Sheaf — The Mathematics of Local-to-Global Consistency for Multi-Agent AI

Jack Widman, PhD

A technical whitepaper · 2026

Abstract

Multi-agent and multi-source AI systems are *locally competent and globally blind*: each agent, model, or source does its part well, but nothing measures whether the parts actually **cohere** into a single trustworthy whole. This is not a prompt-engineering problem — it is a **mathematical** one, and it has a name. **Sheaf theory** is the branch of mathematics that studies precisely how local data on overlapping pieces of a problem glues into a consistent global whole, and how to detect — and quantify — when it cannot. Sheaf operationalizes this: it assembles independent local views into one answer and returns a measured account of where they **agree** (the consistent core, H^0) and where they **irreducibly disagree** (the obstruction, H^1). This paper states the problem, the mathematics, and the construction that computes these invariants **literally** on structured outputs — down to the cohomology of contextuality — and the genuine first-cohomology computation that ships today. Our flagship instantiation is a five-discipline financial-markets desk; the engine is domain-general.

---

1. The problem: local competence, global incoherence

Put several capable agents on a task — five analysts, five language models, a pipeline of tools, a stack of source documents — and each one is **locally accurate and globally blind**. Each sees a slice, does its part well, and hands off. Across those handoffs, small local accuracies compound into **global drift**: contradiction, double-counting, silent disagreement that no single agent is positioned to notice. It is the telephone game at machine scale.

The current generation of "orchestration" tooling is **plumbing**. It routes messages between agents, retries on failure, and draws a flowchart. It answers *did the messages arrive?* It does not answer the only question that matters in a high-stakes workflow:

> **Do the pieces actually cohere — and if not, exactly where, and how much?**

In low-stakes settings you can ignore this; a wrong group-chat answer costs nothing. In **finance, law, compliance, intelligence, and scientific research**, "the agents mostly agreed" is not an answer. The cost of undetected incoherence is the whole risk. These domains do not need a more confident black box; they need a **measured, auditable** one.

---

2. The shape of the problem is the shape of a sheaf

The structure above — many **local** views, defined on overlapping parts of a shared problem, that must be assembled into one **global** answer — is not a loose analogy to sheaf theory. It is, formally, the setting sheaf theory was built for. The correspondence becomes exact **once we fix a cover and restriction maps** — these are modeling choices we make, not structure the workflow hands us for free (§2 caveats; §6):

Multi-agent system	Sheaf-theoretic object
the task / shared context	the base space X
sub-tasks, the domains assigned to agents	an open cover $\{U_i\}$ of X
each agent's local output	a **section** $s_i \in \mathcal{F}(U_i)$
handing an agent only its slice of context	the **restriction maps** res^U_V
"do the overlapping outputs agree?"	compatibility on $U_i \cap U_j$
assembling one coherent global answer	**gluing** to a global section
the irreducible disagreement that blocks gluing	H^1 — the **obstruction**
what everyone consistently agrees on	H^0 — the **global sections**

The value of taking the mathematics seriously is that it does not merely *describe* the problem — it comes with the **right invariants** (H^0, H^1) for measuring coherence, and a century of theory for computing with them.

Two things we make precise up front. First, X is **finite and discrete** — the agents' domains are finitely many context windows, not open sets of a manifold — so the point-set language above is only a bridge; the object we actually compute on is **combinatorial**: the **nerve** of the cover (§3), a finite simplicial/cell complex on which Čech and cellular cohomology are exactly linear algebra. Second, the cover $\{U_i\}$ is **not fixed by the architecture** — which claims overlap is *extracted* from the outputs (§4). A different extraction yields a different nerve, so H^0/H^1 are computed *relative to* the extracted structure, and their stability across runs is itself a quantity to measure, not assume (§6).

---

3. Sheaf theory, precisely enough

(Standard material, stated for a technical but non-specialist reader.)

A **presheaf** $\mathcal{F}$ on a space X assigns to every open set U a set (or vector space) $\mathcal{F}(U)$ — the **sections over** U — together with **restriction maps** $\mathrm{res}^U_V \colon \mathcal{F}(U) \to \mathcal{F}(V)$ for every inclusion $V \subseteq U$, compatible under composition. Sections are "local data"; restriction is "look at the data on a smaller region."

