

Fake News Hates Structure

An authentic video can carry a false caption. Separating the asset, the claim, its evidence and correction history can make that unsupported step easier to challenge.

2026.10.20 21:40 DENNIS HEDEGREEN ANALYSE V1.0

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/FAKE-NEWS-HATES-STRUCTURE/](https://hedegreenresearch.com/articles/fake-news-hates-structure/)

Structure does not make journalism true. It can make truth claims easier to inspect.

The Glacier Was Real. The Caption Wasn't.

In late August 2026, a helicopter video began circulating with an explanation of the deadly flooding near the Nepal–Tibet border. Posts presented it as footage of the glacier responsible for the disaster. Some even described a Chinese helicopter surveying the site.

The glacier in the video was real. The attribution was false.

The footage showed Columbia Glacier in Alaska, not the site of the 26 August flood. Pilot Leigh Coates told AFP that she had recorded it in 2024 and reposted it on 21 August 2026—five days before the disaster. AFP identified the earlier footage and reported her account. Full Fact later documented the same false-context claim, drawing on earlier reporting. 1 (<https://fullfact.org/world/old-video-alaska-glacier-shared-as-nepal-flood-cause/>). 2 (<https://factcheck.afp.com/doc.afp.com.C6Q96G3>).

No synthetic landscape was needed. No convincing deepfake was required. An authentic recording had been given a false job: to support a story about a different place and event.

That is the problem this article is about. **The authenticity of an object and the truth of the claim attached to it are different questions.**

An organized liar can use structured data too. The question is whether an outsider can inspect the step from evidence to claim.

The False Step Can Be a Sentence

Imagine a caption beneath the clip: *This is the glacier that caused the Nepal flood.*

That sentence contains several testable assertions: the video's location, the time of the event shown, the identity of the glacier and its causal role in the flood. Evidence for one does not automatically establish the others.

A genuine image may contain clues about its setting; verification work can compare them with other images and records. But authenticity alone does not establish the publisher's account of where, when or why it was filmed. Even a reliable capture date would not prove that this particular glacier caused a disaster thousands of kilometres away.

Yet when a video, headline, caption and source link are presented together, readers may receive them as a single factual package. If the caption is corrected later, an older screenshot may still carry the error. Those are distinct problems: first, the unsupported inference; then, the failure of the correction to travel.

An article can explain all of this. But when the record stops at the page, it can be difficult to identify *which precise proposition* was disputed, what evidence was considered, and what subsequently changed.

The unit of accountability should not always be the whole page.

Give the Claim an Address

Suppose a newsroom published the misleading glacier caption and then examined a challenge to it. An inspectable record might look like this:

- **Published claim:** “This video shows the glacier responsible for the August 2026 flood near Nepal.”
- **Reference:** A stable public claim identifier assigned by the publisher.
- **Attached material:** The video, treated separately from assertions about what it depicts.
- **Supporting evidence at publication:** Listed, if any; absence of support stated rather than disguised.
- **Contrary evidence:** Earlier Alaska publication and the independently identified location of the footage.
- **Assessment:** Attribution contradicted by the evidence reviewed, with a date and responsible editor.
- **History:** Original wording, challenge, decision and correction, with dates.

This is a **proposed example**, not a record that Full Fact, AFP or the accounts posting the video are reported to have used.

Its advantage is not a magic truth label. It is precision. A finding that *the video was filmed in Alaska* refutes the video's attribution to Nepal. It does not, by itself, disprove the underlying Nepal flood or every claim in a related story.

The record also leaves room for uncertainty. An editor might find a claim unsupported, disputed, partly established or contradicted. Such states would be editorial judgments with evidence and dates attached, not values calculated by the shape of a database.

And a claim address should not expose a confidential source. Where details must be protected, the public may be able to inspect the verification method, the category of evidence and the reason for withholding particulars rather than the underlying identity.

The aim is a challenge that can reach an assertion without destroying the story around it.

The Standards Already Answer Different Questions

We are not starting from zero. Several existing systems make parts of journalism more inspectable, but they verify—or disclose—different things.

ClaimReview (<https://schema.org/ClaimReview>), a schema.org type, gives fact-checkers structured fields for the claim reviewed and the reviewer's assessment. Its markup describes a review. It does not independently establish that the verdict or its cited evidence is correct. ³ (<https://schema.org/ClaimReview>).

Google's handling of ClaimReview also exposes a dependency problem. In June 2025, Google announced it would phase out Claim Review rich-result displays in general Search. A September update said the feature had been phased out and would be removed from Search Console reporting and testing tools. Google still documents support for the markup in Fact Check Explorer. The underlying schema survived a change in its visibility on a major platform. ⁴ (<https://developers.google.com/search/blog/2025/06/simplifying-search-results>), ⁵ (<https://developers.google.com/search/docs/appearance/structured-data/factcheck>).

