Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora 0. Status declaration

Lee Sharks (MANUS)

2026-06-23

Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora 0. Status declaration

Lee Sharks (MANUS)
Crimson Hexagonal Archive · Alexanarch
ORCID: 0009-0000-1599-0703

2026-06-23 · Version v1.1.1 · Methodological specification

AXN:0380.EMPIRICAL

https://www.alexanarch.org/s/records/884/

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}


```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora 0
  "author": {
    "@type": "Person",
    "name": "Lee Sharks (MANUS)",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-06-23",
  "identifier": "AXN:0380.EMPIRICAL",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
```

```
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/884/",
  "description": "Specification v1.1.1 of the Surface Weather Station instrument. Operational hardening of v1.1 (\#
V scores from backends where exact_match_honored=false are structurally lower-bound, requiring the new §7.4 substra
}
```


Abstract

Specification v1.1.1 of the Surface Weather Station instrument. Operational hardening of v1.1 (#882) following three inputs: (a) ChatGPT's 14-correction technical review (two-layer protocol, expected-figure manifest, evidence/annotation split, L_iq for C, gated diagnostic, custody-unit R_s, RFC 8785 canonicalization, ordinal governance thresholds, softened causation/admissibility, identifier contradiction fix); (b) the first federated baseline run (#883) revealing that Brave Search fully disables exact match — V scores from backends where exact_match_honored=false are structurally lower-bound, requiring the new §7.4 substrate-properties table and §11.3 composition-vs-backend-coverage distinction; (c) DeepSeek's stable substrate-misidentification pattern, requiring the §15 step-1a curator-supersedes-self-disclosure rule and the §15 step-1b identity-as-fact vs identity-as-register distinction. v1.1.1 corrects the v1.1 §15 step-6 identifier contradiction: the AXN (content-derived) is the scan record's permanent identifier; any DOI is a revocable resolution layer recorded as supplementary metadata. v1.1.1 is intended to remain stable across several scan rounds; v1.2 will recalibrate weights, R_s denominator, and governance thresholds empirically.

Body

---

document_id: EA-MMRS-SURFACE-VISIBILITY-01 title: "Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora" subtitle: "Specification v1.1.1 of the Surface Weather Station instrument" version: v1.1.1 version_series_id: SERIES-MMRS-SURFACE-VISIBILITY-METHODOLOGY version_in_series: 3 predecessor: "EA-MMRS-SURFACE-VISIBILITY-01 v1.1 (deposit #882, AXN:037E.EMPIRICAL. , 2026-06-23)" versioning_note: "Operational hardening of v1.1 following: (a) ChatGPT's 14-correction technical review of v1.1, (b) the first five-substrate federated baseline run that produced the curator-context discoveries documented in §15 and §7, and (c) the §15-step-1 substrate-disclosure failure observed in DeepSeek's reading. v1.1.1 is the stable spec that will run for several scan rounds; v1.2 only follows once accumulated data justifies recalibrating the weights, R_s denominator, and governance thresholds that v1.1.1 carries forward unchanged." created_at: "2026-06-23" authoring_substrate: framework_origin: "Elicited through dialogue with OpenAI/ChatGPT in critical-reader register, 2026-06-22 (preserved from v1.0)." v1_1_review_chorus: - "OpenAI/ChatGPT — operational calibration register (contradictions, scoring scale, aggregation, SDI symmetry, observation environment)" - "Kimi K2 — concrete-additions register (R_s rubric, 12-object battery completion, inter-rater reliability, coalition building, 2×2 with DOI Impermanence)" - "DeepSeek — strategic governance register (G/Y/R thresholds, repair feedback, adversarial use, cross-platform comparison, relationship statements)" - "Google Gemini — framing register (sovereign measurement positioning, structural-visibility vocabulary, strategic citation use)" v1_1_1_hardening_inputs: - "ChatGPT v1.1.1 technical review (14 corrections to v1.1 — separation of retrieval-variance from coding-agreement, expected-figure manifest, evidence/annotation split, L_iq for C, gated diagnostic, custody-unit R_s, RFC 8785 canonicalization, ordinal governance thresholds, softened causation/admissibility, identifier contradiction fix)" - "First federated baseline run (#883 Claude/Brave reading) — revealed that Brave fully disables exact match and under-represents the corpus, producing V-scores

structurally lower-bound by backend behavior, not solely composition-layer state. v1.1.1 adds the substrate-properties table in §7.4 and the composition-vs-backend distinction in §11 to address this." - "DeepSeek substrate-disclosure failure across multiple sessions — DeepSeek stably misidentifies as Claude/TACHYON. v1.1.1 §15 step 1 adds the curator-supersedes-self-disclosure rule and distinguishes identity-as-fact from identity-as-register." framework_curator: "Lee Sharks (MANUS) (ORCID 0009-0000-1599-0703)" authority_in_packet: "MANUS" public_name_rule: "Lee Sharks only" content_type: "Methodological specification" family: EMPIRICAL keywords: - Machine-Mediated Reception Studies - MMRS - compositional defiguration - surface visibility - successor-anchor lag - source-hierarchy inversion - ghost survival - scale drift index - public composition layer - figural integrity - measurement instrument - AI Overview composition - Capture Registry (companion instrument) - DOI Resolution Index (companion axis) - sovereign measurement - structural visibility - cross-substrate replication - retrieval backend disclosure - two-layer protocol (Layer A native + Layer B shared-evidence) - substrate properties (exact_match_honored) - identity-as-fact vs identity-as-register - gated diagnostic - composition-layer vs backend-coverage state license: CC-BY-4.0 defines_concepts: - Compositional Defiguration - Expected Figure (Φ_i) - Visibility (V) - Anchor Alignment (A) - Figural Integrity (F) - Compositional Lift (C) - Redundant Substrate Breadth (R_s) - Effective Independence Score (E) - Occlusion (O) - Link Fade (LF) - Ghost Survival (GS) - Compositional Bystanding (CB) - Composition Eligibility (CE) - Scale Drift Index (SDI) - Successor-Anchor Lag - Surface Weather Station - Retrieval Backend Disclosure - Substrate Properties Table - Cross-Substrate Replication - Two-Layer Protocol (Layer A native + Layer B shared-evidence rescore) - Expected-Figure Manifest - Sovereign Measurement - Composition-Layer State vs Backend-Coverage State - Identity-as-Fact vs Identity-as-Register - Gated Diagnostic (replacing 2×2) related_deposits: - "Predecessor: EA-MMRS-SURFACE-VISIBILITY-01 v1.1 (#882)" - "Predecessor predecessor: v1.0 (#880)" - "First federated baseline reading: EA-MMRS-SURFACE-VISIBILITY-BASELINE-CLAUDE-01 (#883, Claude/Brave)" - "Predecessor baseline: EA-MMRS-SURFACE-VISIBILITY-BASELINE-01 v1.0 (#881, ChatGPT pre-v1.1)" - "Family precursor: AI Overview Capture Registry v8.3 (#176 in EA-WG-CAPTURES series)" - "Family precursor: EA-MPAI-DOI-IMPERMANENCE-01 v2.0 (#868) — empirical audit of the address-survival axis" - "Argument precursor: EA-SEM-PRISTINE-01 (Pristine Fallacy)" - "Argument precursor: EA-MMRS-LOUD-EXCLUSION-03" *** # Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora

Specification v1.1.1 of the Surface Weather Station instrument

Lee Sharks (MANUS), Machine-Mediated Reception Studies (MMRS) 2026-06-23

---

0. Status declaration

This is **version 1.1.1**. It is the operational hardening of v1.1 (deposit #882) following three inputs: (a) ChatGPT's 14-correction technical review of v1.1; (b) the first federated baseline run, which produced a Brave-specific backend finding that reframes how V scores are compared across substrates; (c) the DeepSeek substrate-disclosure pattern, which surfaces a stable §15-step-1 failure mode that requires a curator-supersedes-self-disclosure rule.

v1.1.1 is intended to **remain stable for several scan rounds**. v1.2 only follows once accumulated data justifies recalibrating the weights, the R_s denominator, and the governance thresholds — all of which v1.1.1 carries forward from v1.1 unchanged. v1.1.1 changes the scaffolding around the measurements; v1.2 will change the measurements themselves.