A presheaf is a **sheaf** when local data that is *mutually compatible* assembles uniquely into global data — the two sheaf axioms over any open cover $\{U_i\}$ of U:

- **Locality (separation).** If two global sections restrict to the same thing on every U_i, they are equal. (No hidden global degrees of freedom.)
- **Gluing.** If local sections $s_i \in \mathcal{F}(U_i)$ agree on every overlap ($s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$), there is a **unique** global section s restricting to each s_i.

The deep point: agent outputs generally form a **presheaf**, and a presheaf's compatible local data *need not amalgamate* — local agreement on overlaps does not, by itself, assemble into a single global section. (For a genuine sheaf the gluing axiom guarantees it does; the gap between presheaf and sheaf is exactly where inconsistency lives.) The precise measure of that gap is **cohomology**:

- $H^0(\mathcal{F})$ — the **global sections**: the data that assembles into one globally consistent whole — the **consistent core**. (Strictly this is the space of global sections, which can carry structure no single view states outright; informally, what the local views genuinely cohere into.)
- $H^1(\mathcal{F})$ — the **first cohomology**: the **obstruction** to gluing. Nonzero H^1 means there is local, overlap-compatible data that *cannot* be reconciled into a single global whole. This is **inconsistency made into a number.**

Computationally, for an open cover this is **Čech cohomology**, computed on the **nerve** of the cover — the simplicial complex whose vertices are the cover sets U_i and whose simplices are their nonempty intersections. The nerve turns "which agents' domains overlap" into a combinatorial object on which H^0 and H^1 are linear algebra.

---

4. From theory to a computable signal: the $\mathbf{H^1}$ Inconsistency Index *(the genuine hybrid we ship today)*

There is no canonical sheaf on *raw* free text — "the sheaf of an LLM's prose" is not a mathematical object until one fixes a representation. So the boundary is **not** *whether* we can compute sheaf cohomology: once the outputs carry a structured representation (§5), the cohomology is **literal and exact** — and for distributional outputs (§5.5) it is the literal first cohomology of contextuality, a nonzero H^1. The boundary is that the **representation-extraction step is model-mediated**. What we *ship today* is a **genuine first-cohomology computation** on the agent-overlap graph (below): a literal nonzero H^1 when the comparative claims extracted from the agents' answers form a directed preference cycle that cannot glue into a consistent global ranking, and the deterministic H^0 disagreement when they do glue — genuine cohomology either way. The deeper cellular-sheaf and contextuality construction of §5 extends this to richer structure, not a different aspiration.

The pipeline:

1. **Local sections.** N independent agents (models, or discipline-specific analysts) each answer the same question from their own context — these are the local sections s_i.
2. **Overlaps and compatibility.** A semantic analyzer identifies the **claims the agents share** (the overlaps $U_i \cap U_j$) and, for each, whether the agents **agree** or **conflict** on it.
3. **The deterministic index.** Over the resulting agent-overlap graph we compute, *not* with a model but with arithmetic:

$$H^1 \;=\; 100 \times \frac{\#\{\text{conflicting overlap pairs}\}}{\#\{\text{total overlap pairs}\}}, \qquad \text{Coherence} = 100 - H^1.$$

   The semantics (what overlaps, agree-or-conflict) are detected by a model; **the invariant itself is deterministic**, so the number is a measured quantity, not a model guessing a score. We write H^1 here for continuity with the product's name, but state the precise mathematics: this conflict-pair formula — like the absolute-scalar measure of §5.4 — *is* the genuine H^0 **disagreement** measure (a deterministic distance from consensus over the overlap graph), labeled as H^0, not a stand-in for H^1. When the extracted comparative claims form a directed preference cycle that cannot glue into a single consistent ranking, the obstruction is a literal nonzero H^1, which we compute and report with the offending cycle. §5.6 fixes the terminology; we reserve literal H^1 for exactly these genuine obstructions (here and in §5.4–§5.5).
4. **The two outputs that matter.** A **consistent core** (H^0 analogue — what the agents cohere on, the high-conviction result) and a **conflict map** (H^1 — exactly where, and how much, they irreducibly disagree).

Scope of the shipping signal. This *shipping* signal is **genuine cohomology** — a literal H^1 when the comparative claims cycle, the deterministic H^0 disagreement when they glue; the cellular-sheaf and contextuality construction of §5 deepens it, not a different goal. It inherits the quality of the semantic analyzer (claim extraction and agreement-detection are the hard upstream steps). It yields a **relative signal**, not an oracle: it tells you *where to look and how much to worry* — exactly what a risk or research function needs. The arithmetic is reproducible; the semantics are model-mediated.

