

The Structural Spectrometer:

A Formal Specification and Runnable Handoff for Measuring the Dimension of Cross-Domain Method Space

Benjamin Dakota Horn
Prometheus7 Research Institute

Handoff document for continuation by an autonomous coding agent

June 5, 2026

Abstract

This document specifies, in full and with all equations, a working instrument for measuring whether the methods of science across all fields collapse onto a small set of shared *structural prerequisites*. The central object is the *failure-signature*: a method’s pattern of survival-versus-dissolution against a battery of maximally-structureless reference problems (“walls”), each of which denies one structural axis. The count of distinct failure-signatures over a corpus is the *measured dimension of method space*; the program’s falsifiable core is whether this count plateaus (a small basis exists – evidence for a unified relational field) or grows without bound (no small basis – the instrument is a thesaurus). A reference implementation exists and is described file-by-file; this paper provides the formal framework so that an autonomous coding agent can continue the program from the current state. We give the operator algebra, the wall-denial semantics, the signature and dimension definitions, the composition calculus with viability filtering, the calibration conditions (negative control, saturation), and a prioritized list of the remaining build tasks with their exact mechanics. The decisive next experiment is specified precisely: run the saturation measurement on a real heterogeneous corpus (e.g. an arXiv dump) and determine whether the dimension plateaus.

1 Purpose and reading guide for the continuing agent

This is a handoff. A reference implementation already runs (Section 7); the formal content below (Sections 2–6) defines what the code computes, and Section 8 lists what remains, in priority order, with mechanics precise enough to implement directly. The agent should: (i) read Section 7 to map equations to files; (ii) treat Section 6 as acceptance tests that must keep passing after any change; (iii) execute Section 8 in order; (iv) treat all extracted operators as *candidates with provenance*, never as discoveries, per Section 9.

2 Objects

2.1 Structural primitives

Let Π be a finite controlled vocabulary of *structural primitives*: substrate properties a method may require to function. The reference set ($|\Pi| = 12$):

$$\Pi = \{\text{continuity, linearity, homomorphism_to_easy, sparsity,} \\ \text{spectral_support, analytic_continuation, correlation, smoothness,} \\ \text{metric, low_rank, monotonicity, compactness}\}.$$

2.2 Operators

A *bare operator* extracted from a concept or paper is a typed object

$$o = \langle \tau_{\text{in}}, \tau_{\text{out}}, \text{op}, R(o), c(o), \text{prov}(o) \rangle,$$

where $\tau_{\text{in}}, \tau_{\text{out}}$ are abstract type-kinds, op is a domain-neutral description of the transformation, $R(o) \subseteq \Pi$ is the set of structural primitives the operator *requires*, $c(o) \in [0, 1]$ is extraction confidence, and $\text{prov}(o)$ is provenance (source paper / extraction trace). Write $\mathcal{O}$ for a corpus of operators.

2.3 Walls (clean assays)

A *wall* is a reference problem whose hardness is a single, well-understood structural deficiency. Formally a wall W is given by the set of primitives it *denies*,

$$D(W) \subseteq \Pi,$$

meaning: any operator requiring a denied primitive cannot function against W . The reference battery ($k = 5$ walls) and their denial sets:

$$\begin{aligned} W_1 = \text{ECDLP} : \quad & D(W_1) = \{\text{continuity, smoothness, homomorphism_to_easy,} \\ & \text{sparsity, spectral_support, analytic_continuation,} \\ & \text{correlation, metric, monotonicity}\}, \\ W_2 = \text{RandomOracle} : \quad & D(W_2) = \{\text{continuity, smoothness, linearity, low_rank,} \\ & \text{spectral_support, correlation, analytic_continuation}\}, \\ W_3 = \text{HighRankGeneric} : \quad & D(W_3) = \{\text{low_rank, sparsity, linearity, compactness}\}, \\ W_4 = \text{NonsmoothDiscrete} : \quad & D(W_4) = \{\text{continuity, smoothness, metric,} \\ & \text{monotonicity, analytic_continuation}\}, \\ W_5 = \text{NoSpectralStructure} : \quad & D(W_5) = \{\text{spectral_support, analytic_continuation, linearity}\}. \end{aligned}$$

3 The failure-signature and the dimension

3.1 Dissolution

An operator o *dissolves* against wall W iff it requires something the wall denies:

$$\text{diss}(o, W) = [R(o) \cap D(W) \neq \emptyset].$$

Equivalently, o survives W iff $R(o) \cap D(W) = \emptyset$. This is the formalization of “failure as a coordinate”: the pattern of dissolution is the measurement.