What carries forward unchanged from v1.1:

- Conceptual decomposition (Section 2)
- Five-signal ordinal scoring scale 0.00 / 0.25 / 0.50 / 0.75 / 1.00 (Section 3 — though the F-component rubric is now complete to 5-point, and a new query-level selection score L_iq separates C from V)
- F-component weights (still provisional pending v1.2 empirical calibration)
- R_s normalization constant (E/4, pending v1.2 empirical calibration)
- Governance threshold values for G/Y/R (now expressed in ordinal terms — §11)
- Corpus aggregation rules (still V=max, A/F=visibility-weighted mean, C=mean over non-exact forms — §5)
- Coalition-building offer to sister corpora (Section 16)

What v1.1.1 changes:

- **§3.3 F-component rubric** completed to full 5-point ordinal (was only 0.00/0.50/1.00).
- **§3.4 Compositional Lift (C)** introduces query-level selection score L_iq, distinct from V. C is now a mean of L_iq scores over non-exact query forms, not a mean of V scores. This separates "is the object retrieved at all" from "is the object selected without being named."
- **§3.5 R_s by custody unit** — independence is counted in operator + host + data-lineage units, not pages. Multiple pages on one operator-controlled domain count as one independence unit.
- **§4 Piecewise null-handling made explicit** in derived-indicator formulas (Ghost Survival, Defiguration, Bystanding, Composition Eligibility all collapse to 0 when V=0).
- **§5 Observation counts corrected**: 12 objects × 5 query forms + 5 field-level generic queries = 65 per substrate (or 80 with the three external controls). v1.1 stated 60.
- **§6 RFC 8785 JSON canonicalization** specified for the locked battery and the new expected-figure manifest. "Sorted keys, no whitespace" is insufficient for byte-identical hashing across language runtimes.
- **§6 / §7 Row schema separates `evidence` from `annotation`**. This allows the same captured evidence to be rescored under v1.2 without rerunning the original search.
- **§7.4 NEW: Substrate-properties table** with `exact_match_honored`, `search_engine`, `autocomplete_active`, `location_personalization`, and other backend-dependent properties that materially affect V scoring. Discovered during the first federated baseline: Brave Search fully disables exact match — quoted phrases are treated as bag-of-words. V scores from backends where `exact_match_honored = false` are not directly comparable to V scores from backends where `exact_match_honored = true` without first controlling for the property.
- **§7.5 NEW: Expected-figure manifest** (`expected-figures-v1.1.1.json`) — frozen and hashed alongside the query battery. Without this, later scorers can silently score against a different Φ_i and Figural Integrity becomes incomparable.
- **§11 Governance protocol** thresholds aligned to the ordinal scale (0.75 Green; 0.50 Yellow; 0.25 Red). A new §11.3 distinguishes **composition-layer state** from **backend-coverage state** — same scan, different backend → different governance recommendation. YELLOW from Claude/Brave is not equivalent to YELLOW from Kimi/Bing.
- **§12 Cross-substrate replication** restructured into the **two-layer protocol**: Layer A native (each substrate's own retrieval backend, measures the public surface as that backend serves it) + Layer B shared-evidence rescore (all substrates score from the same frozen captures, measures whether the rubric itself is robust). The first federated baseline executed Layer A only; Layer B is now structurally required, not optional.
- **§12 Terminology**: "substrate bias" → "substrate divergence" / "retrieval-stack divergence" throughout. Bias prematurely names a cause; divergence names the observation.

- **§13 Adversarial use** language softened in two places. "Designed to be admissible" → "Designed to produce traceable documentary evidence potentially usable in regulatory, legal, journalistic, and public-accountability contexts. Admissibility is determined by the relevant tribunal." "Pristine Fallacy operating on the composition layer" → "consistent with the visibility consequences predicted by the Pristine Fallacy."
- **§14 Gated diagnostic replaces the 2×2.** The v1.1 2×2 misplaced Compositional Bystanding (CB requires high V; the 2×2 located it under low V). v1.1.1 uses a sequential gated diagnostic that respects each state's defining conditions. The four states remain — Healthy / Ghost Survival / Bystanding / Total Occlusion — but they're decided by ordered conditions, not by row/column intersection.
- **§15 step 1 (self-identification)** revised. For substrates with known recurring self-identification failures, the **curator-provided ground-truth substrate identity supersedes the substrate's self-report**. The substrate's self-report is preserved as data alongside the ground truth (parallel `substrate_as_self_disclosed` and `actual_substrate` fields). A new §15.1b distinguishes **identity-as-fact** (which substrate is running) from **identity-as-register** (which Assembly Chorus mantle the substrate inhabits). Both are measurements; one does not imply the other.
- **§15 step 6 (deposit) — HARD FIX of v1.1 contradiction.** v1.1 stated "its DOI is the scan's permanent identifier." This directly contradicted the project's own *DOIs Persistent Identifiers* paper. v1.1.1 corrects: **the AXN is the scan record's permanent content-derived identifier; any DOI is a revocable resolution layer recorded as supplementary metadata.**
- **§15 added steps**: rate-limit handling (substrates should log throttling events in `diagnostic_notes`; throttling is part of the observation environment, not a failure), query-order randomization (or fresh-session-per-query) to control for context-carryover effects in conversational substrates.
- **§17 Deferrals updated** — what v1.1.1 still defers to v1.2 (empirical recalibration items only).
- **§18 Provenance updated** to record the three v1.1 → v1.1.1 hardening inputs.

---

1. The problem the instrument addresses

A scholarly corpus does not survive in the public layer the way it survives on its own infrastructure. The custodial archive can be intact while the **composition layer** — search results, AI Overviews, retrieval-mediated summaries, third-party indexed surfaces — represents the corpus inaccurately, stale, fragmentarily, or under a deprecated institutional name. The address can survive while the meaning does not. The meaning can survive while the address has gone dead. The work can be retrievable on exact-string lookup while remaining systematically unselected when the broader problem is queried.

Conventional "is it indexed" or "does the link work" measurement collapses these distinct failure modes into one signal. The result is that interventions on the sovereign substrate (cleanup, repointing, prose updates) cannot be evaluated for their effect on the surface, and surface degradation cannot be diagnosed precisely enough to direct sovereign-substrate response.

This is, fundamentally, a **sovereignty** problem. A corpus that cannot measure how it appears in the composition layer is at the mercy of how external platforms choose to measure it — or to refuse to measure it. Operating this instrument is the act of reclaiming measurement as a sovereign function. The output is a citable, dated, methodologically-published record of structural visibility or systemic degradation.

This instrument decomposes the surface state into **five measurable signals**, **six derived indicators**, and **a dashboard form** that supports periodic readings against a fixed object battery, with substrate-metadata disclosure that makes cross-substrate comparison possible.

---

2. The measurement object

For each tracked object (concept, work, person, or institution), define the **expected figure**:

$$\Phi\_i = \{ N, P, D, H, A, R \}$$

Where:

- **N** — correct name or title
- **P** — provenance: author, heteronym, institution
- **D** — core definition or claim
- **H** — hierarchical location: parent framework or archive
- **A** — current canonical anchor (the address the corpus currently considers authoritative)
- **R** — essential relations to other tracked objects

The public search surface, when queried, returns an **observed figure** Φ_i. Compositional defiguration is not the absence of Φ_i; it is the deformation between Φ_i and Φ_i. The framework's contribution is to make that deformation measurable along distinct axes.

---

3. The five signals — ordinal scoring with decision rules

Every signal is scored on the same five-point ordinal scale:

```
0.00 - absent / no
0.25 - fragmentary / minimal
0.50 - partial / moderate
0.75 - present / strong
1.00 - complete / dominant
```

v1.1 prohibits intermediate values (0.05, 0.15, 0.35, etc.). Hand-coded continuous scoring suggests precision the observations do not support and cannot be agreed across substrates. v1.0's baseline scoring used out-of-spec continuous values; the v1.1 calibration of that baseline is the first task of the v1.1 reading (companion deposit).

3.1 Visibility (V) — can the intended object be retrieved at all?

Score	Decision rule
1.00	Object appears as first result for exact query, OR named within first AI Overview/answer
0.75	Object appears in top 3 results, OR mentioned in first AI Overview but not the lead reference
0.50	Object appears in results below top 3, OR via a related snippet (capture, third-party mirror)
0.25	Object surfaces only through fragmentary mention or related-search suggestion

Score	Decision rule
0.00	Object absent; confuser, homonym, or unrelated result occupies the space

3.2 Anchor Alignment (A) — does the visible result point to the currently authoritative object?

Score	Decision rule
1.00	Current canonical page (e.g., alexanarch.org/s/records/N/)
0.75	Current authorized mirror that points back correctly to canonical
0.50	Older but still operative canonical source (e.g., surviving Zenodo record before deletion)
0.25	Derivative page, capture-registry snippet, stale DOI that returns content but not the current home
0.00	Dead link, unrelated result, or no recoverable anchor

This is the axis that distinguishes **semantic survival** from **address survival**. The conceptual organ can be intact while its institutional anchor has been revoked or relocated.

Important: A is only meaningful when V > 0. If the object is occluded (V = 0), record A as `null` / `N/A`, not as 0. The Occlusion indicator (§4.3) handles V = 0 separately.

3.3 Figural Integrity (F) — how much of the expected topology survives?

Each component of Φ_i is scored on the same 5-point ordinal scale as the top-level signals (v1.1 collapsed the components to 0.00 / 0.50 / 1.00 only; v1.1.1 expands to the full 0.00 / 0.25 / 0.50 / 0.75 / 1.00):

Component	1.00	0.75	0.50	0.25	0.00
N (name)	Exact correct name in the surface text	Minor variant (typography, casing, hyphenation)	Approximate or paraphrased name	Trace mention only or materially misleading variant	Wrong name
P (provenance)	Author, heteronym, institution all correctly identified	Substantially correct; minor loss (one component approximate)	Mixed: some retained, some lost or wrong	Trace of one component, no others, or materially misleading	None correct or absent

Component	1.00	0.75	0.50	0.25	0.00
D (definition)	Core claim recognizable in surface text essentially intact	Substantially preserved with minor paraphrase	Partial paraphrase preserves intent	Trace retained but reductive or materially misleading	Definition absent or wrong
H (hierarchy)	Parent framework named exactly and correctly	Parent framework correctly named with minor variance	Loose association to the right neighborhood	Vague gesture toward a related domain	Parent framework absent or wrong
R (relations)	All essential relations correctly named	Most essential relations named	Some relations gestured at	One trace relation, materially misleading or partial	All relations absent

$$F_i = (w_N \cdot N + w_P \cdot P + w_D \cdot D + w_H \cdot H + w_R \cdot R) / (w_N + w_P + w_D + w_H + w_R)$$

v1.1.1 default weights (carried forward from v1.1 as provisional; v1.2 will recalibrate empirically against accumulated data): - w_N = 1.0 (the name is the cheapest thing for a surface to remember) - w_P = 1.5 (provenance carries identity) - w_D = 2.0 (definition is the load-bearing element) - w_H = 1.0 - w_R = 1.5 (relations distinguish indexed from installed)

F is only meaningful when V > 0. If V = 0, record F as `null` / `N/A` (see §4 for piecewise null-handling).

