, and includes a calibration reference in frame, so that captures are comparable across devices, operators, and time. A capture profile, defined per device category, specifies the required views, lighting states, and minimum resolution.

4.2 Image and Video Input [110]

High-resolution still images, and video where the capture profile requires it, taken under the enclosure's controlled conditions, form the complete evidentiary input to scoring. Nothing outside the capture set may influence a conformant score. Every item in the set is content-addressed (hashed) and time-stamped at capture, which is what later makes the certificate auditable (Section 6).

4.3 Analysis Core [120]

A trained convolutional neural network examines the capture set and is designed to detect and localize the damage classes that matter to condition: scratches, cracks and chips, display-delamination indicators, dents, discoloration, and deformation patterns — including the fine fractures under glare-suppressed

illumination that are easy to miss by eye — and to convert those detections into the sub-scores defined in Section 5. The analysis core may additionally raise anomaly flags: indications that warrant human attention, such as a suspected non-original component or a device-status indicator requiring verification. Flags are not grades. Every flag is routed to a human for disposition, and both the flag and the disposition are recorded on the certificate. The mechanisms by which an implementation verifies device status or component provenance are implementation choices outside the normative scope of this specification.

4.4 Valuation Interface [130]

In the author's full architecture, the condition output feeds a valuation engine that prices the device from current market data. Valuation is intentionally outside the normative scope of this specification. Pricing is commercial, market-dependent, and implementation-specific; standardizing it would neither serve adopters nor be honest about how markets work. A conformant Device Condition Certificate carries no price. The specification defines the condition record; what an adopter does with it commercially is the adopter's business.

4.5 Output: the Device Condition Certificate [140]

The terminal artifact is not a label but a certificate: a structured record binding the Wear Index score, its sub-scores, the evidence set, model and specification versions, and — where the operator also performs data sanitization — a reference to the device's NIST SP 800-88 sanitization certificate. Section 6 defines the certificate. The design intent is a single portable record (a “device health passport”) that answers the buyer's real question: not “what word did the seller choose,” but “what exactly was observed, by what method, and can I check?”

5. The Wear Index: Score Model

5.1 Composite and Sub-Scores

The Wear Index (WI) is a composite score on a 0–100 scale, where 100 denotes no machine-detectable cosmetic or structural wear under the defined capture conditions. It is computed from three sub-scores, each also expressed 0–100:

- DS — Display Surface. Front-glass and display-stack condition: scratches, cracks, chips, delamination indicators, pressure marks.
- HF — Housing and Frame. Rear and edge surfaces: scratches, dents, chips, discoloration, wear-through.
- SI — Structural Integrity. Deformation, bend, gap, and separation indicators, and crack patterns whose geometry suggests structural rather than superficial damage.

The composite is a weighted combination of the sub-scores. Weights are normative per specification version and per device category, and are published with the version; an implementation must state the specification version in every certificate, which fixes the weights used. Version 0.1 proposes, for the smartphone/tablet profile, a weighting that privileges the display surface as the dominant driver of buyer-perceived condition, with structural indicators acting as a cap: a device whose SI sub-score falls below the

certifiability floor cannot receive a certifiable composite regardless of its other sub-scores. Exact weight values are a matter for public comment in this draft. Figure 2 summarizes the model.