3.2 Signature

The *failure-signature* of o against battery $\mathcal{W} = (W_1, \dots, W_k)$ is the Boolean vector of survivals,

$$\sigma(o) = (\neg \text{diss}(o, W_1), \dots, \neg \text{diss}(o, W_k)) \in \{0, 1\}^k.$$

There are at most 2^k distinct signatures (reference: $2^5 = 32$).

3.3 Equivalence classes and dimension

Operators with equal signature are *structurally equivalent*; the classes partition $\mathcal{O}$:

$$[\sigma] = \{o \in \mathcal{O} : \sigma(o) = \sigma\}.$$

The **measured dimension of method space** over corpus $\mathcal{O}$ is the number of populated classes,

$$\boxed{\dim(\mathcal{O}) = |\{\sigma(o) : o \in \mathcal{O}\}|.}$$

The *collapse ratio* is $\dim(\mathcal{O})/|\mathcal{O}|$ and the *top- m concentration* is the fraction of operators in the m largest classes,

$$\kappa_m(\mathcal{O}) = \frac{1}{|\mathcal{O}|} \sum_{j=1}^m |[\sigma_{(j)}]|, \quad |[\sigma_{(1)}]| \geq |[\sigma_{(2)}]| \geq \dots$$

4 Cross-domain bridges

A *candidate cross-domain identity* (bridge) is a populated class containing operators from ≥ 2 distinct sources:

$$\text{Bridge}(\sigma) \iff |\{\text{prov}(o) : o \in [\sigma]\}| \geq 2.$$

The *shared structural source* of a bridge is the intersection of requirements,

$$S(\sigma) = \bigcap_{o \in [\sigma]} R(o),$$

i.e. the primitives all members genuinely share (may be a strict subset of what the common signature implies). Bridges between distant fields are the high-value output; see Task T6.

5 Composition calculus

5.1 Type compatibility

Let $\text{kinds}(\tau) \subseteq K$ map a concrete type to abstract kinds. Operators chain $B \circ A$ iff their type-kinds intersect:

$$A \rightsquigarrow B \iff \text{kinds}(\tau_{\text{out}}(A)) \cap \text{kinds}(\tau_{\text{in}}(B)) \neq \emptyset.$$

5.2 Composite operator

A composite $B \circ A$ inherits the union of requirements (it needs everything both stages need):

$$R(B \circ A) = R(A) \cup R(B), \quad \tau_{\text{in}}(B \circ A) = \tau_{\text{in}}(A), \quad \tau_{\text{out}}(B \circ A) = \tau_{\text{out}}(B).$$

To keep the composite set bounded, **compose at the class level**: one representative per signature class, so the number of candidate composites is $O(\dim(\mathcal{O})^2)$ rather than $O(|\mathcal{O}|^2)$.

5.3 Viability filter

A composite is *structurally viable* iff some structured substrate could supply all its requirements — equivalently, it is *not* dissolved by every wall:

$$\text{viable}(o) \iff \neg \bigwedge_{j=1}^k \text{diss}(o, W_j) \iff \exists W_j : R(o) \cap D(W_j) = \emptyset.$$

Viable composites are *manufactured candidate methods* (Task T5). They re-enter the dimension count: $\dim(\mathcal{O} \cup \mathcal{C})$ where $\mathcal{C}$ is the viable composite set.

6 Calibration conditions (acceptance tests)

These must hold after any change to extraction or the battery.