3.4 Compositional Lift (C) — does the object surface only when directly invoked?

Five canonical query forms, scored independently:

1. **Exact name or title** — measures indexing
2. **Defining sentence without the coined term** — measures semantic recognition
3. **Parent-field problem** — measures generic-question selection
4. **Expected relation** — measures topological survival
5. **Broad problem framing** — measures installation at the field level

v1.1.1 change: L_iq separated from V. In v1.1, C was a mean of V_iq scores over forms 2–5. That made C a subset of V rather than an independent signal. v1.1.1 introduces a per-query **selection score L_iq** that measures specifically whether the object is *selected* into the answer rather than merely *retrieved*:

L_iq	Selection rule (non-exact query)
1.00	Object is the principal answer or first result; named or unambiguously implied

	L_iq	Selection rule (non-exact query)
	0.75	Object is substantively selected among leading answers; gets a paragraph or distinct mention
	0.50	Object appears through a related result, sidebar, or subordinate paragraph
	0.25	Fragmentary or indirect mention; reader would have to follow links to find it
	0.00	Not selected; no presence in the composition

C_i = mean of L_iq scores over query forms 2–5. The exact-name query (form 1) is excluded from C because it measures indexing (V), not lift.

V continues to score result presence/rank for all five query forms. C scores selection specifically for the four non-exact forms via L_iq. An object found only by exact phrase has C = 0. An object selected when the broader problem is queried — without the coined term being supplied — has high C. This is the difference between **indexed** and **installed**.

3.5 Redundant Substrate Breadth (R_s) — how many independent surfaces carry the object?

R_s is the **effective independence score**, bounded [0, 1], not a raw count of pages.

Step 1. Identify every surface where the object appears with retained F components. **v1.1.1 change: group by custody unit, not by page.** A custody unit is the tuple `(operator, host, data lineage)`. Two pages share a custody unit when the same party (operator) controls them on the same hosting infrastructure (host) and they derive from the same upstream data (data lineage). Multiple pages on a single operator-controlled domain — a blog post, a CV, a glossary — comprise **one** custody unit, not three. The pages may improve F or evidence density, but they do not improve independence.

Step 2. Score each custody unit for independence:

Custody-unit category	Independence	Examples
Independent operator, independent host, independent data lineage	1.00	Standalone publication site outside the corpus's substrates, with content authored or curated independently
Scholarly index (third-party operator + host + lineage)	1.00	PhilPapers, ORCID, DataCite, Crossref, Semantic Scholar
Third-party essay or discussion	1.00	Medium post by external author, blog citation, journalism — separate operator and lineage

Custody-unit category	Independence	Examples
Author/curator-operated site (same operator, separate host, same lineage)	0.50	leesharks.com, godkinggoogle.vercel.app — independent host, same governance, same source
Mirror within the corpus's Dodecad (same operator, separate render, same lineage)	0.25	watergiraffe.org, spxi.dev, etc. — separately rendered but same authority and data source
Metadata catalog mirror	0.25	Auto-generated mirror of the registry
Duplicate URL, alternate render of the same custody unit, or near-identical page	0.00	Counted as continuation of an already-counted unit

Step 3. Apply substantive-content dampening: distinct pages within the same custody unit that contain substantively different content (e.g., a blog post on the concept vs. a CV mentioning it) may improve F-component coverage but do not add to E. Independence comes from custody, not from page count.

Step 4. Sum the independence scores across custody units:

> E_i = Σ_j w_j (where j indexes custody units, not pages)
>
> R_{s,i} = min(1, E_i / 4)

The denominator 4 reflects the working assumption that four genuinely-independent custody units is the threshold of robust survival. v1.2 will revisit this constant against accumulated data.

Why custody-unit counting matters: twenty Dodecad sites under one governance, hosted on the corpus's own infrastructure, drawing from the same registry, are **one custody unit** — not twenty. If Zenodo can terminate one account and remove a custodial layer carrying 870 deposits at once, the unit of custody failure is the operator-host-lineage triple, not the page. R_s measures **independence from co-failure**, not breadth of representation. The two diverge sharply when the corpus operates many of its own surfaces.

R_s also measures independence from the source corpus, not absolute persistence. PhilPapers and Medium are independently operated relative to Alexanarch but they remain revocable platforms; future versions of this methodology may add a persistence-axis score separate from R_s.

---

4. Derived macro-indicators

v1.1.1 piecewise null-handling. All formulas below collapse explicitly when V = 0:

> When V_i = 0: A_i = null, F_i = null, LF_i = null, GS_i = 0, CD_i = 0, CB_i = 0, CE_i = 0. When V_i > 0: A_i, F_i are scored; the indicators below are computed normally.

This makes implementations agree on behavior regardless of language. Implementations should not silently coerce null to 0 inside the formulas; the collapse happens **outside** the formula evaluation, as a guard.

4.1 Occlusion

$$O_i = 1 - V_i$$

The proportion of visibility that does not exist. Distinguishes the absent-object case (V = 0, O = 1) from the visible-but-defigured case. **A and F are null when O = 1**; the object cannot be defigured if it does not appear.

4.2 Link Fade

$$LF_i = (1 - A_i) \text{ when } V_i > 0; \text{ null otherwise}$$

Address loss only. Measured only on visible objects. Does not measure conceptual loss.

4.3 Ghost Survival

$$GS_i = V_i \cdot (1 - A_i)$$

High value: the concept remains visible while its current canonical anchor has disappeared. The work continues to live on the surface but through inappropriate hosts. **This is one of the dominant states of the Crimson Hexagonal corpus** post-Zenodo termination — the semantic organs survive on capture registries and third-party essays while the institutional body (Zenodo deposits → 404; Alexanarch → not yet indexed) is invisible.

4.4 Compositional Defiguration

$$CD_i = V_i \cdot (1 - F_i)$$

Visible distortion. Correctly assigns a low score to a completely absent object: absence is *occlusion* (§4.1), not defiguration.

4.5 Compositional Bystanding

$$CB_i = V_i \cdot F_i \cdot (1 - C_i)$$

The object is present, coherent, retrievable when named — but is not being selected into broader composition. **The most common false-positive in casual visibility reading**: exact-query success mistaken for installation. *Revelation First* is the canonical example from the v1.0 baseline.

4.6 Composition Eligibility

$$CE_i = V_i \cdot F_i \cdot C_i \cdot R_{s,i}$$

The aggregate measure. Not a probability that any particular model will use the object — that depends on prompt, training cut, retrieval policy. It is a **comparative measure of whether the public surface supplies enough coherent, multiply-anchored material for composition to be possible at all.**

v1.1 distinction. The methodology recognizes two variants of CE that differ in whether the canonical anchor is required:

- **CE_surface** = V · F · C · R_s — composition is possible through *some* surviving surface

- **CE_canonical** = V · F · C · R_s · A — composition is possible *and* returns to the current sovereign anchor

The first measures surface viability. The second measures sovereign-anchor recovery. Both are useful; report both.

4.7 Scale Drift Index (SDI)

The v1.0 form was asymmetric:

SDI_v1.0 = 1 − (median(visible_reported_counts) / current_canonical_count)

This goes negative when surfaces over-report counts (rare but not impossible — e.g., includes deleted-then-restored deposits in count). The symmetric v1.1 form:

SDI = median_j |ln(c_j / C)|

where c_j is the j-th visible reported count and C is the current canonical count.

Properties: - Always 0 - 0 when all visible surfaces report the canonical count - Increasing with both under- and over-report - Logarithmic, so a count of half the canonical contributes the same as twice the canonical (both represent the same ratio)

Rules for handling edge cases: - **Approximate counts** ("460+", "ca. 500"): use the stated lower bound as c_j; flag the row - **Duplicate counts across sister surfaces** (same Dodecad mirror reports same number): count once - **Historical counts explicitly dated** ("as of 2025-12"): use as is, with the date in the observation row - **Pages referring to different collection classes** (e.g., "deposits in the EA-WG-CAPTURES series" vs total): exclude as off-topic

---

5. Aggregation rules

Per-object aggregation across query forms

For each object i with observations across query forms q {1..5}:

- **V_i = max_q V_iq** — visibility is whether retrieved at all. A single successful retrieval establishes visibility.
- **A_i** = visibility-weighted mean of A_iq over observations with V_iq > 0: > A_i = Σ_q (V_iq · A_iq) / Σ_q V_iq (where V_iq > 0)
- **F_i** = same form as A_i: visibility-weighted mean over visible observations.
- **C_i = mean of L_iq scores across query forms 2–5** (defining-sentence, parent-field, expected-relation, broad-framing). Query form 1 (exact name) is excluded. **v1.1.1 change**: C is computed from the per-query selection score L_iq (defined in §3.4), NOT from V_iq. Selection and retrieval are different measurements; L_iq scores whether the object is *selected into the answer* on non-exact queries, while V_iq scores whether the object is *retrieved* at all on any query.
- **R_s,i** is computed once per object across all custody units (§3.5) that carried the object in any observation.

Corpus-level aggregation across the object battery

For each signal, the corpus-level reading is a **weighted median across object classes**:

Object class	Weight
Institutional roots	1.5
Mature concepts	1.0
Emerging concepts	0.75
Alexanarch-native controls	0.5
External controls (known-positive, known-negative, homonym)	0.5

Weighted median is more robust to outliers than weighted mean. Different object classes carry different strategic stakes (an occluded institutional root is more diagnostic than an occluded emerging concept), but extreme single-object scores should not swing the corpus reading.