The load-bearing dependency, stated plainly. Making the math exact does *not* make the answer exact: the cohomology is exact *given the extracted nerve and stalk data*, but a mis-extraction — the analyzer mapping two unrelated claims to the same variable, or missing a real overlap — yields a precise, beautiful, and **fictional** score. The precision is bottlenecked by the **semantic alignment error of the upstream parser**, and no amount of downstream linear algebra repairs it. This is why the extracted structure (claims, overlaps, weights) is surfaced for **audit** rather than hidden, and why extraction stability — across reruns and paraphrases — is a first-class metric, not an afterthought. The geometry is only as trustworthy as the parse; we instrument the parse accordingly.

---

5. The rigorous construction: cellular sheaves and the sheaf Laplacian

The genuine cohomology of §4 gains its full cellular-sheaf form the moment agent outputs are **structured** rather than free text. We require, per shared claim, a small piece of structure from each participating agent; on that data the modern computational theory of **cellular sheaves** (Hansen–Ghrist, *Toward a Spectral Theory of Cellular Sheaves*) applies exactly, and consistency becomes linear algebra over a real sheaf.

5.1 Structuring the outputs. The analyzer extracts the shared claims $\mathcal{C}$ and, for each agent v addressing a claim c, a structured datum in one of three representations: an **absolute position** $x_{v,c} \in \mathbb{R}$ on a normalized axis (bearish -1 … bullish $+1$, or a standardized estimate); a **relational judgment** $y_{(u,v),c} \in \mathbb{R}$ comparing two agents' views; or — the richest — a **distribution** $p_{v,c}$ over the claim's possible outcomes (a forecast). Each carries a **confidence** $w \geq 0$. All three yield *genuine* cohomology, but they place the content in different degrees: scalars and relations drive the **spectral** construction of §5.2–5.4 (the coherence score and the consensus answer), while distributions drive the **contextuality** construction of §5.5 — which is where a literal, nonzero H^1 lives. The choice is the central modeling decision (§5.6).

5.2 The sheaf. Form the **overlap graph** $G = (V, E)$: vertices = agents; for each claim c and each pair (u, v) both addressing it, an edge $e = (u, v, c)$. A cellular sheaf $\mathcal{F}$ assigns a stalk to each cell — the **vertex stalk** $\mathcal{F}(v) = \mathbb{R}^{\mathcal{C}_v}$ (agent v's positions across the claims it addresses) and the **edge stalk** $\mathcal{F}(e) = \mathbb{R}$ (the single value the two views must reconcile) — with **restriction maps** $\mathcal{F}_{v \trianglelefteq e}$ the confidence-weighted projection onto the shared coordinate, $x \mapsto \sqrt{w_{v,c}}\, x_{v,c}$.

5.3 Complex, coboundary, Laplacian. The data forms the two-term cochain complex $0 \to C^0(\mathcal{F}) \xrightarrow{\delta} C^1(\mathcal{F}) \to 0$ with coboundary $(\delta x)(u, v, c) = \sqrt{w_{u,c} w_{v,c}}\, (x_{v,c} - x_{u,c})$ and **sheaf Laplacian** $L = \delta^\top \delta$, positive semidefinite, with

$$x^\top L x \;=\; \|\delta x\|^2 \;=\; \sum_{(u,v,c) \in E} w_{uv,c}\, (x_{u,c} - x_{v,c})^2.$$

5.4 The Hodge decomposition — the spine of the construction. Because the complex has no 2-cells, edge cochains split orthogonally

$$C^1 \;=\; \underbrace{\operatorname{im} \delta}_{\text{gradient — explainable}} \;\oplus\; \underbrace{\ker \delta^\top}_{\text{harmonic } \cong\, H^1 \text{ — irreducible}}.$$

By finite-dimensional Hodge theory for cellular sheaves, harmonic cochains are *canonical* class representatives — $H^1 \cong \ker \delta^\top = \ker L_1$ for the up-Laplacian $L_1 = \delta \delta^\top$ — so the splitting is well-defined, not a choice of basis. Every pattern of pairwise disagreement is then uniquely a **gradient part** (disagreement fully accounted for by *some* global assignment of positions — a potential) plus a **harmonic part** (cyclic frustration that *no* global assignment can remove). This one decomposition yields both numbers:

- **Coherence — the H^0 side, the headline.** With *absolute* data the agents' positions $x^\star$ are themselves a potential, so $\delta x^\star$ is entirely gradient and

  $$\text{Coherence} \;=\; 1 - \frac{x^{\star\top} L x^\star}{E_{\max}} \;\in [0, 1]$$

  measures *how far these analysts sit from a consistent whole.* Boundedness is by construction: each claim's axis is normalized to $[-1, 1]$, so an edge contributes at most $\Delta_{\max}^2 = 4$ and $E_{\max} = \sum_e w_e\, \Delta_{\max}^2$ caps the energy. It is confidence-weighted and **differentiable** (so $\partial\,\text{Coherence}/\partial(\text{claim})$ ranks which claims carry the disagreement). It is a **weighted continuous analogue** of §4 — *not* an exact generalization: the quadratic energy and the conflict-pair count induce different geometries (one large disagreement vs. many small ones score differently). The **reconciliation** $\hat{x} = \operatorname{proj}_{\ker L} x^\star$ is the global section **nearest the data** (orthogonal projection onto $\ker L$ — note every global section has zero energy, so the right object is the *closest* one, not a "least-energy" one); on each connected claim it is the consensus value the desk's views collapse to — *the sheaf's own answer.*
- **Irreducible inconsistency — the H^1 side.** With *relational* data y, the harmonic component $h = y - \delta \hat{x} \in \ker \delta^\top$ is generally **nonzero**: the part of the disagreement that is cyclically frustrated

(A over B, B over C, C over A) — genuine, data-carrying H^1. Then $\|h\|^2/\|y\|^2$ is a *literally cohomological* inconsistency index, and the support of h localizes the obstruction to specific claim-loops. The distributional construction of §5.5 sharpens this to the literal first cohomology of *contextuality*.

- **Robustness — secondary.** $\dim\ker L = \dim H^0$ counts independent consistency components; the **spectral gap** λ_2 (smallest nonzero eigenvalue, or its normalized form) measures how robustly a unique consensus is enforced — the convergence rate of the diffusion $\dot{x} = -Lx$, and by Cheeger-type bounds how near the constraint structure sits to disconnection. Crucially λ_2 is a property of the *structure* L, largely independent of the answers $x^\star$ — it asks "is the agreement structure tight or fragile," not "did they agree" — so we surface it as a labeled robustness read, never as the coherence score.

5.5 The distributional construction: contextuality and a literal H^1. When agents output **distributions** rather than point positions, a second, sharper construction applies — and it is where H^1 becomes literally nonzero and genuinely meaningful. We import the **sheaf-theoretic theory of contextuality** (Abramsky–Brandenburger 2011; Abramsky–Mansfield–Barbosa 2012) directly.

Setup. Let X be the shared variables underlying the claims. Each agent v addresses a **context** $C_v \subseteq X$ and reports a distribution p_v over the joint outcomes of C_v — its **local section**. The contexts $\{C_v\}$ form a cover; their intersection pattern is the measurement scenario. The whole construction presupposes a well-posed extraction map **text** $\mapsto (X, \{C_v\}, \{p_v\})$ — shared variables with **agreed outcome spaces**, the contexts, and the per-agent distributions. This is the binding modeling assumption: LLM prose does not satisfy it for free, the assumption carries real weight, and where it is loose the contextuality verdict is only as meaningful as that map. We treat the map as part of the model — surfaced and audited (cf. §4), not assumed away.

The sheaf and the gluing question. The assignment $U \mapsto \{\text{distributions over outcomes of } U\}$, with **marginalization** as restriction, is a presheaf $\mathcal{D}$ of distributions. The family $\{p_v\}$ is **compatible** when agents agree on overlaps — equal marginals on shared variables, $p_u|_{C_u \cap C_v} = p_v|_{C_u \cap C_v}$ — and the system is **consistent** iff that compatible family **glues** to one **global section**: a joint distribution p over all of X whose marginal on each C_v is exactly p_v.

The obstruction is literal cohomology. When the local data is compatible on every overlap yet **no global joint exists**, the system is **contextual** — pairwise-consistent, globally impossible. This is precisely a failure of the sheaf gluing axiom, and Abramsky–Mansfield–Barbosa make it cohomological: a Čech complex (coefficients in the presheaf of $\mathbb{Z}$-modules generated by the outcomes) carries a class whose non-vanishing in $\check{H}^1$ **certifies** the obstruction. That class is the literal H^1 inconsistency — the same machinery that renders Bell/Kochen–Specker non-locality a cohomological obstruction, applied to a desk of analysts.

What we actually compute — two complementary tests.

