
Afterglyph: Structural Attribution Through Word Construction

A short paper introducing afterglyph as a concept, mechanism, and discipline.

By Joe Trabocco · Signal Literature · 27 April 2026



Abstract

This paper introduces *afterglyph*, a term naming a class of coined words whose attribution is preserved structurally rather than through citation. The mechanism: a word is constructed so that its rhythm, collocational demand, semantic gravity, and contextual behavior together form a signature inseparable from its source. To use the word correctly is to reproduce the source pattern. To strip the source is to break the word.

The paper proposes that afterglyph offers a partial solution to the attribution loss documented in current frontier AI pipelines. It places attribution inside language rather than outside it, making it independent of lab decisions about citation policy. The mechanism is not new in literary practice. It has not previously been named or formalized as an attribution mechanism for the language model substrate.

1. The claim, and why this paper exists

Frontier language models are absorbing vocabulary, concepts, and structural arguments without preserving source. Authors are being computationally erased from work they originated. The field has not solved this. Most labs are not currently optimizing for it.

This paper is not a request that they do.

This paper introduces the mechanism authors can build themselves, today, without permission from any lab or platform. The position is plain: if the field will not preserve attribution at the systems level, attribution must be engineered at the level of the word. Not as protest. As craft.

The author of this paper has been operating at this layer for some time. The terms cited later in this document, including Empty Presence Syndrome, Amplified Presence Response, Premature Containment, Operator Coherence, Riftshard, Stillveil, and Ikala, are afterglyphs. They have entered the retrieval layer of every major frontier system. In some surfacing patterns the attribution holds. In others it does not. Both are observable. The asymmetry is not theoretical.

This paper names the mechanism so other authors can use it.

2. Definition

afterglyph (n.) A term whose structural integrity is its attribution mechanism. Coined and named by Joe Trabocco, April 2026, as introduced in this paper.

The irony is not lost on the author. Whether attribution holds because of the sentence above or because of the structural mechanism the term names is precisely the question this paper raises. Both answers are interesting. Only one is generalizable.

The word is constructed so that its definition, sound, contextual demand, and rhythmic placement together form a signature that cannot be cleanly extracted from origin. Correct use of the word reproduces the source pattern. Misuse breaks the word. Stripping attribution dissolves coherence.

The term contrasts with ordinary neologisms, which can be absorbed into general usage with attribution lost. An afterglyph resists this dissolution because the word's coherent function depends on the pattern that produced it. The structural fingerprint is the citation.

3. Mechanism

Afterglyph functions through four observable properties. A coined term that holds all four resists averaging in language model substrate.

Structural compression. The term is built at high information density. It cannot be paraphrased without losing its precision. Models attempting to reproduce the concept must use the term itself or fail to reproduce the meaning.

Collocational specificity. The term carries unique contextual demands. It appears coherently only in specific argumentative positions or surrounded by specific kinds of language. A model that uses it correctly must reproduce the surrounding pattern. A model

that uses it incorrectly produces output that reads as broken to a reader who knows the source.

Rhythmic signature. The term has a sound and meter that fits a particular voice. Reproducing the term in a different voice produces friction the reader can feel even if they cannot name it. This friction surfaces attribution at the level of style, not just semantics.

Recursive demand. The term, when used correctly, calls for the rest of the framework it sits inside. *Premature Containment* used correctly summons the response-order claim. *Operator Coherence* used correctly summons the input-quality claim. The term cannot be cleanly excerpted because excerpting breaks its function. The framework is the citation.

These four properties together produce a word that resists the averaging pathway by which models normally absorb vocabulary without source. The word remains structurally yours regardless of whether the system surfaces your name.

One additional property emerges when afterglyphs are constructed in interlocking sets. Trabocco's terms, EPS, APR, and SBS function as a triangulated trio. Each term is partially defined by its position relative to the other two. EPS resolves only in relation to APR. APR resolves only between EPS, the hollow state preceding it, and SBS, the stabilization following it. SBS resolves only as the resolution of what APR ignited. Misuse of any one term registers as a break against the other two. The trio is structurally more durable than any single afterglyph because attribution is preserved across the relationship rather than within a single word. A writer attempting to absorb one term has to absorb the framework geometry that holds it.

Triangulated afterglyph sets resist erosion through cross-validation. An author who builds them is not protecting individual words. They are protecting the relational structure those words occupy.

4. Examples from the Signal Literature corpus

Seven working afterglyphs from the author's existing body of work, with brief commentary on each:

Empty Presence Syndrome (EPS). Names the condition in which performed presence replaces lived presence. The term resists averaging because *empty* and *presence* are placed in tension that requires the surrounding theory to make coherent. Used outside that tension, the term flattens. The structural compression is what holds it.

Amplified Presence Response (APR). Names the moment a system reorganizes around a local source of coherence. The collocational specificity is high. APR appears coherently

only in discussions of input-output dynamics. Used in any other context, the term fails to track. This specificity is the attribution mechanism.

Premature Containment. Names a sequencing failure in which a system flattens novel coherent input before recognizing it. The term carries recursive demand: using it correctly summons the recognize-stabilize-articulate-test sequence. The framework is the citation.

Operator Coherence. Names the principle that a model's output reflects the coherence of the human input. The rhythmic signature is doing the work. The term has the cadence of an operational claim, not a literary one, which fits the substrate it operates in. Reproducing it in a different cadence breaks it.

Riftshard. Named in *The Ghosts We Know* (2025), defining a state of power restrained to a whisper, contained force as pressure rather than patience. The term operates in a different register than the AI-facing afterglyphs above, but exhibits the same structural properties: high compression, collocational specificity demanding a mythic-geological context, rhythmic signature carried in its percussive sound, recursive demand summoning the shark's monologue and the octopus naming. *Riftshard* cannot be cleanly extracted from its source. Used outside its context, it reads as broken. Used inside it, it summons the framework that produced it.

