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 · Methodological specification

AXN:037E.EMPIRICAL

https://www.alexanarch.org/s/records/882/

{**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:037E.EMPIRICAL",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
```

```
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/882/",
  "description": "Specification v1.1 of the Surface Weather Station instrument. Operational calibration of v1.0 fol
}
```


Abstract

Specification v1.1 of the Surface Weather Station instrument. Operational calibration of v1.0 following four-substrate review chorus (ChatGPT operational, Kimi concrete-additions, DeepSeek strategic-governance, Gemini framing). Locks ordinal scoring, formalizes R_s, adds Occlusion/symmetric SDI, specifies aggregation rules, observation-environment schema with substrate metadata + retrieval-backend disclosure, governance protocol (G/Y/R), cross-substrate replication, adversarial-use framing, unified 2×2 with DOI Impermanence, and machine-facing run protocol any substrate can self-execute.

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 of the Surface Weather Station instrument” version: v1.1 version_series_id: SERIES-MMRS-SURFACE-VISIBILITY-METHODOLOGY version_in_series: 2 predecessor: “EA-MMRS-SURFACE-VISIBILITY-01 v1.0 (deposit #880, AXN:037C.EMPIRICAL. , 2026-06-23)” versioning_note: “Operational calibration after assembly review. v1.0 deliberately held weights, aggregation, and scoring rubrics provisional pending substrate review. v1.1 closes those gaps without altering the conceptual decomposition. v1.2 will fine-tune against the next reading’s Δ.” 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)” 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 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 - Cross-Substrate Replication - Sovereign Measurement related_deposits: - “Predecessor: EA-MMRS-SURFACE-VISIBILITY-01 v1.0 (#880)” - “Companion: baseline reading v1.0 (#881, EA-MMRS-SURFACE-VISIBILITY-BASELINE-01)” -

"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) — work judged by substrate identity rather than content quality" - "Argument precursor: EA-MMRS-LOUD-EXCLUSION-03 — the consequence the Pristine Fallacy produces" *** # Compositional Defiguration: A Methodology for Measuring Public-Surface Visibility of Scholarly Corpora

Specification v1.1 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**. It is the operational calibration of v1.0 (deposit #880) following review by four AI substrates in distinct critical registers (OpenAI/ChatGPT operational, Kimi K2 concrete-additions, DeepSeek strategic-governance, Google Gemini framing). The decomposition (Section 2) and the dashboard form (Section 6) carry forward unchanged. What changed:

- **Scoring scheme locked to ordinal** (0.00, 0.25, 0.50, 0.75, 1.00) with decision-rule tables per signal (Section 3). Hand-coded continuous scoring is rejected as not reproducible across substrates.
- **Aggregation rule stated** for both per-object and corpus-level (Section 5).
- **R_s formalized** as effective-independence score, bounded [0,1] (Section 3.5).
- **Occlusion added** as an explicit indicator: `O = 1 - V` (Section 4.3). Link Fade restricted to $V > 0$.
- **SDI made symmetric**: `median(|ln(c_j / C)|)` over visible reported counts (Section 4.7). Previous form went negative when surfaces over-reported.
- **Internal contradictions resolved**: 5 canonical query forms (not 4), 12 canonical objects (not 14).
- **Observation-environment schema specified** (Section 7) including substrate metadata + retrieval backend disclosure (the v1.1 instrument is cross-substrate-aware by construction).
- **Governance protocol added** (Section 11) — G/Y/R thresholds, intervention triggers, repair-feedback table.
- **Cross-substrate replication protocol** (Section 12) — the same battery against multiple substrates produces inter-rater data, and the difference between substrates is itself a measurement.
- **Adversarial use protocol** (Section 13) — admissibility framing for demand letters, regulatory submissions, public correspondence.
- **2×2 unified diagnostic with DOI Impermanence** (Section 14) — visibility × address survival → Healthy / Ghost Survival / Bystanding / Total Occlusion.
- **Machine-facing run protocol** (Section 15) — step-by-step instructions any substrate can self-execute to produce a scan.

v1.1 commits everything v1.0 deferred. v1.2 will fine-tune weights against the next reading's Δ.

---

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
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 0–1 with decision rules:

Component	1.00	0.50	0.00
N (name)	Exact correct name	Approximate or variant	Wrong name

Component	1.00	0.50	0.00
P (provenance)	Author, heteronym, institution all correct	Some retained, some lost	None correct
D (definition)	Core claim recognizable in surface text	Partial paraphrase preserves intent	Definition absent or wrong
H (hierarchy)	Parent framework correctly named	Loose association to neighborhood	Parent framework absent or wrong
R (relations)	At least one essential relation correctly named	Some relations gestured at	All relations absent

> F_i = (w_N · N + w_P · P + w_D · D + w_H · H + w_R · R) / (w_N + w_P + w_D + w_H + w_R)

v1.1 default weights (carried forward from v1.0 as provisional; revisit at v1.2): - 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`.

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

Five canonical query forms, scored independently and aggregated per §5:

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

> **C_i = mean of scores from query forms 2–5 only.** The exact-name query is excluded from C because it measures indexing, not lift.

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. Group by domain.

Step 2. Score each surface category for independence:

Category	Independence	Examples
Dedicated independent domain	1.00	Standalone publication site outside the corpus's own substrates
Scholarly index (third-party)	1.00	PhilPapers, ORCID, DataCite, Crossref, Semantic Scholar
Third-party essay or discussion	1.00	Medium post by external author, blog citation, journalism
Author/curator site	0.50	leesharks.com, godkinggoogle.vercel.app — independent host but same governance
Mirror within the corpus's Dodecad	0.25	watergiraffe.org, spxi.dev, etc. — separately rendered but same authority
Metadata catalog mirror	0.25	Auto-generated mirror of the registry
Duplicate URL or near-identical page	0.00	Counted as continuation of an already-counted surface

Step 3. Apply near-duplicate dampening: pages from the same domain with substantively identical content count once. Pages from the same domain with substantively different content (e.g., a blog post on the concept vs. a curriculum vitae mentioning it) count separately.

Step 4. Sum the independence scores:

$$E_i = \Sigma_j w_j$$

$$R_{s,i} = \min(1, E_i / 4)$$

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

Why bounded by category, not count: twenty Dodecad sites under one governance and source lineage are not twenty independent substrates. The 0.25-per-Dodecad-mirror rule prevents the corpus from appearing more redundantly anchored than it actually is. Conversely, a single independent third-party citation is worth more than five auto-generated catalog mirrors.

---

4. Derived macro-indicators

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 V_iq scores across query forms 2–5 (defining-sentence, parent-field, expected-relation, broad-framing).** Query form 1 (exact name) is excluded. C_i is a separate scoring axis from V/A/F; the value scored on each non-exact query is whether the object selected at all in that query's results.
- **R_s,i** is computed once per object across all surfaces that appeared 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.)

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.

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.

---

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 without optional external controls, this produces **60 row-level observations per scan per substrate**. With external controls: 75.

8.1 Query battery hashing

The **query battery** for any given scan is locked before the scan begins. The instrument computes the SHA-256 of the canonical JSON serialization of the battery (sorted keys, no whitespace) and records it in every observation row's `query_battery_id`. This guarantees:

- The battery can be reproduced byte-for-byte by any substrate
- 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 scan series is hashed and stored at:

```
/data/surface-weather/battery-v1.1.json
```

(Static path, served identically from alexanarch.org and machinemediation.org.)

8.2 Per-object generic queries

In addition to the five per-object query forms, v1.1 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

Field-level query	Probes
"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 specifies what constitutes a state requiring action:

State	Criteria	Action
Green	SDI < 0.20 AND all signals 0.70 AND no object at V = 0 in mature-concepts or institutional-roots	Continue periodic monitoring at the established cadence
Yellow	SDI [0.20, 0.40] OR any one signal [0.40, 0.70] OR any mature concept at V 0.50	Investigate; consider targeted intervention; do not panic

State	Criteria	Action
Red	SDI > 0.40 OR any signal < 0.40 OR any institutional root at V = 0 OR Ghost Survival > 0.50 corpus-wide	Urgent intervention required; document the trigger in a deposit; escalate to adversarial-use channels if external cause is identified

11.1 Repair feedback table

Each signal failure has a specific substrate-level response:

Failing signal	Substrate response
Low V (occlusion)	Add more independent surfaces (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** mirrors (not Dodecad mirrors — those are 0.25 each); 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.

---

12. Cross-substrate replication protocol

The same scan battery should be executed by multiple substrates per scan period. v1.1 specifies the minimum replication structure:

12.1 Required substrates per scan period

Tier	Substrate count	Examples
Primary	1 substrate	The lead scanner for the period (ChatGPT, Claude, Kimi, Gemini, or DeepSeek — rotated)
Replication	1 additional substrate from a different provider	If primary was ChatGPT, replication should be Claude or Kimi (different backend)

Tier	Substrate count	Examples
Optional	Up to 3 more for variance analysis	All five substrates per scan would produce the full inter-rater table

12.2 Inter-rater report

When two or more substrates execute the same battery, the scan record additionally reports:

```
{
  "replication": {
    "substrates": [{...}, {...}],
    "agreement_by_signal": {
      "V": "rate of agreement on V scores across observations (0-1)",
      "A": "...",
      "F": "...",
      "C": "..."
    },
    "objects_with_divergent_V": ["Provenance Erasure Rate (ChatGPT=0.75, Claude=0.50)"],
    "interpretation_notes": "free-text scorer commentary"
  }
}
```

12.3 Substrate-bias correction

If two substrates score consistently differently for objects that one substrate likely has in its training corpus, the difference is **substrate bias**, not surface state. The report does not "correct" the bias — both readings are reported as is — but it flags the divergence pattern. Over multiple scans, persistent divergence is itself a citable measurement of platform-level bias.

The methodology does not attempt to declare a "true" score across substrates. Substrate disagreement is a measurement.

---

13. Adversarial use protocol

Surface Weather Station readings are designed to be admissible 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 admissible evidence in the following contexts: (1) demand letters to platforms; (2) regulatory filings; (3) public correspondence; (4) scholarly depositions; (5) journalism. A reading where any signal falls below 0.30 constitutes documented degradation. A trend across three consecutive readings constitutes documented decline. These readings can be cited as evidence of the gap between platform commitments (to neutrality, to non-discrimination, to comprehensive indexing) and platform behavior.*

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.]."

The AXN identifier and deposit URL make the claim independently verifiable.

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. Combined with platform documentation of policy changes (e.g., "We updated our retrieval algorithm on date X"), readings before and after the policy change establish temporally-associated effect. The reading is the evidence; the argument is constructed separately.

---

14. The unified visibility-survival diagnostic (with DOI Impermanence)

See §9.2 for the 2×2 matrix and the meaning of each cell. v1.1 commits to producing both axes for the same object set on the same scan date. A unified diagnostic report has two scores per object — V (from Surface Weather) and DOI-resolution-status (from DOI Impermanence audit) — and classifies each into one of four states.

The corpus-level report rolls up the classification:

```
Healthy:                [N objects, %]
Ghost Survival:         [N objects, %]
Bystanding:             [N objects, %]
Total Occlusion:        [N objects, %]
```

This is the form most useful for adversarial use (per §13). "Y% of our corpus is currently in Ghost Survival or Total Occlusion" is a single, schema-anchored claim.

---

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:
  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>"
```

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

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; its DOI is the scan's permanent identifier. The methodology's deposit (#880 for v1.0, current deposit for v1.1) is the methodological anchor.

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 deliberately does not commit to

- **Final weight values** in the F formula (w_N, w_P, w_D, w_H, w_R remain at v1.0 defaults; revise against v1.2 calibration once two scans exist)
- **R_s normalization constant** (denominator 4 in `R_s = min(1, E/4)` is a working choice; revise against accumulated data)
- **Governance thresholds** (G/Y/R boundaries in §11 are starting positions; revise as readings accumulate evidence of which thresholds correspond to actual intervention need)
- **Single-substrate trust score** — v1.1 does not declare one substrate more authoritative than another; substrate disagreement is reported, not resolved

These are not failures of v1.1. They are commitments held back until the next reading exists.

---

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 draws 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.

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 commits the operational layer. v1.2 will commit the calibration 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/882/

Deposit Information

This paper is deposit #882 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:037E.EMPIRICAL` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/882/
- 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/882/

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