6.1 Negative control

Let $\mathcal{N}$ be a structureless corpus (random common text, no method vocabulary) and $\mathcal{P}$ a structured corpus. With extraction rate $\rho(\cdot) = |\text{operators extracted}|/|\text{papers}|$, require

$$\rho(\mathcal{N}) \leq 0.15 \rho(\mathcal{P}).$$

Interpretation: the instrument must read *near-zero structure* on noise, or its positives are untrustworthy. (Reference implementation achieves $\rho(\mathcal{N}) = 0$.)

6.2 Saturation (the falsifiable core)

For nested corpora $\mathcal{O}_{n_1} \subset \mathcal{O}_{n_2} \subset \dots$ of growing size, define the dimension curve $n \mapsto \dim(\mathcal{O}_n)$. The *small-basis hypothesis* is

$$\boxed{\exists N, d^* : \forall n > N, \dim(\mathcal{O}_n) = d^*}.$$

Conjecture 1 (Small basis). *On a real heterogeneous corpus, $\dim(\mathcal{O}_n)$ plateaus at a small d^* . If it does, d^* is the measured dimension of the relational field; if $\dim$ grows with n without plateau, no small basis exists at this resolution.*

Caveat: $d^* \leq 2^k$ always, so a plateau on synthetic data generated from $|\Pi|$ primitives is *expected and uninformative*; the test is only meaningful on data the instrument did not author (Task T1).

7 Reference implementation

The following files exist and run. Equation $\rightarrow$ file mapping:

- `core.py` — Π (STRUCTURAL_PRIMITIVES), walls W_j and $D(W_j)$ (WALLS), `Operator`, `diss` (`Wall.dissolves`), σ (signature).
- `ingest.py` — corpus loaders (JSONL/JSON/text-dir/PDF) $\rightarrow$ `PaperRecord`; `auto_ingest` dispatches by format.
- `extract_precise.py` — the precision extractor: weighted phrase tiers (strong/medium/weak), per-primitive evidence threshold, negative cues. Implements $o \mapsto R(o), c(o)$ from text. Use this as the default heuristic extractor.
- `extract_rich.py` — LLM extraction (one model call per paper) for vocabulary-independent reading; same `Operator` schema. Requires API access.
- `pipeline.py` — $[\sigma]$ (`equivalence_classes`), bridges, $A \rightsquigarrow B$ (`admissible_compositions`), viable (`composite_viable`), `dim` (`measure_dimension`); `TYPE_KINDS` is `kinds()`.
- `belt.py` — end-to-end driver: `ingest` $\rightarrow$ `extract` $\rightarrow$ `signature` $\rightarrow$ `dimension/concentration` $\rightarrow$ `bridges` $\rightarrow$ `class-level composition` $\rightarrow$ `JSON export`. Run: `python3 belt.py <path>`.
- `belt_tests.py` — negative control + saturation harness (Section 6).
- `seed_library.py` — hand-characterized operators used as a validation set; the matcher reproduces the manually-derived identities (phase-gradient $\equiv$ swarm-coherence via continuity; spectral $\equiv$ BSD-growth via analytic_continuation).

Run order:

```
python3 core.py           # primitives + wall battery
python3 seed_library.py   # operators + signatures (validation)
python3 belt.py           # full belt on seed set, writes belt_output/
python3 belt_tests.py     # negative control + saturation (synthetic)
python3 belt.py <arxiv>   # full belt on a real corpus
```

8 Remaining build tasks (priority order)

T1 — Real-corpus saturation (decisive). Replace the synthetic generator in `belt_tests.py::saturation_corpus` with `auto_ingest(path)` over a real heterogeneous dump (arXiv metadata snapshot recommended: one JSON/line with `title`, `abstract`, `categories`). Compute $\dim(\mathcal{O}_n)$ for $n \in \{10^2, 10^3, 10^4, 10^5\}$, plot, and report whether it plateaus. This is the experiment the whole program exists to run. Acceptance: a dimension-vs- n curve on real data with a plateau-or-not verdict.

T2 — Rich extraction at scale on the interesting subset. The heuristic extractor has a known recall gap across field-specific vocabularies (e.g. biology’s “minimize Gibbs free energy” vs. ML’s “gradient descent”; both are smoothness/compactness). For any class flagged interesting by T1/T6, re-extract its papers with `extract_rich.py` (vocabulary-independent) and recompute signatures. Acceptance: rich-extracted operators raise recall on a held-out cross-domain set (measure recall on the 9 reference abstracts where the heuristic dropped 4).