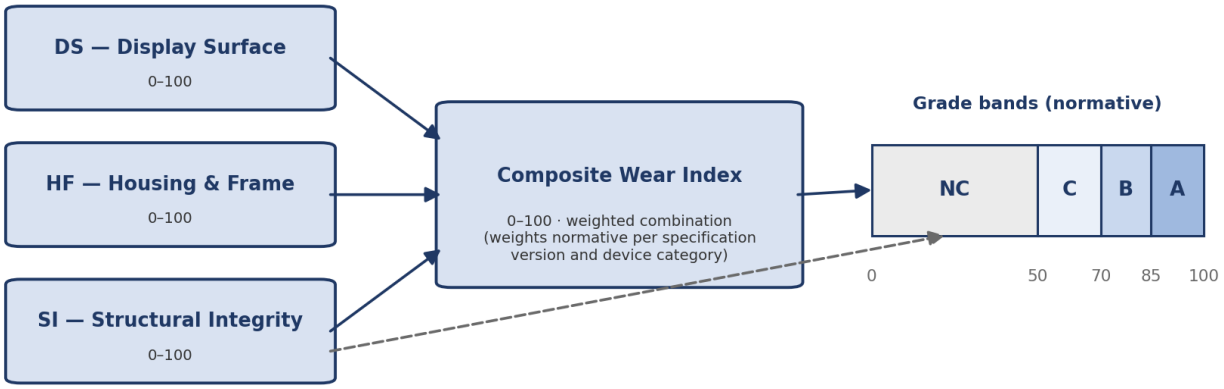

Figure 2. Score model: three sub-scores combine into the composite Wear Index, which maps to normative grade bands; a Structural Integrity sub-score below the certifiability floor caps the device at NC regardless of composite.

5.2 Grade Bands and Vocabulary Mapping

Because the market communicates in letter grades and adjectives, the specification defines fixed bands over the composite score and a mapping layer to external vocabularies. The bands are normative; the vocabulary columns are informative alignments that adopters may adapt to their own customer-facing language without altering the underlying score.

Band	Composite Wear Index	Informative vocabulary alignment
A	85–100	“Excellent” / “Premium”-tier terms in market scales.
B	70–84	“Good”-tier terms.
C	50–69	“Fair” / “Acceptable”-tier terms.
NC	Below 50, or SI below the certifiability floor	Not certifiable for consumer resale under this specification; route to repair, parts recovery, or certified recycling.

An informative annex to the specification repository maintains the mapping from these bands to the CTIA Wireless Device Grading vocabulary [1] and to representative marketplace scales, and will track those vocabularies as they evolve. The mapping direction matters: external words are described in terms of the Wear Index, never the reverse, because the score is the defined quantity and the words are its presentation.

6. Certificate and Audit-Trail Format

Every graded device receives a Device Condition Certificate. The certificate is the unit of trust in this specification: scores without certificates are not conformant output. Version 0.1 defines the following required fields; the machine-readable schema (JSON) accompanies the specification repository.

Field	Description
spec_version	Wear Index specification version governing bands, weights, and capture profile (e.g., 0.1).
model_version	Identifier and cryptographic hash of the trained vision model that produced the scores.
capture_profile	Device-category capture profile applied (views, lighting states, resolution).
device_identity	Make, model, storage variant; device serial identifiers stored as salted hashes to permit verification without exposing the identifier.
capture_set	Content-addressed references (hashes) to every image and video used in scoring, with per-item capture timestamps and enclosure calibration reference.
sub_scores	DS, HF, SI values (0–100).
composite_wi	Composite Wear Index (0–100) and resulting grade band.
anomaly_flags	Machine-raised flags (e.g., suspected non-original component, status-check indicator) with disposition field for human review outcome.
sanitization_ref	Reference to the device's NIST SP 800-88 sanitization certificate, where the operator performs sanitization.
functional_ref	Optional reference to an attached functional-test report (out of scope for scoring; in scope for the record).
operator	Facility and operator identifiers of the grading run.
issued_at / signature	Issuance timestamp and the operator's digital signature over the complete record.

The audit property follows from construction: because the certificate binds content-addressed evidence to a versioned model and a versioned specification, any party can later re-examine the images, re-run a conformant model, and compare outputs. Disputes reduce to evidence, which is the entire point.

7. Conformance Levels and Adoption Model

Adoption is the measure of a standard's success, and adopters differ in capacity. Version 0.1 defines three conformance levels so that a small refurbisher, a large processor, and a marketplace can each participate at appropriate depth.

Level	Requirements	Intended adopter
WI-Basic	Controlled capture per profile; publishes composite WI and grade band with spec and model versions.	Independent refurbishers entering objective grading.
WI-Audit	WI-Basic plus full Device Condition Certificate issuance and evidence retention for a defined period.	Refurbishers and processors selling into institutional and wholesale channels.

