



EVIDENCE COVERAGE AND UNRESOLVED CLAIMS IN DIGITAL FILE INTEGRITY ASSESSMENT

B. B. Turdibekov

Tashkent University of Information Technologies named after Muhammad al-Khwarizmi, Uzbekistan

Abstract. Integrity assessment can be incomplete even when every executed check returns a plausible result. This paper models the coverage of a requested claim using applicable evidence, unavailable inputs and contradictory findings. A three valued support relation distinguishes supported, contradicted and unresolved requirements. A coverage matrix connects authenticated digest comparison, provenance and file-specific observations to distinct questions. The analysis derives a stopping condition for a bounded claim and demonstrates why missing reference data cannot be compensated by additional visual indicators. The resulting framework specifies both the conclusion and the precise evidence still required for review.

Keywords: digital file integrity, metadata, structural analysis, visual forensics, trusted baseline.

Introduction


Digital files can change through deliberate editing, routine conversion, recompression, or storage errors. The observable change alone does not establish whether an action was authorized. A hash comparison can determine byte agreement with a trusted reference [1], while forensic examination considers additional evidence about the file and the sequence of actions affecting it [2]. These questions require different forms of evidence.

This distinction matters for mixed document collections. A PDF can contain text, images, object references, and metadata governed by a structured format [3]. A JPEG image has compression characteristics that affect the traces visible after processing [4]. Visual inspection may reveal a suspicious region, but a consistent appearance cannot establish exact byte equality.

The objective of this paper is to compare evidence layers and identify how they complement each other in file integrity assessment. The framework links claim coverage to method applicability and explicit unresolved evidence.

A qualitative comparison was performed using the Secure Hash Standard, forensic guidance, document format specifications, and image analysis literature. The selected sources describe mechanisms relevant to the question of what each check observes. The comparison used four criteria: required baseline, directly observable property, interpretation supported by that property, and main ambiguity.

Five evidence layers were considered: cryptographic hash, metadata, internal file structure, visual or pixel features, and compression traces. The analysis followed two





MODERN PROBLEMS IN EDUCATION AND THEIR SCIENTIFIC SOLUTIONS

scenarios. In the first, an authenticated baseline was available. In the second, only the questioned file was available. The resulting mapping is summarized in Table 1.

Let E denote the existence of an authenticated baseline digest, and let $H(F)$ be the SHA-256 digest of the examined bytes. The comparison result is a partial function: a missing baseline is not assigned a numerical zero or interpreted as a match [1]. This avoids conflating an absent check with a negative observation.

$$B(F) = \begin{aligned} &\text{match if } E \wedge H(F) = H_0; \text{ difference if} \\ &E \wedge H(F) \neq H_0; \text{ undefined if } \neg E. \end{aligned} \quad (1)$$

For contextual evidence, define $G = (V, A)$, where each node in V is an observation and an edge in A records a plausible shared processing cause. G is a reasoning aid, not a fitted probabilistic model. A claim of independent corroboration requires evidence beyond nodes that may all arise from the same export event. The provenance relation between entity, activity and agent in PROV-DM offers a way to document that common cause [5].


Table 1.

Evidence layers and interpretation boundaries

Layer	Required input	Supported observation	Interpretation boundary
Hash	File and trusted reference digest	Byte agreement or disagreement	Does not identify intent or edited region
Metadata	Service attributes and contextual expectations	Inconsistent times, software or provenance fields	Fields can be changed or removed
Structure	Parsed file and format rules	Object or container anomalies	Legitimate export can alter structure
Visual features	Image or rendered page	Local visual inconsistencies	Sensitive to acquisition and image quality
Compression traces	Encoded image and analysis settings	Differences related to recompression	Benign processing may produce similar traces

With a trusted baseline, the hash layer supplies a direct equality test. The remaining layers can explain what changed or where additional examination should be concentrated. When the baseline is absent, structural and visual findings can indicate inconsistency, but they do not reconstruct an original file by themselves.

Metadata is useful for contextual comparison. For example, a software field or editing timestamp can conflict with the expected workflow. Such a conflict supports further examination, while a plausible field value is not sufficient evidence of authenticity. File structure supplies another layer because an unchanged visible page may still contain different objects or service information.