T3 — Confidence-weighted dimension. Recompute dim counting only operators with $c(o) \geq \theta$, sweeping $\theta \in [0.4, 0.9]$. Plot dim vs. θ . If the small-basis result persists as θ rises (keeping only high-confidence operators), it is robust; if it appears only at low θ , it is extraction noise. Acceptance: a $\text{dim}(\theta)$ curve.

T4 — Plural-wall expansion. Add walls denying axes not currently covered (candidates: temporal/causal-order, group/algebraic-closure, locality, stationarity), each as a $D(W) \subseteq \Pi$ (extend Π if a genuinely new axis is needed). Re-run T1. Test: does the plateau d^* stay stable as k rises (real clustering) or does dim track 2^k (artifact of coarse walls)? Acceptance: d^* as a function of k .

T5 — Composite semantic-coherence filter. Structural viability (cheap) over-admits. For each viable composite, call `extract_rich`-style judgment: is the composite *operation* a coherent method? This is the two-stage filter: structural (Section 6.3) then semantic. Acceptance: a ranked list of composites that pass both filters, each with provenance of its two source classes.

T6 — Cross-category bridge bias. Using arXiv `categories`, measure whether bridges connect *distant* categories (e.g. `q-bio` $\leftrightarrow$ `cs`) more or less than chance. Distant-field bridges are the valuable output. Acceptance: an observed-vs-expected distant-bridge rate.

T7 — Primitive co-occurrence / effective dimension. Build the graph on Π with edges weighted by co-occurrence in $R(o)$ across $\mathcal{O}$; find tight clusters. If primitives co-travel (e.g. continuity+smoothness+correlation), the *effective* primitive count $< |\Pi|$, a deeper basis measurement. Acceptance: a clustered primitive graph and an effective-dimension estimate.

T8 — Reduction-replay validation. For operators whose signature dissolves against ECDLP, verify the verdict matches the hand-derived reductions from the originating research (phase-gradient, solvent/homomorphism, sparsity, binaural, bundle-slope all reduce to discrete log). Acceptance: agreement between belt verdicts and the documented reductions on the known cases.

T9 — Paraphrase stability. Run `extract_rich` on the same paper phrased three ways (title-only / abstract / full body) and check σ is invariant. Flag operators whose signature is phrasing-dependent as low-reliability. Acceptance: a per-operator stability score.

T10 — Export a queryable field graph. Emit the final artifact: nodes = operators (with signature, requirements, provenance), edges = bridges (shared-signature) and viable compositions; serialize to a graph format (e.g. JSON-graph / GraphML) for interactive query. Acceptance: a loadable graph whose highlighted subgraph is the bridges + viable composites.

9 Honesty constraints (must hold throughout)

- Every operator is a *candidate with provenance*, never a discovery. Outputs are claims to be verified.

- Extraction fidelity is the quality gate and is model-judgment-laden; report confidence and keep provenance so mis-extractions are traceable.
- A plateau on synthetic/self-authored data is *uninformative* ($d^* \leq 2^k$ by construction); only real-corpus saturation (T1) tests the hypothesis.
- The battery reads only the axes its walls deny; a measured d^* is the dimension of method-space *as projected through* $(\Pi, \mathcal{W})$, not a domain-independent absolute. State this whenever reporting d^* .
- Public release is recommended as adversarial validation; provenance-per-classification converts criticism into debugging.

10 Summary for the agent

The instrument computes, from a corpus of papers, the number of distinct structural failure-signatures (Eq. in Section 3), which is the measured dimension of method space. The code runs end-to-end now on a validation set and synthetic data, passing the negative-control and saturation acceptance tests. The decisive open task is T1: run the saturation measurement on a real heterogeneous corpus and report whether the dimension plateaus. Execute T1 first; then T2–T10 in order, holding the honesty constraints throughout. All equations above are implemented or directly implementable against the named files.