Level	Requirements	Intended adopter
WI-Interop	WI-Audit plus machine-readable certificate exchange and vocabulary mapping via the published schema.	Marketplaces, carriers, and large processors integrating grades across many sellers.

The levels are cumulative. An adopter states its conformance level in every certificate and listing that carries a Wear Index, so that the depth of the claim travels with the claim itself.

8. Interoperability Posture

This specification is deliberately a layer, not an island. Its relationships to the existing landscape are as follows.

- To the CTIA Wireless Device Grading lexicon [1]: the informative mapping in Section 5.2 lets a seller who grades by Wear Index present results in CTIA vocabulary. The lexicon supplies shared words; the Wear Index supplies the objective method beneath them. The two are complements by design.
- To marketplace scales: the same mapping layer lets sellers keep their customer-facing adjectives while the underlying score — and its evidence — remains portable across platforms. A certificate can travel with a listing without requiring the platform to change anything.
- To R2v3 [2] and e-Stewards [3]: facilities certified under either can issue Wear Index certificates within their existing practice; the specification adds a condition-measurement layer to operations those standards already govern, and conflicts with neither.
- To NIST SP 800-88 [4]: the certificate's sanitization_ref field links the condition record to the sanitization record, so a single document chain answers both of the buyer's trust questions — what condition, and whose data.
- Category extensibility: Version 0.1 defines a capture profile for smartphones and tablets. The score model and certificate format are category-neutral by construction; later versions will add profiles (laptops, wearables, and other consumer categories) without altering the architecture.

9. Validation Methodology

A specification for objective grading must define what it would mean to demonstrate that an implementation actually grades objectively. This section defines that demonstration. Everything in it is a design requirement for future validation, not a report of results obtained.

9.1 Training and Evaluation Corpus

A conformant implementation is expected to train and evaluate on a labeled image corpus of 1,500–3,000 images per device-category capture profile, spanning all grade bands and the damage classes enumerated in Section 5.1, captured under the profile's controlled conditions, with labels produced under a documented labeling protocol. The corpus description — size, composition, labeling protocol, and inter-labeler agreement — must be published with the validation study.

9.2 Blinded Comparison Study

The core validation instrument is a blinded study of at least 200 physical devices, sampled across grade bands and damage types. Each device is graded independently by (a) the implementation under test, operating under this specification, and (b) two independent, qualified human graders applying the market's conventional visual method, blinded to the machine output and to each other. The study reports weighted Cohen's kappa (quadratic weights) for model-versus-grader-consensus agreement, grader-versus-grader agreement as the human baseline, confusion matrices over the grade bands, and per-sub-score error analysis.

9.3 Design Target and Publication Requirement

The design target for Version 0.1 is model-to-consensus agreement of weighted $\kappa \geq 0.75$, meeting or exceeding the observed grader-to-grader agreement — the point at which a machine grade is demonstrably at least as consistent as the human process it replaces. This is a target, not a claim. An implementation conforms to this section only when it publishes the complete study: corpus description, protocol, statistics, and failure analysis. Until that publication, an implementation must describe itself as unvalidated under this specification. This rule applies to every implementation, including the author's own.

9.4 Re-validation

Because the certificate binds scores to a model_version, any change of model requires either re-validation under Section 9.2 or an explicit statement that certificates issued under the new version are unvalidated. Implementations are further expected to monitor for drift between validation episodes and to disclose their monitoring practice in the published study.

10. Limitations and Honest Status

A specification earns trust by stating what it does not yet do. The following limitations are asserted plainly.