MODERN PROBLEMS IN EDUCATION AND THEIR SCIENTIFIC SOLUTIONS

ELA and other compression based checks require careful interpretation. The analysis of image error levels can highlight regions with different processing histories, yet differences may also result from ordinary saving operations [6]. Broader media forensics literature similarly treats detection as dependent on manipulation type and available evidence [7]. Combining the layers therefore supports explanation and triage, provided their limitations remain visible.

The comparison yields a reporting rule that separates three outcomes: a demonstrated match to an authenticated version, a demonstrated byte difference from that version, and an assessment based only on indicators because no authenticated version exists. For each indicator, the report should name the observed feature, the applicable method, a plausible benign cause, and the next check that could distinguish competing explanations. For instance, reordered PDF objects may be consistent with a routine export; an authenticated version history can test that explanation. This makes the path from observation to conclusion inspectable.





Figure 1. Evidence layers used to interpret file integrity

Let Q be the set of requirements needed for a requested claim and E the observations available to an examiner. Define $S(e, q)$ as supported, contradicted or unresolved. These values describe the relation between a specific observation and a specific requirement. A plausible producer field cannot support the requirement of byte equality; its relation to that requirement is unresolved even when the field contains a familiar software name.

A coverage record stores, for each q , the evidence that supports it, any contradiction and the prerequisite still missing. The claim is accepted only if every required q has a valid supporting path and none has an unresolved contradiction within the declared scope. This creates a conjunction over requirements rather than a vote over tools.

$$Accept(Q, E) \Leftrightarrow \forall q \in Q: Supported(q, E) \wedge \neg Conflict(q, E) \quad (2)$$

The equation depends on the meaning of Supported. For equality to an approved byte sequence, support requires an authenticated reference and a matching digest. For



**MODERN PROBLEMS IN EDUCATION AND THEIR SCIENTIFIC SOLUTIONS**

authorization, support requires a relevant policy and approval record. For a local alteration hypothesis, support may identify a region and a plausible transformation, while the exact original remains unresolved. Such a finding is reported at the hypothesis level instead of satisfying an unrelated equality requirement.

Table 2.

Coverage states for a requested assessment

Required property	Available evidence	Coverage state	Missing input
Approved byte version	Current digest only	Unresolved	Authenticated reference
Recorded byte version	Matching trusted digest	Supported	None for this scope
Authorized revision	Digest mismatch only	Unresolved	Approval and version link
Common processing cause	Metadata and recompression traces	Hypothesis	Independent process record

Consider two examinations with three executed tools each. Examination A has metadata extraction, image inspection and recompression analysis but no authenticated earlier copy. Examination B has format identification, a trusted digest comparison and an approval record. A raw count assigns both the same coverage. The requirements differ: A may identify processing clues, while B can support correspondence to a recorded version and its approval context. The number of tools is therefore not a measure of claim completeness.

The reverse problem is also possible. A matching digest can settle the narrow equality question while leaving provenance unresolved. Running more visual tools on the same bytes adds observations but does not establish who approved the registered reference. The next action should target the uncovered requirement. This gives an adaptive rule: obtain the evidence source needed for the unresolved claim before repeating checks that already address covered properties.

An apparent contradiction must be located at the same scope before it is escalated. For example, a whole-file digest mismatch and a valid signature over an earlier signed byte range can coexist in a PDF with later additions. The observations refer to different extents. The examiner records the signed range and the appended content, then asks whether those additions were permitted [3]. Labeling the two observations as a direct conflict would discard their scope information.

A justified stopping point is reached when the requested claim is covered, or when the missing prerequisite cannot be obtained and its effect has been stated. The second outcome is an unresolved result with an explicit next evidence request. The framework thus supports reproducible triage: another examiner can see not only which tools ran but why the process





MODERN PROBLEMS IN EDUCATION AND THEIR SCIENTIFIC SOLUTIONS

ended. This is particularly useful when a file passes between technical analysts and an authority responsible for approval decisions.

The baseline scenario gives a useful way to organize the examination. Suppose an organization has retained a trusted digest of an approved document. A newly submitted file can be hashed and compared with that digest. Equality answers the limited question of byte correspondence to the approved version. If the digests differ, the examiner can then use metadata and structure to locate an explanation. The mismatch by itself says nothing about whether the change was permitted. A routine conversion and an unauthorized edit may both produce a new digest.