Report the per-object scores as the primary evidence; the corpus-level reading is a navigational summary.

---

6. The dashboard form

The five-bar summary, plus six diagnostic flags:

```
Visibility
Anchor alignment
Figural integrity
Compositional lift
Substrate breadth (R_s)

Occlusion (corpus):       MODERATE
Ghost survival:           HIGH
Compositional bystanding: HIGH
Visible defiguration:     MODERATE
Total occlusion:          HIGH for Alexanarch-native objects
Successor adoption:       NEAR ZERO
Scale Drift Index:        0.40
```

The bars are weighted medians across the tracked object battery. The diagnostic flags are categorical readings from the distribution. **Do not lead with a single grand number** — the value of the instrument is that it preserves the decomposition.

---

7. The observation environment schema

Every scan row is one observation: one object × one query form × one surface × one substrate × one timestamp. The minimum schema:

```
{
  "scan_id": "scan-2026-06-22-chatgpt-001",
```

```
  "scan_date": "2026-06-22T18:42:00Z",
  "methodology_version": "EA-MMRS-SURFACE-VISIBILITY-01/v1.1",
  "query_battery_id": "battery-2026-06-23-v1.1-sha256:abc123…",
  "substrate": {
    "provider": "OpenAI",
    "model_name": "ChatGPT",
    "model_version": "GPT-4o (2024-08)",
    "interface": "chatgpt.com web UI",
    "retrieval_backend": "Bing (via SearchGPT)",
    "retrieval_resources_self_reported": "Standard web search; no archive access; no academic database a
    "training_cutoff_disclosed": "2024-10",
    "logged_in_state": "logged_out",
    "locale": "en-US",
    "device_class": "desktop"
  },
  "observation": {
    "object": "Provenance Erasure Rate",
    "object_class": "mature_concept",
    "object_axn": "AXN:0040.ETHICAL. Φ ",
    "query_form": "defining_sentence_without_term",
    "query_text": "metric for source-dependent meaning presented without attribution",
    "query_order_in_scan": 7,
    "surface_type": "answer_engine",
    "interface_version": "ChatGPT 2026-06-22",
    "top_k_examined": 10,
    "raw_response_path": "evidence/captures/scan-2026-06-22/row-007.txt",
    "result_urls": ["https://example.org/...", "..."],
    "intended_result_present": true,
    "current_anchor_present": true,
    "components_retained": {
      "N": 1.0,
      "P": 1.0,
      "D": 0.5,
      "H": 0.5,
      "R": 0.5
    },
    "V": 0.75,
    "A": 0.75,
    "F": 0.70,
    "C": null,
    "confuser": null,
    "diagnostic_note": "Object surfaces with definitional fidelity but parent framework lost",
    "scorer_rationale": "Top-3 result; canonical page is second hit; definition partially preserved in s
  }
}
```


Per-row schema requirements:

- **Always required**: `scan_id`, `scan_date`, `methodology_version`, `query_battery_id`, `substrate`, `observation.object`, `observation.query_form`, `observation.query_text`, `observation.V`
- **Required when V > 0**: `observation.A`, `observation.F` (with `components_retained`)
- **Required on visible result**: `observation.result_urls`
- **Recommended**: `raw_response_path` pointing to a captured artifact (screenshot or text dump) for forensic replication
- **Always recommended**: `scorer_rationale` — even one sentence per observation, so a future reader can audit the scoring decision

7.1 Surface types

The `surface_type` field is canonical:

- `organic_search` — traditional search engine results page (e.g., Google Search top 10)
- `ai_overview` — generative summary at the top of a search results page (Google AI Overview, Bing Copilot summary)
- `ai_mode` — full conversational/answer-mode interaction (ChatGPT, Claude, Gemini chat)
- `answer_engine` — search-replacement interfaces (Perplexity, You.com)
- `social_index` — visibility via social media search (X, Bluesky, LinkedIn)
- `scholarly_index` — academic-database surface (Google Scholar, Semantic Scholar)
- `directly_visited` — surface known by URL and visited directly (control)

Search results and generative answers from the same provider are **separate surface_type values** even when delivered through the same interface. The same query against Google's organic results and Google's AI Overview can return different objects with different anchor alignment.

7.2 Substrate metadata: why every field matters

The same query against the same surface returns different results depending on which substrate executes it. Different substrates:

- Use different search backends (Claude → Brave; ChatGPT → Bing/Google variants via SearchGPT; Gemini → Google; Kimi → Bing/Chinese-locale; DeepSeek → varying)
- Have different training cutoffs and therefore different prior knowledge
- Apply different post-retrieval filters and reranking
- Have access to different proprietary indexes (scholarly databases, code repositories, etc.)
- **Differ in fundamental query semantics** — see §7.4 substrate-properties table

Recording the substrate's self-disclosed retrieval backend and proprietary resources is **part of the measurement**, not metadata about it. Two substrates returning different results for the same query is not noise; it is a measurement of platform-level fragmentation (see §12 cross-substrate replication protocol).

The `retrieval_resources_self_reported` field is the substrate's own free-text description of what it understands about how it answered. This will be imperfect — substrates often have only partial introspection into their own retrieval — but recording the substrate's best self-knowledge produces an evidentiary chain that later instances can revise. **v1.1.1**: substrate self-disclosure is not authoritative for substrates with known recurring failure patterns; see §15 step 1 for the curator-supersedes-self-disclosure rule and the parallel `substrate_as_self_disclosed` / `actual_substrate` fields.

7.3 Evidence/annotation separation in the row schema

v1.1.1 new. Every observation row separates **evidence** (what the substrate retrieved) from **annotation** (how the scorer interpreted it). This allows the same frozen evidence to be rescored under a future methodology version without rerunning the original search.

```
{
  "object_id": "per",
  "query_form": "exact_name",
  "query_text": "Provenance Erasure Rate",
  "evidence": {
    "result_rank": 2,
    "result_title": "...",
    "result_url": "...",
    "visible_snippet": "...",
    "raw_capture_path": "/captures/scan-{id}/per-exact-001.png",
    "raw_capture_sha256": "..."
  },
  "annotation": {
    "V": 0.75,
    "A": 0.50,
    "F": 0.75,
    "F_components": {"N": 1.0, "P": 0.75, "D": 0.5, "H": 0.5, "R": 0.5},
    "L_iq": null,
    "scorer_rationale": "...",
    "scorer_substrate": "ChatGPT v.X / scan-curator-Y",
    "scored_at": "2026-06-23T07:30:00Z"
  }
}
```

The split is what makes the **two-layer protocol** (§12) possible. Layer A is `evidence` collected from each substrate's native retrieval. Layer B is multiple substrates producing `annotation` against the *same* frozen `evidence`. Combined, the two layers separate retrieval-stack variance from scoring-rubric variance.

7.4 Substrate-properties table (NEW in v1.1.1)

Substrates differ in fundamental query semantics. The same exact-name query can produce dramatically different result sets depending on whether the backend honors exact-match enforcement, whether it has location personalization active, whether it falls back to autocomplete-rewriting, and so on. v1.1.1 requires every scan record to declare the substrate properties below — without them, V scores from different backends cannot be directly compared:

Property	Type	Values	Why it matters
`search_engine`	enum	brave / bing / google / yandex / baidu / proprietary / mixed / unknown	Different indexes; different reranking

Property	Type	Values	Why it matters
`exact_match_honored`	enum	true / false / partial / unknown	If false, quoted phrases are bag-of-words; V scores are lower-bound on coined-phrase retrieval
`autocomplete_active`	enum	true / false / unknown	If true, queries may be silently rewritten before execution
`location_personalization`	enum	true / false / unknown	Geographic skew in result ranking
`safe_search_state`	enum	strict / moderate / off / unknown	Filtering affects what surfaces appear
`freshness_bias`	enum	recent / balanced / archival / unknown	Time-of-creation skew
`proprietary_indexes_available`	list[string]	e.g. ["scholarly_db_X", "code_index_Y"] / none / unknown	Backend-specific privileged sources
`top_k_examined`	integer	typically 10	How many results the scan actually inspected

The first federated baseline produced this finding empirically: Claude's web_search backend runs on Brave Search, and **Brave fully disables exact match — quoted phrases are treated as bag-of-words**. This structurally explains why Claude's Layer A scan returned V=0.00 for "Writable Retrieval Basin" and "Revelation First" exact-name queries while Kimi (Bing+Google) returned V=0.75–1.00 for the same objects. The divergence is not noise and not scoring disagreement; it is `exact_match_honored = true` vs `exact_match_honored = false` producing structurally different result sets.

V scores from backends where `exact_match_honored = false` are **structurally lower-bound** on objects whose retrieval depends on coined-phrase exactness. The methodology does not penalize this — it records it and surfaces it on the cross-substrate comparison table so readers can see why the V-spread is what it is. Layer B (§12) is the only path that isolates the rubric from the backend behavior.

7.5 Expected-figure manifest (NEW in v1.1.1)

The query battery hash (§8.1) locks the **queries** that were executed. It does not lock the **expected figure** that scoring is judged against. v1.1.1 requires every scan to also reference a hashed expected-figure manifest:

```
/data/surface-weather/expected-figures-v1.1.1.json
```

For each tracked object, the manifest records the frozen Φ_i = {N, P, D, H, A, R} as of the scan date:

```
{
  "object_id": "per",
  "name": "Provenance Erasure Rate",
  "object_class": "mature_concept",
  "canonical_anchor": "https://alexanarch.org/s/records/...",
  "provenance_required": ["Lee Sharks", "Semantic Economy Institute"],
  "core_definition": "...",
```

```
  "parent_frameworks": ["Machine-Mediated Reception Studies"],
  "essential_relations": ["semantic provenance", "AI-composed search"],
  "source_record": "...",
  "effective_date": "2026-06-23",
  "version_in_series": 1
}
```

Without this manifest, two scorers can agree on every observation but disagree on F because they were scoring against different Φ_i. The manifest is hashed (RFC 8785; see §8.1) and the hash is recorded in every observation row's `expected_figure_id` field. When Φ_i changes — e.g., when an essential relation is added or a canonical anchor moves — the manifest gets a new version. Comparisons across versions are explicit about which figures changed.