- **Feasibility (exact, necessary-and-sufficient).** A joint distribution matching all marginals exists iff a **linear program** is feasible (the marginal-consistency polytope is nonempty). Its infeasibility — or the minimal marginal slack / transport distance to feasibility — is an **exact, continuous** inconsistency measure: a quantitative "how far from a coherent joint world-model."
- **Cohomological certificate (sufficient).** The Čech $\check{H}^1$ class is a rigorous *witness*: when it is nonzero, the data is provably non-gluable. **Known limitation:** the obstruction is **sufficient but not necessary**. It is a *partial* invariant — it depends on the coefficient ring ($\mathbb{Z}$ vs. $\mathbb{Z}_2$ resolve different obstructions) and can miss incompatibility carried in higher-order correlations (higher cohomology), so there are genuine *false negatives*: globally inconsistent data whose $\check{H}^1$ class vanishes. We therefore treat $\check{H}^1$ as a **one-sided certificate** (nonzero $\Rightarrow$ provably inconsistent) and the LP feasibility/slack below as the **complete** measure; a vanishing class is never read as a clean bill of health.

This makes "H^1 Inconsistency Index" **literal** — a genuine first cohomology class, not a heuristic — on a decade of peer-reviewed applied sheaf theory, with an exact continuous companion. It is **computable within an explicit bound.** The feasibility LP ranges over the joint distribution on the shared variables X, whose dimension is $\prod_{x\in X} |\text{outcomes}(x)|$ — **exponential in** $|X|$ in the worst case. It is tractable precisely when the shared-variable set is **small and bounded-cardinality** (or the context structure has low treewidth, so the joint factors over the nerve's cliques rather than being materialized whole). In production we therefore **bound context size** — a desk reconciles a handful of shared variables per call, not its entire state — and discretize continuous outcomes to low cardinality; we do *not* claim feasibility over, say, a twenty-variable continuous joint. This is the construction we build toward to earn the name literally.

5.6 Scope — and the decision that earns the name. Given the structured data, every invariant above is *exact* cohomology of a real sheaf: measured, reproducible, model-free arithmetic. Two qualifications we will not paper over. First, **the extraction is model-mediated** — turning text into $(x_{v,c}, w)$, $y_{(u,v),c}$, or $p_{v,c}$ is the upstream step, and the invariant is only as faithful as it is (the extraction is itself inspectable and improvable). Second, and sharper: **with absolute-position data the harmonic part is structurally zero** — $\delta x^\star$ is always a coboundary, so there is *no* irreducible obstruction, and "inconsistency" is precisely the H^0 **defect** (distance to consensus), measured by energy, not H^1. A literally-nonzero H^1 — what our index is *named* for — requires **relational stalks, non-trivial restriction maps, or (cleanest) the distributional/contextuality construction of §5.5**, where pairwise-compatible views provably fail to glue. So the representation we extract is the decision that determines whether the "H^1 Index" is computed *exactly* or *estimated*; **we build toward the distributional construction so the name is earned literally**, and we state, per release, exactly which we compute.

To keep the mathematics and the label aligned, we name the absolute-scalar and §4 conflict-pair quantity the **Coherence Defect** (equivalently the *Sheaf Energy* — the H^0 distance from consensus), and reserve the H^1 **Inconsistency Index** for the relational and distributional/contextuality constructions, where the harmonic class is genuinely nonzero. A surface that exposes the energy measure under an "H^1" label is, strictly, reporting the Coherence Defect; aligning the product's nomenclature to this distinction is a deliberate near-term change, not a cosmetic one.

5.7 The frontier. Beyond these: **higher-rank vector/operator stalks**, **learned, non-trivial restriction maps** (agents reasoning in different but linearly-relatable frames — exactly where Hansen–Ghrist's spectral theory is richest), and **higher cells** (a claim shared by k agents as a single $(k-1)$-cell, giving genuine $H^{\geq 2}$ and finer obstructions). One scope boundary is worth naming: the symmetric sheaf Laplacian $L = \delta^\top \delta$ models **symmetric comparison** of views — appropriate for a desk that reasons independently and is then compared. It is symmetric PSD by construction, for *any* restriction maps and orientation; what it cannot *represent* is **directed influence** (agent A's output constraining B but not the reverse), which calls for a different object — directed/connection Laplacians on an oriented complex. We treat that as out of scope for a *consistency* measure, and a frontier item, not a defect of the present construction. These are real research; we describe them as such.