The second scenario begins without a trusted earlier state. A file still has format rules and observable features, so the examiner can ask whether its internal structure is plausible and whether its metadata agrees with the surrounding workflow. An image region with unusual compression behavior may be selected for closer inspection. These observations can prioritize investigation, yet their absence is not a proof that no manipulation occurred. The proposed framework keeps the word “indicator” tied to such observations and reserves “match” for a comparison against a known reference.

An important practical question is whether apparent corroboration really comes from different causes. A PDF export can simultaneously update its creation software field, rebuild object offsets, and recompress embedded images. Those three observations are related consequences of one operation. Treating them as independent votes would overstate the strength of the evidence. A report should record the originating module, the relevant file region, and a plausible common process that may have produced the observations.

The choice of checks also changes with the file format. PDF parsing can reveal internal objects that a page preview does not display. In a JPG scan, compression and pixel features are more directly relevant, while document metadata may provide weak contextual support. A PNG image should not be assessed as though it had JPEG compression artifacts. Selecting only applicable checks avoids creating a spurious result when a method has no meaningful input for that format.

The resulting selection rule has two steps. First, establish whether an authenticated reference exists. Second, select checks that apply to the file representation and the question under examination. A visual heuristic then retains the role of an indicator, while reference based comparison answers a narrower equality question.

Coverage depends on the available reference and the recognized representation of the file. An unsupported format or missing approval history leaves the associated requirement unresolved and determines the next evidence request.

The assessment can be documented as a short evidence record. Such a record identifies the examined file, the format recognized from its contents, the existence or absence of a trusted reference, the checks that were applicable, and the observed outcomes. It should preserve the difference between a failed check and a check that could not be performed. For example, an unavailable reference digest means hash correspondence is unknown; it





MODERN PROBLEMS IN EDUCATION AND THEIR SCIENTIFIC SOLUTIONS

does not mean that correspondence has been established. Explicitly recording this distinction makes later review possible.

The same principle applies to the interpretation of a suspicious region. A highlighted image area may guide an analyst to compare nearby pixels, compression history, and document context. It does not, on its own, identify an editor or a time of editing. A separate source document, transaction history, or authenticated earlier version may be needed to answer those questions. Describing the limits of each observation improves the practical usefulness of a combined report.

Conclusion

There is also an ordering benefit in the proposed framework. Cheap identification and format checks can be performed first, then computationally or manually demanding inspection can be concentrated on files whose initial observations merit it. This ordering is a workflow recommendation derived from the distinct roles of the evidence layers. The ordering remains conditional on the cost and relevance of each check. A trusted baseline determines the interpretation available from hash comparison. Where no baseline exists, metadata, structural, and image analysis supply qualified indicators. The proposed analytical mapping supports complementary selection of checks and transparent reporting of their evidential boundaries.

References

- [1] National Institute of Standards and Technology, Secure Hash Standard (SHS), FIPS PUB 180-4, 2015. DOI: 10.6028/NIST.FIPS.180-4.
- [2] K. Kent, S. Chevalier, T. Grance and H. Dang, Guide to Integrating Forensic Techniques into Incident Response, NIST SP 800-86, 2006. DOI: 10.6028/NIST.SP.800-86.
- [3] ISO 32000-2:2020, Document management – Portable document format – Part 2: PDF 2.0. Geneva: ISO, 2020.
- [4] ITU-T, Information technology – Digital compression and coding of continuous-tone still images – Requirements and guidelines, Recommendation T.81, 1992.
- [5] W3C, PROV-DM: The PROV Data Model, W3C Recommendation, April 30, 2013. <https://www.w3.org/TR/prov-dm/>.
- [6] N. Krawetz, A Picture's Worth: Digital Image Analysis and Forensics. Hacker Factor Solutions, 2007.
- [7] L. Verdoliva, "Media Forensics and DeepFakes: An Overview," IEEE Journal of Selected Topics in Signal Processing, vol. 14, no. 5, pp. 910–932, 2020. DOI: 10.1109/JSTSP.2020.3002101.