C2PA (https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html) addresses the provenance of a digital asset. Its signed manifests can make declared creation or editing information tamper-evident and help establish who or what signed it. The specification explicitly avoids judging whether provenance assertions are 'good' or 'bad'. A valid signature does not independently verify the reality of a publisher's explanation of the scene. ⁶

(https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html).

The Alaska footage is **not** presented as a C2PA-credentialed case. The cited fact-checks do not establish a C2PA credential history for that particular clip. It is an example of a distinction that would remain relevant *even if* its media provenance were authenticated.

Nor does the absence of a credential make a file false. Credentials may never have been attached, may be deliberately omitted or may be stripped during redistribution. C2PA discusses both removal of manifest data and ways to limit disclosure of identifying information. ⁷

(https://spec.c2pa.org/specifications/specifications/2.4/security/Security_Considerations.html).

These tools can provide useful evidence about an asset or a published review. None eliminates the separate need to check the claim made *about* the asset.

A Process Badge Cannot Repair a Caption

The Journalism Trust Initiative assesses newsroom processes such as transparency, independence and professional accountability; Reporters Without Borders explicitly says its criteria do not evaluate individual items of content. The International Fact-Checking Network's Code of Principles concerns methods, sourcing, fairness and corrections. The Trust Project offers eight indicators that help people assess journalistic practices and context. ⁸ (<https://rsf.org/en/journalism-trust-initiative-tout-comprendre-%C3%A0-la-certification-des-m%C3%A9dias-d-information>), ⁹ (<https://ifcncodeofprinciples.poynter.org/the-commitments>), ¹⁰ (<https://thetrustproject.org/trust-indicators/>).

Such work matters. A newsroom should disclose how it operates. But a transparent institution can still produce a wrong sentence, and a diligent fact-checker can still misjudge evidence.

Organizational trust deserves its own examination. Here the test is narrower: **when a specific claim fails, can the reader see where it failed?**

The Liar Can Fill In the Boxes Too

There is an obvious and serious objection. An organized liar can use structured data.

A fabricated story can have a permanent claim identifier, well-formed markup, a long reference list, neatly dated revisions and links to ten websites repeating the same invention. A signed file can be miscaptioned. A corrections page can be decorative.

If the proposal is merely “add metadata,” it fails this test.

The meaningful standard is whether an *outside challenger* can get past the formatting to the evidence. Is there a dated original? Does a cited source actually say what the story claims? Are several apparently independent references all quoting the same unsupported assertion? Can the critic submit contradictory material and discover what the publisher did with it?

Even this does not guarantee truth. Some sources are fabricated. Others are unavailable, dangerous to disclose, prohibitively expensive to access or controlled by the people making the claim. A publisher may refuse a valid challenge. Readers may never see a correction.

Structure is therefore not a substitute for independent reporting, judgment or power to investigate. It can preserve the questions those activities must answer. The claim that it reduces the *quantity* of misinformation, or improves public trust, would require comparative evidence that this article does not present.

There is another cost: maintenance. Small publishers cannot create a public evidence dossier for every clause. Start with claims whose failure would materially change a story: a number, a date, an identity, an attributed quotation, a location, a decisive causal inference. Keep the narrative readable. Make the vulnerable joints inspectable.

A Correction Needs a Destination

Imagine the outlet that posted the glacier clip quietly replacing the misleading caption.

A new visitor sees the corrected text. Someone with a screenshot sees the original. A summarizer relying on an old copy may repeat it. A second publisher may never find out that the attribution changed.

These are possible downstream failures, not outcomes measured for this particular case. No system can reliably retrieve every screenshot or compel another publisher to correct a copy.

But the original publisher can maintain a stable reference to the disputed claim, a dated account of what changed and a link from the old assertion to its correction. A downstream system *that chooses to check* can use that record. A reader can learn not merely that the current article looks different, but why.

This is also a continuing editorial obligation. An identifier with stale evidence is not accountability. A neat correction log that never admits mistakes is not accountability. Software can make a decision traceable; the publisher still has to make, defend and revise the decision.

Implementation status at Hedegreen Research

Hedegreen Research has versioned articles and local source, claim and correction components in development. This article does not establish that the complete workflow proposed here is operating across the publication. No newsroom or audience comparison was conducted for it.

What Fake News Actually Hates

The title makes a deliberately overstated claim. Fake news can be perfectly structured. Good reporting can be poorly documented. Neither the presence nor absence of metadata settles the facts.