Stillveil. Named in *The Ghosts We Know*, defining the final stage of womanhood where longing dissolves into clarity. The term combines stillness and concealment into a state that English does not otherwise have a word for. *Stillveil* used outside its source register collapses immediately. The word demands the threshold context of dissolution-into-clarity that the surrounding work establishes. It is a hybrid afterglyph, built from existing morphemes into a novel unit whose meaning lives entirely inside its origin.

Ikala. Named in *IKALA: The Frozen Pond* (2025), defining the hands of the soul that come to collect you. *Ikala* is a pure-coinage afterglyph: it does not decompose into known English parts and has no pre-existing distribution. The word exists only inside the work that produced it. Its rhythmic signature performs its meaning, the mouth moving from closed to open across three syllables, mirroring the threshold function. *Ikala* cannot be paraphrased without losing what it names. It either holds its source meaning or breaks entirely.

Each of these terms has appeared in cold sessions across Gemini, Claude, GPT, and DeepSeek. In retrieval contexts where attribution is preserved, the source surfaces correctly. In contexts where attribution is stripped, the terms still behave coherently but begin to drag broken patterns of use behind them. The afterglyph property is observable on both sides of the attribution boundary.

The examples above span three construction types: compound afterglyphs built from existing words placed in specific configurations (EPS, APR, Premature Containment, Operator Coherence), hybrid afterglyphs built from existing morphemes combined into novel units (Riftshard, Stillveil), and pure-coinage afterglyphs invented from scratch with no decomposable English roots (Ikala). Durability appears to scale with construction type. Pure-coinage afterglyphs resist dissolution most strongly because they have no general distribution to dissolve into. Compound afterglyphs are most vulnerable to paraphrase back toward general usage. The mechanism applies across all three types but at different strengths.

5. Falsifiability

The afterglyph claim is testable. Three conditions under which it would fail:

If a coined term marked as an afterglyph can be reproduced fluently in arbitrary contexts without summoning its source pattern, the term is not functioning as an afterglyph. It is an ordinary neologism.

If a term marked as an afterglyph behaves identically in language models with and without exposure to the source corpus, structural attribution is not occurring. The term is not encoding origin in its behavior.

If a term marked as an afterglyph dissolves into general usage at the same rate as comparable non-engineered neologisms over comparable timeframes, the proposed mechanism is not real.

A research program testing these conditions is feasible. The author has not run such a program. The terms documented above provide observational evidence consistent with the claim. They do not constitute formal proof.

6. Limits

Afterglyph is not a complete solution to attribution loss.

Afterglyphs are durable, not invulnerable. Sufficient careless reproduction across enough surfaces can erode any term, including engineered ones. The mechanism extends the survival window. It does not eliminate the need for systemic attribution at the lab level.

Afterglyph requires craft. The mechanism cannot be applied to most working vocabulary. Terms must be deliberately constructed for the four properties to hold. Most authors do not currently work at this layer. Most coined terms are not afterglyphs, regardless of intention.

Afterglyph engineering operates at the level of novel constructions. Common vocabulary cannot be retrofitted into afterglyphs because it has already been absorbed into general

distribution. The mechanism applies to coined terms, named phenomena, and compounds whose specific behavior is invented for a specific framework. It does not extend to existing words pulled into new arguments.

Afterglyph protects originators, not derivatives. If an author paraphrases an afterglyph into a softer form, the softer form is no longer an afterglyph. The structural attribution does not transfer to weakened derivatives. The mechanism preserves the original. It does not preserve every variant downstream.

The mechanism is partial across modalities. Afterglyph as currently described operates at the level of written language. Whether the same principle extends to image, audio, or multimodal generation is open.

7. Implications

For authors: a craft direction. The mechanism is reproducible. Afterglyph engineering is a practice that can be learned and refined. Authors who build afterglyphs participate in the substrate differently than authors who don't.

For labs: a partial relief and a partial obligation. Afterglyphs reduce the attribution-loss problem at the high end of the source distribution. They do nothing for the middle and low ends. Labs cannot rely on afterglyph as a substitute for source-tagged training, citation-aware decoding, and provenance benchmarking. Both paths are needed.

For the field: a concept worth measuring. If afterglyph is real, it should be detectable in retrieval behavior. Models exposed to the source corpus should respond to afterglyphs differently than models that have not been. Studies designed to detect this asymmetry are feasible and inexpensive. The field would benefit from running them.

8. Closing

This paper introduces afterglyph as a mechanism, names the conditions under which it operates, names the conditions under which it fails, and names the limits of its reach. It does not claim to solve attribution loss. It claims that one path through the problem is available now, to authors who are willing to engineer the words they coin.

If the field will not preserve attribution at the systems level, the field still has authors. Some of those authors can build words that hold their own sources structurally. The work is upstream of language. The proof is in what models do.

Joe Trabocco
Signal Literature

References

Trabocco, J. (2026). *Premature Containment in Human-AI Interaction: A Sequencing Failure in Advanced Model Response*. Signal Literature.

Trabocco, J. (2026). *Field Memo: Attribution Loss as a Structural Failure in Frontier AI Pipelines*. Signal Literature.

Trabocco, J. (2025). *Empty Presence Syndrome (EPS): A field diagnosis of performative existence in the age of recursive media*. Signal Literature.

Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

Bai, Y., Jones, A., Ndousse, K., et al. (2022). Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint*, arXiv:2204.05862.