§5.1–5.5 is **implementable today** — linear algebra for the spectral construction, plus a linear program for the contextuality test — and is the specification for our reference computation; §5.6 names the modeling choice we make; §5.7 is the open horizon. This is the defensible technical core and the natural home of a **provisional patent** and an **open-source reference implementation** — a working invariant in production with a clearly-marked research frontier, not vaporware.

---

6. Why this is defensible

- **It is real mathematics, not a metaphor.** The invariants H^0/H^1 are the *correct* tools for the problem, not a borrowed aesthetic. Most "multi-model" products stop at ensembling or voting; none measure the obstruction to coherence.
- **Precision is the moat (§9).** The buyers and investors this attracts will check the math. Being exactly right — and explicit about what is computed and what is frontier — is a durable advantage that a competitor cannot fake by adding the word "sheaf" to a landing page.
- **Founder–field fit.** Sheaf is built by a researcher in topology — the field sheaf theory belongs to. The rigor is not the marketing; it is the training.

---

7. Flagship application: financial markets

Finance is the **hardest version** of the local-to-global problem: high stakes, genuine and legitimate disagreement, and a hard requirement — regulatory and cultural — to **show your work**.

Sheaf's flagship is a **five-discipline desk**: Fundamental, Macro, Quant, Technical, and Risk. Each is a local section reasoning from its own framework; disagreement **emerges from the frameworks**, never an assigned bull/bear stance. The system returns the **consistent core** (the high-conviction view the desk coheres on) and the **conflict map** with an H^1 score (where and how much they diverge — on direction, magnitude, timing, or which risk dominates). For a risk or investment-committee function this is a **model-risk and governance instrument**: a measured, auditable answer to "where might our multi-model view be quietly wrong?"

---

8. Generalization: any local-to-global problem

The engine is domain-general — it applies wherever many local sources must cohere into one trustworthy whole and you need the disagreement made explicit:

- **Research synthesis** — many papers/sources → one defensible finding + every contradiction surfaced.
- **Legal & compliance** — clauses, precedents, reviewers, policies → a coherent position + the conflicts.
- **Intelligence analysis** — multiple reports → one assessment with confidence and dissent kept visible.
- **Multi-model AI / QA** — many models → one synthesized answer + a number for how much they agree.
- **Due diligence** — disparate documents and signals → one verdict + the disagreements worth a second look.
- **Data reconciliation** — local stores that should agree → a unified whole + a map of where they don't.

---

9. Why precision is the moat

One discipline runs through the product and this paper: **claim exactly the rigor we deliver — no more, no less.** We compute genuine cohomology and call it exactly that — a literal H^1 where a cyclic obstruction exists, the genuine H^0 disagreement where it does not — and we mark the boundary between the shipped hybrid and the deeper construction of §5 precisely (§4–§5). This is not modesty; it is strategy. The entire value proposition is *measured trust* — and a consistency product that overclaims its own consistency detonates the moment a sharp reviewer, exactly the buyer we want, checks the math. Checkable rigor is the one advantage a competitor cannot fake by adding the word "sheaf" to a landing page. The rigor is not the marketing; it is the moat.

---

10. Conclusion

Multi-agent AI gave us many local minds. It did not give us a way to know whether they add up to one coherent whole — or where they quietly fall apart. That is a mathematical question with a mathematical answer. **Sheaf brings many local views into one coherent whole, and measures exactly where they agree and where they don't.** Proven where it is hardest — financial markets — and built for any domain.

> *Where others route, we glue — and we measure the seams.*

Appendix A — References

- J. Curry, *Sheaves, Cosheaves and Applications* (PhD thesis) — applied sheaf theory foundations.

- J. Hansen & R. Ghrist, *Toward a Spectral Theory of Cellular Sheaves* — the sheaf Laplacian; the computational backbone of §5.2–5.4.
- S. Abramsky & A. Brandenburger, *The Sheaf-Theoretic Structure of Non-Locality and Contextuality* — the presheaf-of-distributions / no-global-section framework underpinning §5.5.
- S. Abramsky, S. Mansfield & R. S. Barbosa, *The Cohomology of Non-Locality and Contextuality* — the Čech $\check{H}^1$ obstruction (and its sufficient-not-necessary status) we build on in §5.5.
- M. Robinson, *Topological Signal Processing* — sheaves for heterogeneous data fusion (directly on-thesis for §8).
- R. Ghrist, *Elementary Applied Topology* — accessible entry point.