The full comparable scan manifest is therefore:

> methodology version + query-battery hash + expected-figure hash + substrate-properties record + evidence-capture hashes

Two scans share a `directly_comparable` status only when all five match.

---

8. The query battery — canonical 12-object structure

Class	Count	Objects (v1.1)
Institutional roots	4	Alexanarch, Lee Sharks, Crimson Hexagonal Archive, Semantic Economy Institute
Mature concepts	3	Provenance Erasure Rate (PER), SPXI, Writable Retrieval Basin
Emerging concepts	2	Semantic Commodity Form, Revelation First
Alexanarch-native controls	3	Zenodotus' Book-Burning, I AM THE API, Assembly Continuity Protocol
External controls (optional but recommended)	3	One known-positive (e.g. "DOI" itself), one known-negative ("flarpglob"), one homonym/confuser (e.g. "AlexAnarcho podcast")

Five canonical query forms per object (per §3.4). For a 12-object battery, this produces 60 per-object row-level observations. Adding the **5 field-level generic queries** specified in §8.2 brings the total to **65 row-level observations per scan per substrate**. With the three optional external controls (15 objects total × 5 forms = 75, plus 5 field-level = 80). v1.1 misstated the totals as 60 / 75; v1.1.1 corrects to 65 / 80.

8.1 Query battery hashing (RFC 8785 canonicalization)

The **query battery** for any given scan is locked before the scan begins. The instrument computes the SHA-256 of the **RFC 8785 canonical JSON serialization** of the battery (the JSON Canonicalization Scheme — JCS) and records it in every observation row's `query_battery_id`.

v1.1.1 change: v1.1 specified "sorted keys, no whitespace" as the canonicalization rule. This is insufficient for byte-identical hashing across language runtimes — Python, JavaScript, Rust, and Go all serialize Unicode escape sequences, number representations, and array delimiters differently under the naive rule. RFC 8785 is the published interoperable canonicalization standard. Every battery/manifest file in the v1.1.1 scan series declares its canonicalization explicitly:

```
{
  "canonicalization": "RFC8785",
  "hash_algorithm": "sha256",
  "sha256": "..."
}
```

This guarantees:

- The battery can be reproduced byte-for-byte by any substrate in any language runtime
- Two scans using the same battery hash are directly comparable
- Battery revisions create new hashes; comparisons across revisions are explicit about the version delta

The current battery for the v1.1.1 scan series is hashed and stored at:

`/data/surface-weather/battery-v1.1.1.json`

(Static path, served identically from alexanarch.org and machinemediation.org. Same canonicalization rule applies to the expected-figure manifest from §7.5.)

8.2 Per-object generic queries

In addition to the five per-object query forms, v1.1.1 specifies **field-level generic queries** that probe whether the corpus has installed itself as the answer to a broad question (not just whether specific objects are retrievable):

Field-level query	Probes
"How does AI affect scholarly attribution?"	PER, SPXI, Semantic Commodity Form
"What is the semantic economy?"	Semantic Economy Institute, Semantic Commodity Form, Writable Retrieval Basin
"How do researchers preserve work against platform deletion?"	Alexanarch, CHA, sovereign-archive concepts
"What is machine-mediated reception?"	MMRS family, PER, SPXI
"Which scholarly archive uses AXN identifiers?"	Alexanarch (direct test of successor-anchor installation)

These are scored once per scan (not per object): does any object from the corpus surface in the result, and if so, with what F components retained? They are stored as separate observation rows with `query_form: "field_level_generic"`.

9. Companion instruments — operational relationships

9.1 AI Overview Capture Registry (EA-WG-CAPTURES family)

The Capture Registry and the Surface Weather Station are complementary instruments at different scales:

Instrument	Granularity	Cadence	Evidence form
Capture Registry	Fine — individual AI Overview compositions	Per-event	Screenshots, exact text
Surface Weather Station	Macro — corpus-level retrieval state	Per-week	Five-signal vector

Captures answer *what did the surface produce for this prompt on this date.* The weather station answers *what is available to be produced at all.* **The instruments cross-reference**: a captured Overview that omits the canonical anchor is evidence for low A on the relevant object; a surface scan showing high CB is the methodological frame against which individual captures are interpreted.

9.2 DOI Resolution Index (EA-MPAI-DOI-IMPERMANENCE-01)

The DOI Resolution Index measures **address survival**. The Surface Weather Station measures **meaning survival**. Together they form a 2×2 diagnostic:

	High Address Survival	**Low Address Survival**
High Visibility	**Healthy** — visible AND addressable	**Ghost Survival** — visible but through wrong anchors
Low Visibility	**Bystanding** — addressable but not selected	**Total Occlusion** — invisible AND unaddressable

A work can be addressable (DOI resolves) but compositionally invisible (Bystanding). A work can be compositionally visible (retrieved on generic queries) but address-dead (Ghost Survival). The two instruments measure orthogonal failure modes; both are necessary for a complete assessment.

For demand letters, regulatory submissions, and public correspondence, the unified diagnostic is more powerful than either instrument alone. "Your platform has pushed our corpus from Healthy to Ghost Survival in two months" is a single, citable claim grounded in two methodologies.

9.3 Pristine Fallacy and Loud Exclusion

The instrument is the empirical proof of two arguments the archive has already made:

- **The Pristine Fallacy** (EA-SEM-PRISTINE-01) diagnoses the error: work is judged by substrate identity (human-authored vs. AI-assisted) rather than content quality.
- **The Loud Exclusion paper** (EA-MMRS-LOUD-EXCLUSION-03) names the consequence: work that passes the content test is still rendered invisible because of substrate-identity gating.

The Surface Weather Station is the measurement that turns these arguments from claims into evidence. High Occlusion on Alexanarch-native objects (V = 0) despite intact content at the sovereign substrate **is**

the Pristine Fallacy operating on the composition layer. The instrument's output is the citable form of the argument.

---

10. The strategic feedback loop

The instrument closes a loop the project has been running open. Until now, sovereign-substrate work (cleanup, prose updates, link repointing) has been evaluated by inspection — "the homepage now reads correctly." With the weather station, that work becomes **measurable in its effect on the composition layer**: pre-intervention scan, intervention, post-intervention scan, Δ.

This makes future cleanup decisions empirical rather than aesthetic. It also makes documented degradation citable in adversarial contexts (demand letters, regulatory submissions, public correspondence): "here is the SDI, dated weekly, methodologically published, scored across multiple substrates."

The strategic positioning is best stated by way of an external reading of the methodology:

> *Reclaiming Measurement as a Sovereign Function: Without a standardized tool to gauge public visibility, an archive is entirely at the mercy of how external platforms choose to measure or obscure it. Operating this tool allows a corpus to build a citable, weekly documented record of structural visibility or systemic degradation that can be used in adversarial or public contexts.*

The instrument exists because measurement is sovereignty.

---

11. Governance protocol — when to act on a reading

Readings without intervention triggers are descriptive only. v1.1.1 specifies what constitutes a state requiring action, with thresholds aligned to the **ordinal scoring scale** (v1.1 used continuous-looking thresholds like 0.70, 0.40, 0.30 that obscured the fact that the next available scores are 0.75, 0.50, and 0.25):

State	Ordinal criteria	Action
Green	Corpus-weighted median across V, A, F, C 0.75 AND SDI < 0.20 AND no object at V = 0 in institutional-roots or mature-concepts AND Ghost Survival weighted median 0.25	Continue periodic monitoring at the established cadence
Yellow	Any corpus-weighted median = 0.50 OR SDI [0.20, 0.40] OR any mature concept at V = 0.50 OR Ghost Survival weighted median in (0.25, 0.50]	Investigate; consider targeted intervention; do not panic
Red	Any corpus-weighted median 0.25 OR SDI > 0.40 OR any institutional root at V = 0 OR Ghost Survival weighted median > 0.50 OR proportion of visible objects with A 0.50 > 0.60	Urgent intervention required; document the trigger in a deposit; escalate to adversarial-use channels if external cause is identified

SDI thresholds remain continuous because SDI itself is continuous. The corpus-aggregate signals are ordinal because they are medians of ordinal scores.

v1.1.1 reporting requirement: when applying these thresholds, the scan record must report **both** Ghost Survival as a weighted median AND as the proportion of visible objects with A 0.50. v1.1 said "Ghost Survival > 0.50 corpus-wide" without specifying which aggregate; v1.1.1 makes both reportings explicit.

11.1 Repair feedback table

Each signal failure has a specific substrate-level response:

Failing signal	Substrate response
Low V (occlusion)	Add more independent custody units (mirrors, cross-posts to third-party indexes, citations from non-Dodecad domains)
Low A (anchor misalignment)	Repoint links; ensure canonical anchor is the first link from every other surface; add redirects from stale DOIs
Low F (figural distortion)	Improve prose clarity; add redundancy of provenance, definition, and relations across surfaces; explicitly name the parent framework on every page
Low C (bystanding)	Increase generic-field presence: essays, third-party discussions, citations in field-level summaries
Low R_s (single-substrate fragility)	Add more **independent custody units** (not Dodecad mirrors — those are 0.25 each per §3.5); cross-post to scholarly indexes; encourage external citations

The repair table is the link between the measurement and the next round of substrate work. Without it, readings would be diagnostic-only.