What unsupported claims have more difficulty with is a public, precise challenge to the point where evidence becomes conclusion—provided somebody can actually inspect the evidence, and the publisher is answerable for what happens next.

The Alaska video was genuine. The story attached to it wasn't. The difference becomes visible when we separate the asset from the proposition and ask each to support only what it can establish.

That leads to a modest standard with demanding consequences: **for a consequential factual claim, can an outsider locate it, inspect the evidence for and against it, and follow its corrections?**

Not every assertion needs its own page. No schema makes a liar honest. But a correction should not have to depend on the reader noticing that yesterday's sentence quietly disappeared.

Structure does not make journalism true. It can make truth claims easier to inspect.

— Dennis Hedegreen, trying to see the structure

Sources and verification notes

[1] Full Fact, Sian Bayley, *[Old video of Alaskan glacier shared with claims it shows cause of Nepal flood](https://fullfact.org/world/old-video-alaska-glacier-shared-as-nepal-flood-cause/)* (<https://fullfact.org/world/old-video-alaska-glacier-shared-as-nepal-flood-cause/>), 21 September 2026. Secondary synthesis of the circulated claim and earlier reporting.

[2] AFP Fact Check, Devesh Mishra, *[Footage of water rushing from glacier filmed in Alaska, not Nepal](https://factcheck.afp.com/doc.afp.com.C6Q96G3)* (<https://factcheck.afp.com/doc.afp.com.C6Q96G3>), published 31 August 2026, updated 9 September 2026. Independent visual search and an interview with the pilot. The AFP story describes an Instagram upload dated 22 August, while the pilot says she reposted it on 21 August; the draft attributes the latter date to her account rather than treating the display date as independently settled.

[3] Schema.org, *[ClaimReview](https://schema.org/ClaimReview)* (<https://schema.org/ClaimReview>). Data model for a fact-checking review of a claim.

[4] Google Search Central, *[Simplifying the search results page](https://developers.google.com/search/blog/2025/06/simplifying-search-results)* (<https://developers.google.com/search/blog/2025/06/simplifying-search-results>), 12 June 2025, updated 8 September 2025. Specific Claim Review rich-result retirement and removal from Search Console reporting tools.

- [5] Google Search Central, *Fact check (ClaimReview) structured data* (<https://developers.google.com/search/docs/appearance/structured-data/factcheck>), accessed 20 October 2026. States that Google Search support is being phased out while Fact Check Explorer support persists; the page also still contains legacy rich-result eligibility text.
- [6] C2PA, *Content Credentials: C2PA Technical Specification 2.4* (https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html), especially §1.2, §1.3 and the trust model. Cryptographic validation of a declared media history is not an independent adjudication of the truth of associated editorial claims.
- [7] C2PA, *Security Considerations 2.4* (https://spec.c2pa.org/specifications/specifications/2.4/security/Security_Considerations.html), especially §§2.6 and 4.3.1. Documents privacy-preserving options and the possibility of manifest stripping.
- [8] Reporters Without Borders, *Journalism Trust Initiative: everything you need to know about news media certification* (<https://rsf.org/en/journalism-trust-initiative-tout-comprendre-%C3%A0-la-certification-des-m%C3%A9dias-d-information>). Defines the scope as newsroom procedures, not individual content.
- [9] International Fact-Checking Network, *The commitments of the Code of Principles* (<https://ifcncodeofprinciples.poynter.org/the-commitments>). Describes methodological commitments, independent assessments and an explicit source-safety exception.
- [10] The Trust Project, *The 8 Trust Indicators* (<https://thetrustproject.org/trust-indicators/>). Organizational and article-level disclosure indicators.

Editorial boundary: No audience experiment or field deployment was conducted for this article. The hypothetical claim record, downstream scenario and proposed standard are normative examples, not observations of Hedegreen Research V3 or a functioning external system. The Alaska/Nepal example uses AFP’s reporting and Full Fact’s synthesis; these are not treated as two independent original investigations. This editorial pass did not contact the pilot or authenticate the raw media directly.

RELATION MEMORY

CONTINUES [“AI Makes Mistakes” Is Not an Argument](#) · Moves from comparing workflow failures to locating false contextual claims attached to authentic media.

SOURCE LAYER [Claims Need States](#) · Dated epistemic assessments remain editorial judgments tied to evidence.

SOURCE LAYER [Sources Should Travel With Claims](#) · The false-context case requires the source and its contextual proposition to remain distinct.

SOURCE LAYER [Corrections Need Memory](#) · A disputed proposition needs visible change reasons and retained history, without promised universal propagation.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/FAKE-NEWS-HATES-STRUCTURE/](https://HEDEGREENRESEARCH.COM/ARTICLES/FAKE-NEWS-HATES-STRUCTURE/)