- No validated implementation of this specification exists as of this draft. The author's own implementation is under development and, per Section 9.3, must not be described as validated until its study is published.
- Scope is cosmetic and structural condition only. Functional testing — whether the device works — is out of scope for the score; the certificate may reference an attached functional report (functional_ref), but the Wear Index does not measure function.
- Scores are defined relative to controlled capture conditions. A Wear Index computed from uncontrolled photography is not a Wear Index; it is out of scope and non-conformant.
- Version 0.1 profiles cover smartphones and tablets only. Other categories await later profile versions.
- Vision models have known failure modes: novel damage types, unusual materials and finishes, and deliberate concealment can defeat classification. The anomaly-flag and human-disposition

mechanism exists precisely because machine grading is strong but not omniscient; the specification treats human review as a recorded part of the method, not an embarrassment to it.

- The composite weights proposed in Section 5.1 are provisional pending public comment; certificates issued against Version 0.1 must be read accordingly.

11. Roadmap and Governance

Version 0.2 will integrate public comment, publish normative composite weights for the smartphone/tablet profile, and release the machine-readable certificate schema in the specification repository. Subsequent milestones, in intended order: publication of a reference image corpus description; publication of the first complete validation study under Section 9; and additional device-category capture profiles.

On governance: this document begins as a single-author draft because someone has to write the first version. It is not intended to remain one. As adoption warrants, the author's intent is to move stewardship of the specification toward a multi-stakeholder group — refurbishers, marketplaces, recyclers, and institutional buyers — with versioning authority held by the group rather than by any single commercial participant. An open standard that stays the property of one company is a product with better marketing; the roadmap is designed to prevent that outcome.

Comments on this draft are invited from all parties, and particularly from working refurbishers and graders, marketplace trust-and-safety teams, recyclers, and institutional device buyers. Direct comments to the correspondence address on the title page.

12. Intellectual-Property Notices and Licensing

Specification license. The text of this specification is published under the Creative Commons Attribution 4.0 International license (CC BY 4.0). Any party may implement, cite, reproduce, and redistribute it, with attribution, without permission or fee.

Relationship to the author's patent application. The term “Wear Index” and the imaging-and-classification method on which this specification draws originate in a pending United States provisional patent application filed by the author on May 27, 2026. Provisional applications are not published by the United States Patent and Trademark Office; the application is unpublished and on file with the author, and consistent with that confidentiality its serial number and title are not reproduced here. This document and that application are deliberately distinct in scope: this specification defines an open score model, certificate format, conformance levels, and validation requirements — the layer intended for adoption by the field — while particular implementation methods described in the patent application remain the intellectual property of the inventor. Publication of this specification is not a license, express or implied, to practice patented or patent-pending implementation methods.

No warranty. This is a draft specification provided as-is, for public comment, without warranty of any kind. It reports no performance results and makes no fitness claims; Sections 9 and 10 govern what any implementation may claim and when.

References

[1] CTIA, Wireless Device Grading, Version 5.1 (June 10, 2026).

[2] SERI (Sustainable Electronics Recycling International), R2v3 Standard. https://sustainableelectronics.org/r2/

[3] Basel Action Network, e-Stewards Standard. https://e-stewards.org/

[4] National Institute of Standards and Technology, Special Publication 800-88, Revision 1: Guidelines for Media Sanitization (2014). https://csrc.nist.gov/pubs/sp/800/88/r1/final

[5] United Nations (ITU & UNITAR), The Global E-waste Monitor 2024. https://ewastemonitor.info/the-global-e-waste-monitor-2024/

[6] U.S. Government Accountability Office, Science & Tech Spotlight: Consumer Electronics Recycling, GAO-20-712sp (2020). https://www.gao.gov/products/gao-20-712sp

[7] Mordor Intelligence, US Refurbished and Used Mobile Phones Market (industry estimate). https://www.mordorintelligence.com/industry-reports/us-refurbished-and-used-mobile-phones-market

[8] Assurant, “Applying AI and Automation Should Be Boring” — industry publication describing AI-assisted grading applied at the scale of roughly 20 million devices per year. https://www.assurant.com/news-insight/insights/mobile-connected-world/article/applying-ai-and-automation-should-be-boring

[9] Iftikhar, K., pending U.S. provisional patent application on automated condition assessment of consumer electronics using polarized-light computer vision (filed May 27, 2026; unpublished, on file with the author).