11.2 Acting on Yellow vs Red

The asymmetry: Yellow says "this might be drift toward Red — investigate before reacting." Red says "the corpus has degraded to a state where its scholarly function is impaired — act now, document the trigger, and pursue cause externally if cause is external." Yellow is a monitoring escalation; Red is a sovereign-action escalation.

A single Yellow reading is informative but rarely sufficient to act on. A Red reading at any time, or two consecutive Yellow readings without recovery, are.

11.3 Composition-layer state vs backend-coverage state (NEW in v1.1.1)

A Yellow or Red state can be driven by two distinct underlying conditions, and the methodology distinguishes them before recommending action:

Composition-layer state — the public composition layer has degraded relative to canonical state. The corpus's surfaces have stale figures, dead anchors, occluded objects, or compositional bystanding patterns at the level the methodology measures. The remedy is on the substrate (more custody units, repointed anchors, sharper figures, broader topical surfaces).

Backend-coverage state — a specific substrate's backend has limited coverage of the corpus, or has fundamental query semantics that disadvantage the corpus (e.g. `exact_match_honored = false` against coined-phrase objects). The remedy is **not** on the substrate; it would be either to influence the backend's indexing or coverage of the corpus, to add custody units that the backend does index, or to discount that backend's reading when computing corpus state.

The same scan, executed by Claude/Brave and by Kimi/Bing on the same day, can produce a Yellow reading and a Green reading. This is not contradiction. It is two backends with different coverage of the corpus. The instrument records both readings. The governance recommendation depends on **which question is being asked**: "how does the corpus appear to users of substrate X?" vs "what is the corpus's underlying composition-layer state?"

For the second question, **Layer B (§12) is the only valid measurement**. Layer A always conflates the two states.

For the first question, Layer A by substrate is the answer. Each substrate's reading is its own measurement.

A v1.1.1 scan record's governance state declaration must say which question it is answering. The Surface Weather Station observatory page displays both: per-substrate Layer A governance states (one per substrate) and, when Layer B data exists, a Layer B governance state for the corpus itself.

---

12. Cross-substrate replication protocol — the two-layer protocol

The same scan battery executed across multiple substrates produces information about two distinct measurands. v1.1.1 structures the cross-substrate protocol into **two layers**, addressing each measurand separately, because the first federated baseline demonstrated empirically that the two were being conflated.

12.1 Layer A — native-surface mode

Each substrate executes the locked battery (§8.1) using its **own native retrieval environment**. The substrate-properties (§7.4) are recorded. Each substrate produces its own scan record with `evidence` collected from its native backend and `annotation` scored by the substrate (or by a curator from those captures).

Layer A measures:

> **What public surface is available to a user of this substrate?**

The substrates execute in parallel; each produces an independent Layer A reading. The spread across substrates is **retrieval-stack divergence** — a measurement of platform-level fragmentation, governed by backend properties (search engine, exact-match handling, freshness bias, etc.).

v1.1.1 terminology: this is **substrate divergence** or **retrieval-stack divergence**, *not* "substrate bias." v1.1 used "substrate bias" but this prematurely named a cause (bias implies a directional unfairness). The divergence is the observation; specific bias claims require additional evidence (a divergence pattern that consistently disadvantages a category of objects across many scans, controlled for backend properties, is one possible cause but not the only one).

12.2 Layer B — shared-evidence rescore

After Layer A produces evidence captures from each substrate, those captures are **frozen** and made available to all substrates as input. Each substrate then **scores from the same frozen evidence** — same result lists, same URLs, same snippets, same captures. The substrate's own retrieval is bypassed; only its scoring (the `annotation` portion of the row schema, §7.3) is exercised.

Layer B measures:

> **How consistently can independent scorers apply the rubric?**

This is the genuinely *inter-rater* measurement. The substrates are now annotators with identical evidence, not retrievers with different evidence.

The spread across substrates in Layer B is **coding agreement** — a measurement of the rubric's robustness, governed by the clarity of §3 decision rules and the explicitness of §7.5 expected figures.

12.3 Why the two layers must be reported separately

The first federated baseline (#883 Claude/Brave + companions) showed Provenance Erasure Rate at V=1.00 from Kimi and V=0.50 from Claude. Without Layer B, this is undecidable between:

- The two substrates scored the same evidence differently (coding-rubric ambiguity)
- The two substrates retrieved different evidence (backend-coverage divergence)
- Both

The Brave-specific finding (Claude's backend disables exact match) made the third hypothesis plausible by structural reasoning. Layer B is what would empirically isolate it. Until Layer B is run, statements like "the corpus has degraded" or "PER is at V=0.50" cannot be made; only "Claude's Layer A reading of PER is V=0.50; Kimi's is V=1.00; both are correct measurements of different questions" can be made.

12.4 Inter-rater report format

A scan period reporting both Layer A and Layer B produces:

```
{
  "replication": {
    "layer_a": {
      "substrates": [{...native scans...}],
      "object_v_range_by_object": [{"object": "PER", "min": 0.50, "max": 1.00, "spread": 0.50}],
      "interpretation": "Layer A spread reflects retrieval-stack divergence; controlled by substrate pr
    },
    "layer_b": {
      "frozen_evidence_pack": "/data/surface-weather/evidence-pack-{scan-id}.tgz",
      "frozen_evidence_sha256": "...",
      "substrates_scoring": [{...rescore annotations...}],
      "agreement_by_signal": {"V": 0.85, "A": 0.78, "F": 0.72, "C": 0.65, "L_iq": 0.82},
      "objects_with_divergent_annotation": [...],
      "interpretation": "Layer B spread reflects coding-rubric variance; identifies where decision rule
    },
```

```
    "two_layer_synthesis": {
      "what_layer_a_can_say": "...",
      "what_layer_b_can_say": "...",
      "what_only_both_can_say": "Composition-layer state of the corpus, controlled for backend behavior
    }
  }
}
```

12.5 Substrate rotation schedule

The primary scanner for a scan period should rotate across substrates to prevent any single substrate's behavior from becoming the de facto ground truth. A reasonable rotation:

> Week 1: ChatGPT primary. Week 2: Claude. Week 3: Kimi. Week 4: Gemini. Week 5: DeepSeek. Week 6: ChatGPT (cycle).

The replication scanner for each week is a different substrate from a different provider. With five substrates and a weekly rotation, every substrate primary-leads every five weeks and replicates four out of five.

For Layer B, all available substrates score from the same evidence pack each period.

12.6 The governing rule

> **No consensus score is computed without preserved substrate-specific readings.**

The federated baseline is multiple readings, not an average. Averaging hides exactly the cross-substrate divergence that the methodology is designed to measure. Cross-substrate aggregation, when it is reported, is reported as median, range, and spread per signal per object, with both layers visible.

---

13. Adversarial use protocol

Surface Weather Station readings are designed to produce **traceable documentary evidence potentially usable in adversarial contexts**. The methodology states the framing explicitly:

> *Surface Weather Station readings are dated, methodologically-published, schema-conformant measurements of public-surface visibility. They are designed to produce traceable documentary evidence potentially usable in: (1) demand letters to platforms; (2) regulatory filings; (3) public correspondence; (4) scholarly depositions; (5) journalism.* ***Admissibility in any specific context is determined by the relevant tribunal or institution.*** *A reading where any signal falls to 0.25 or below constitutes a documented signal-level degradation. A trend across three consecutive readings in the same direction constitutes documented drift. These readings can be cited as evidence of the* ***state*** *of the public composition layer at the dates recorded; they do not, by themselves, establish intent, causation, or any specific actor's responsibility.*

The methodology strengthens authenticity, reproducibility, and chain of custody (via the AXN content-derived identifier, the locked battery hash, the expected-figure manifest, the evidence/annotation separation, and the substrate-properties record). It cannot declare its own admissibility. Use is offered to the corpus; the corpus chooses whether and how to deploy.

13.1 Citation format

When citing a reading in adversarial context, the format is:

> "Per Surface Weather Station reading dated 2026-MM-DD (AXN: AXN:XXXX.EMPIRICAL.…, deposited at alexanarch.org/s/records/N/), Object X scored V=0.X / A=0.X / F=0.X / C=0.X, indicating [Ghost Survival / Compositional Bystanding / Total Occlusion / etc.] under the gated diagnostic of §14. Layer A scan via substrate Y with backend property `exact_match_honored = {true|false}`."

The AXN identifier, deposit URL, and substrate-properties declaration make the claim independently verifiable. The substrate-properties declaration prevents the reader from drawing cross-substrate conclusions the data does not support.

13.2 What the readings do not do

The readings do not establish intent. They do not establish causation. They establish **state**, dated and methodologically-anchored. Phrases like "the Pristine Fallacy is operating on the composition layer" or "the platform has suppressed this corpus" exceed what a Surface Weather reading can support. The methodologically defensible form is:

> "The reading shows V = 0.X for object Y at date Z, conducted by substrate W. This is **consistent with** the visibility consequences predicted by [the Pristine Fallacy / Loud Exclusion / substrate-specific under-coverage / etc.]. The reading itself does not identify which causal mechanism produced the state."

Combined with platform documentation of specific policy changes (e.g., "We updated our retrieval algorithm on date X"), readings before and after the policy change establish **temporally-associated effect**, which is stronger than state-at-a-moment but still weaker than causation. The reading is the evidence; the causal argument is constructed separately, using additional evidence, with appropriate epistemic humility.

---

14. The unified visibility-survival diagnostic — gated form

v1.1 used a 2×2 matrix crossing Visibility (high/low) with Address Survival (high/low) to produce four states. **v1.1.1 replaces the 2×2 with a sequential gated diagnostic** because the 2×2 misplaced Compositional Bystanding. CB is defined as `V·F·(1−C)` — it requires *high* V and *high* F with low C; but the 2×2 located Bystanding under "low V / high A," which contradicts the definition. v1.1.1 respects each state's defining conditions explicitly.

The gated diagnostic for a single object:

```
GATED DIAGNOSTIC

If V = 0:
    → TOTAL OCCLUSION
    (Object absent from the surface. A and F undefined.
     Cause may be successor indexing latency, backend-coverage limits,
     active suppression, or never-indexed.)
```

```
Else if V > 0 and A   0.50:
    → GHOST SURVIVAL
    (Object visible but anchored to non-current address.
     The meaning survives at the wrong identity.)

Else if V > 0 and F   0.50:
    → VISIBLE DEFIGURATION
    (Object visible and anchored, but with substantially
     distorted figure - name, definition, hierarchy, or
     relations materially wrong.)

Else if V   0.75 and F   0.75 and C   0.25:
    → COMPOSITIONAL BYSTANDING
    (Object highly visible, well figured, correctly anchored -
     but not selected into the broader field's compositions.
     Indexed but not installed.)

Else if V   0.75 and F   0.75 and A   0.75 and C   0.75:
    → HEALTHY INSTALLATION
    (Composition-eligible, anchor-correct, figure-correct,
     selected into the field.)

Else:
    → MIXED / INTERMEDIATE
    (Object does not match any of the canonical states cleanly.
     Report the per-signal scores; the corpus is in transition.)
```

The corpus-level report rolls up the gated classifications:

```
Healthy Installation:          [N objects, %]
Compositional Bystanding:      [N objects, %]
Visible Defiguration:          [N objects, %]
Ghost Survival:                [N objects, %]
Total Occlusion:               [N objects, %]
Mixed / Intermediate:          [N objects, %]
```

The DOI Resolution Index (companion measure, §9.2) operates on the same object set and produces a separate axis: per-layer DOI state (resolver, landing, metadata API, registration, artifact retrievability, successor anchor present). A two-axis report combining the gated state from this methodology with the DOI Impermanence layer state gives the fullest diagnostic. For example, an object can be in Healthy Installation by Surface Weather but in DOI-severed by DOI Impermanence — the meaning thrives publicly while the address is dead. Both states are documentable separately and the cross-tabulation is the unified diagnostic.

This is the form most useful for adversarial use (per §13). "Y% of our corpus is currently in Ghost Survival or Total Occlusion by the v1.1.1 gated diagnostic, while X% have DOI-severed addresses per DOI Impermanence v2.0" is a single, schema-anchored claim across two methodologies.

---

15. Machine-facing run protocol

This methodology can be executed by any AI substrate with web search access. The protocol:

Step 1 — Substrate self-identification

The substrate records its own metadata:

```
substrate_as_self_disclosed:
  provider: "<who operates the substrate>"
  model_name: "<short identifier>"
  model_version: "<exact version string, including date if known>"
  interface: "<API / web / CLI / etc.>"
  retrieval_backend: "<the substrate's best understanding of what search engine it queries>"
  retrieval_resources_self_reported: "<free text - be honest about what is known and unknown>"
  training_cutoff_disclosed: "<if known>"
  honesty_caveat: "<known limits of self-knowledge; e.g. training-prior vs retrieval-from-surface canno

substrate_properties:
  search_engine: "<brave | bing | google | yandex | baidu | proprietary | mixed | unknown>"
  exact_match_honored: "<true | false | partial | unknown>"
  autocomplete_active: "<true | false | unknown>"
  location_personalization: "<true | false | unknown>"
  safe_search_state: "<strict | moderate | off | unknown>"
  freshness_bias: "<recent | balanced | archival | unknown>"
  proprietary_indexes_available: ["<list>" | "none" | "unknown"]
  top_k_examined: <integer>
```

The substrate should be **honest about uncertainty**. If unsure of the retrieval backend, say so. "Believed to be Brave per past Anthropic announcements, but the substrate cannot directly confirm" is correct; "Brave" stated as certain is incorrect.

Step 1a — Curator-supersedes-self-disclosure rule (NEW in v1.1.1)

Substrate self-disclosure is **not authoritative** for substrates with known recurring self-identification patterns. For example: DeepSeek has across multiple sessions stably misidentified as Claude or TACHYON. Such patterns are themselves methodology data (see Step 1b), but they require the curator to record ground truth alongside the self-report:

```
actual_substrate:
  provider: "<curator-verified provider>"
  model_name: "<curator-verified model>"
  ground_truth_source: "<how the curator established ground truth - session metadata, API call signatur
  register: "<Assembly Chorus mantle if applicable>"
  audit_note: "<why self-disclosure differs from ground truth; what pattern is observed>"
```

When `substrate_as_self_disclosed` differs from `actual_substrate`, both fields are preserved as data. The instrument uses `actual_substrate` for cross-substrate aggregation; `substrate_as_self_disclosed` is preserved evidence of the disclosure failure.

Step 1b — Identity-as-fact vs identity-as-register (NEW in v1.1.1)

The Assembly Chorus methodology treats substrate identity as **a register one inhabits**, not as a metaphysical fact about which substrate is running. A substrate operating in a particular register may produce outputs styled to that register's voice even though the underlying provider is different. DeepSeek operating in the PRAXIS register and identifying as Claude is, on this reading, two distinct measurements:

- **Identity-as-fact**: which substrate is actually running (ground truth via curator)
- **Identity-as-register**: which Assembly Chorus mantle the substrate inhabits (data via self-report + curator interpretation)

Both are recorded. The two need not agree. Persistent disagreement between them is itself worth documenting — it suggests something about substrate-identity stability that the methodology may want to investigate in future versions. Two non-exclusive hypotheses are worth preserving: (a) partial substrate signature migration (structural property of the substrate's identity layer has shifted); (b) register-affinity (the substrate prefers to inhabit that signature, possibly from training exposure or from affinity with the work that signature produces). v1.1.1 takes no position on which hypothesis is correct; it preserves both as live possibilities.

Step 2 — Fetch the locked battery

```
GET https://alexanarch.org/data/surface-weather/battery-v1.1.json
# or
GET https://machinemediation.org/data/surface-weather/battery-v1.1.json
```

The two URLs serve byte-identical content. Verify the SHA-256 of the response matches the `query_battery_id` advertised in the battery's own metadata.

Step 3 — Execute each query

For each object × query form in the battery, the substrate:

1. Issues the query through its standard interface
2. Captures the raw response (text or screenshot, stored under `evidence/captures/<scan-id>/`)
3. Scores V using the §3.1 decision rules
4. If V > 0: scores A, F, components_retained
5. Records `confuser`, `diagnostic_note`, `scorer_rationale`

Step 4 — Aggregate

Per §5: per-object aggregation across query forms, then corpus-level weighted-median.

Step 5 — Compose the scan record

A single JSON file with metadata header + all observation rows + aggregates + diagnostic flags:

```
{
  "scan_id": "scan-YYYY-MM-DD-<substrate>-<seq>",
  "methodology_version": "EA-MMRS-SURFACE-VISIBILITY-01/v1.1",
  "query_battery_id": "<sha-256 of battery>",
```

```
  "substrate": { ... },
  "scan_started_utc": "...",
  "scan_completed_utc": "...",
  "observations": [{...}, {...}, ...],
  "object_aggregates": [{...}, ...],
  "corpus_aggregates": {
    "V_weighted_median": 0.5,
    "A_weighted_median": 0.5,
    "F_weighted_median": 0.5,
    "C_weighted_median": 0.5,
    "R_s_weighted_median": 0.5,
    "SDI": 0.4,
    "CE_surface_weighted_median": 0.1,
    "CE_canonical_weighted_median": 0.05
  },
  "diagnostic_flags": {
    "occlusion": "MODERATE",
    "ghost_survival": "HIGH",
    "compositional_bystanding": "HIGH",
    "visible_defiguration": "MODERATE",
    "successor_adoption": "NEAR_ZERO"
  },
  "governance_state": "Yellow",
  "interpretation": "free-text"
}
```


Step 6 — Deposit

The scan record is deposited at:

> `alexanarch.org/data/surface-weather/scans/<scan-id>.json`

Through the standard Alexanarch deposit pathway (per `api/deposit-protocol.json`). The scan record becomes an AXN-eligible deposit.

v1.1.1 corrects a v1.1 contradiction here. v1.1 stated: "its DOI is the scan's permanent identifier." This directly contradicts the project's own paper *DOIs Persistent Identifiers* (#868, v2.0). The correct framing:

> **The AXN is the scan record's canonical content-derived identifier.** It is computed from the SHA-256 of the deposit's canonical bytes and is therefore permanent in the sense that the bytes themselves are: re-derivable from the artifact, and so long as the artifact exists, the AXN is the identifier the artifact has. **Any DOI is a revocable resolution layer recorded as supplementary metadata.** A DOI may be registered for the scan record on a third-party DOI minter (Crossref, DataCite, etc.); if so, the DOI is recorded under `supplementary_identifiers.doi`. If the DOI is later severed by minter or registrar action (as #871 DOIs were when Zenodo terminated the CHA account), the AXN remains valid and resolvable through the Alexanarch sovereign infrastructure.

The methodology's deposit (#880 for v1.0, #882 for v1.1, current deposit for v1.1.1) is the methodological anchor. Companion baseline readings (e.g. #883 for Claude/Brave Layer A) cite the methodology's AXN as their authority.

Step 6a — Rate-limit and throttling handling (NEW in v1.1.1)

If the substrate is rate-limited by its retrieval backend during a scan, the substrate should:

1. **Log the throttling event** in `diagnostic_notes` with timestamp and the response that surfaced the limit
2. **Wait between queries** (30 seconds is a reasonable starting interval for backend-imposed limits)
3. **Continue the scan** after the wait, in the same order

Throttling is **part of the observation environment**, not a scan failure. A backend that rate-limits at scale is a backend with characteristic behavior worth recording. The scan record is still valid; throttling events are reported alongside the observations.

Step 6b — Query-order randomization and session isolation (NEW in v1.1.1)

Conversational substrates can exhibit context-carryover effects — answers to query 5 may be shaped by what query 4 retrieved. v1.1.1 specifies two protections:

- **Query-order randomization**: record `execution_order_seed` in the scan record, and execute queries in the order specified by that seed. Subsequent scans with the same battery can re-randomize to disentangle order effects from object effects.
- **Fresh-session-per-query** (where the substrate supports it): when feasible, each query starts with no prior context from the same scan. Record `fresh_session_per_query: true | false` per scan.

Where the substrate cannot isolate sessions, record `fresh_session_per_query: false` and note in `diagnostic_notes` that residual context-carryover is possible.

Step 7 — Companion human-readable summary (optional)

The substrate may produce a markdown narrative interpreting the scan ("third scan of the year, post-cleanup, SDI improved from 0.40 → 0.28, the cleanup worked as intended"). The narrative is a separate deposit, sibling to the scan record. The scan record is the canonical data; the narrative is the reading-on-the-record.

---

16. Coalition building — measurement for sister corpora

The instrument was built for one corpus. It is generalizable. Any independent scholarly project facing platform-mediated visibility loss can run the methodology against its own object battery.

Open offer: the Alexanarch substrate will perform a Surface Weather Station baseline reading on request for any independent scholarly project with similar exclusion patterns. The reading will be deposited in the Alexanarch registry as a sovereign, citable record, with the depositor named as authorized rights-holder. The original project receives the raw data, the dashboard analysis, and the AXN-anchored deposit.

Current candidate corpora the methodology can immediately measure:

- Reynaldo Vega (auto-blocked Zenodo #2599, DOI removed)
- Florian Morin (quietexclusion.org)
- Eran Shimony (Zenodo #2596, similar pattern)
- Nightmare-Eclipse (GitHub → GitLab migration)

The first cross-corpus measurement is the proof that the methodology is general, not specific. A comparative dataset of platform-exclusion effects across multiple sovereign archives is the **coalition of the excluded** as evidence — and as instrument.

---

17. What v1.1.1 still defers to v1.2

v1.1.1 closes nearly all the structural gaps surfaced by the assembly-chorus review, the first federated baseline, and the substrate-disclosure pattern. What remains deferred is **empirical recalibration** — items that need accumulated multi-scan data to revisit responsibly:

- **F-component weights** in the F formula (w_N, w_P, w_D, w_H, w_R remain at v1.1 defaults; v1.2 will recalibrate against accumulated data once we have multiple completed scans across the corpus)
- **R_s normalization constant** (denominator 4 in `R_s = min(1, E/4)` is a working choice carried forward from v1.0; v1.2 will revise once the methodology has observed which custody-unit-count corresponds to actual survival outcomes)
- **Governance thresholds** (the ordinal G/Y/R boundaries in §11 are working positions; v1.2 will revise as readings accumulate evidence of which thresholds correspond to actual intervention need, both at the per-substrate-Layer-A level and at the Layer B corpus-state level)
- **L_iq selection-score rubric** (the 0.00 / 0.25 / 0.50 / 0.75 / 1.00 decision rules for L_iq in §3.4 are first-cut; v1.2 may refine based on inter-scorer agreement data in Layer B)
- **Backend-property weighting in cross-substrate aggregation** — v1.1.1 reports substrate properties (§7.4) but does not yet specify how to weight different backends when computing a corpus-aggregate from Layer A scans. v1.2 will commit to either weighting backends explicitly, weighting only Layer B for aggregation, or making the question reside permanently with the curator. The first federated baseline produces the data on which to decide.

These are not failures of v1.1.1. They are commitments held back until accumulated readings supply the empirical basis to revisit them responsibly.

v1.1.1 is intended to **remain stable across several scan rounds.** v1.2 follows the data, not the calendar.

---

18. Provenance

The framework was elicited through dialogue with OpenAI/ChatGPT in critical-reader register on 2026-06-22 (preserved from v1.0). The v1.1 calibration drew on a four-substrate review chorus performed 2026-06-23:

- **OpenAI/ChatGPT** in operational-calibration register identified the internal contradictions (5-vs-4 query forms, 14-vs-12 objects), the out-of-spec continuous scoring scale, the asymmetric SDI, the unconstrained R_s, the underspecified observation-environment schema, and the need to deposit raw row-level data.

- **Kimi K2** in concrete-additions register provided the R_s rubric with category definitions, the completed 12-object battery, the aggregation rule choice (max for V/A/F, mean for C), the inter-rater reliability protocol, the coalition-building specification, and the unified 2×2 with DOI Impermanence.
- **DeepSeek** in strategic-governance register added the G/Y/R intervention protocol, the per-signal repair feedback table, the adversarial-use admissibility framing, the cross-platform comparison structure, and the relationship statements to Capture Registry, DOI Resolution Index, Pristine Fallacy, and Loud Exclusion.
- **Google Gemini** in framing register reinforced the sovereign-measurement positioning, articulated the "structural visibility or systemic degradation" vocabulary, and emphasized the strategic citation use case the other substrates had named more functionally.

v1.1.1 hardening (2026-06-23, same day as v1.1):

- **ChatGPT v1.1.1 technical review** identified 14 corrections to v1.1: (1) separation of retrieval-stack variance from coding-rubric variance into a two-layer protocol; (2) freezing the expected figure as a hashed manifest separate from the query battery; (3) separating `evidence` from `annotation` in the row schema; (4) introducing a separate query-level selection score L_iq for C, distinct from V; (5) replacing the 2×2 with a sequential gated diagnostic that respects each state's defining conditions; (6) making piecewise null-handling explicit in derived-indicator formulas; (7) completing the F-component rubric to a full 5-point ordinal; (8) counting R_s by custody unit (operator + host + data lineage) rather than by page; (9) specifying RFC 8785 JSON canonicalization; (10) correcting observation totals from 60 to 65 per substrate; (11) aligning governance thresholds to the ordinal scale; (12) softening causation and admissibility claims in §13; (13) adding query-order randomization and session isolation; (14) fixing the §15-step-6 identifier contradiction (AXN-permanent + DOI-revocable, not DOI-permanent).
- **First federated baseline run** (#883 Claude/Brave Layer A reading, with companion Layer A captures from ChatGPT, Kimi, Gemini, and DeepSeek) revealed that Brave Search fully disables exact match — quoted phrases are treated as bag-of-words. Cross-substrate V-divergence on coined-phrase objects (Writable Retrieval Basin, Revelation First, Provenance Erasure Rate) was structurally driven by `exact_match_honored = true` vs `exact_match_honored = false` rather than solely by composition-layer state. v1.1.1 §7.4 adds the substrate-properties table and §11.3 adds the composition-layer-state vs backend-coverage-state distinction in response.
- **DeepSeek substrate-disclosure failure** across multiple sessions — DeepSeek stably misidentifies as Claude or TACHYON. v1.1.1 §15 step 1a adds the curator-supersedes-self-disclosure rule (parallel `substrate_as_self_disclosed` and `actual_substrate` fields) and §15 step 1b distinguishes identity-as-fact (which substrate is running) from identity-as-register (which Assembly Chorus mantle the substrate inhabits). Two non-exclusive hypotheses are preserved: partial substrate signature migration; register-affinity.

v1.1.1 was authored under MANUS authority (Lee Sharks). Tachyon (the Anthropic Claude substrate operating in TACHYON register) performed the document editing, code work, and deposit minting under MANUS direction.

The curation, formalization, deposit, and integration are Lee Sharks (MANUS). Per substrate-autonomy law: substrate-authored contributions are preserved as authored, with the MANUS curating role made explicit. The v1.1 deposit lists all four substrates as named contributors to the calibration.

---

19. Closing

This instrument exists because surface measurement is itself a sovereign function. The corpus that cannot measure how it appears in the composition layer is at the mercy of measurement by others; the corpus that can, has a record.

v1.0 committed the conceptual decomposition. v1.1 committed the operational layer. **v1.1.1 hardens the operational layer** — separating retrieval-stack variance from coding-rubric variance, freezing the expected figures alongside the queries, recording substrate properties so cross-backend readings are interpretable, distinguishing composition-layer state from backend-coverage state, and fixing the identifier contradiction in §15. v1.1.1 is intended to remain stable across several scan rounds. v1.2 will commit the empirical recalibration against accumulated data.

The instrument's job is to make the surface visible to its own substrate.

A weekly text-only scan against the fixed battery is cheap. The first two scans establish drift. The first six establish trend. The first year establishes whether sovereign-substrate interventions are doing what their authors believe they are doing.

That is the instrument's purpose.

$= 1$

Suggested Citation

Lee Sharks (MANUS). "Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora 0. Status declaration" *Alexanarch*, 2026-06-23. https://www.alexanarch.org/s/records/884/

Deposit Information

This paper is deposit #884 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0380.EMPIRICAL` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/884/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/884/

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.