Discourse Complexes: A Sheaf Model of Feature Binding and a Combinatorial Criterion for the Sufficiency of Pairwise Constraints

Bradley Z. Vigil[1]

[1] *Institute for Studies in Pragmaticism*

Abstract

The feature-binding problem asks how a bank of parallel detectors can represent that one feature belongs *with* another rather than merely co-occurring with it. A long tradition answers by *coherence*: a percept is a joint assignment obeying a family of relational constraints, and a binding error is a violation of one of them. Cellular sheaf theory makes coherence exact, with global sections modeling coherent percepts; but the hypergraph sheaf Laplacians in current use reduce each constraint to its pairwise, clique-expansion part, while Peirce's reduction thesis warns that a genuinely triadic relation admits no such reduction *on its own carrier*. We introduce the *discourse complex*, a cellular sheaf on the incidence structure of a hypergraph whose higher cells carry, through a *sum-type* coboundary, linear constraints not generated by binary-supported ones on the same carrier, and use it to recast binding as the existence of a global section. Working in the scalar-stalk setting, we introduce a combinatorial invariant of the bare cell-incidence hypergraph, *binary-decomposability*, and *conjecture* that it characterises when the global-section space admits an equivalent pairwise description. We prove the characterisation for hypergraphs admitting a reduction order that meets a bipartite-matching condition, establish partial necessity under the same condition, and provide computational evidence on several hundred instances for the general conjecture, isolating the exact combinatorial gap that remains. A controlled experiment then illustrates the practical force of the sum-type choice: the triadic obstruction perfectly separates a class of binding errors that the dyadic (clique-expansion) statistic T_{pair} is provably unable to see, and stays discriminative well into the noisy regime. We present the discourse complex as a linear, Peirce-inspired model of one structural aspect of relational non-reduction, not as a complete theory of perception or of Peircean Thirdness.

Contents

1 Introduction **2**

2 Background **3**

3 The discourse complex **4**

4 Non-reduction **6**

5 From Peircean relational graphs to discourse complexes **8**

6 When are pairwise constraints enough? **10**

6.1 Genericity, made precise 10

6.2 The right notion of pairwise expressibility 11

6.3 Two plausible criteria that fail . . . 12
6.4 Binary-decomposability . . . 13
6.5 The characterization . . . 14
6.6 Verification . . . 15

7 An experiment: detecting binding failure **17**
7.1 The generative model . . . 17
7.2 Scene classes . . . 17
7.3 A prediction and its proof . . . 17
7.4 Results . . . 18

8 Discussion **19**

9 Open problems **20**

1 Introduction

Show a subject a red circle beside a green square briefly and away from the center of attention, and ask them to report what they see; on occasion the subject will report a green circle. The features were seen; it is their *pairing* that failed. Such *illusory conjunctions* [12] dramatize a problem that any system of distributed representation must solve: how does a bank of parallel detectors encode not just *which* features are present, but *which belong together*, and why does the encoding sometimes slip? This is the *feature-binding problem*, and it appears wherever local pieces of evidence must be assembled into a structured whole.

A resilient family of proposals solves binding by *coherence*. On this view a percept is not a set of active features but a joint assignment that satisfies a system of relational constraints; a misbinding is exactly a configuration that violates one of them [13]. Coherence accounts are attractive for being compositional and local: global consistency is assembled from many small, checkable relations. They also have a natural mathematical home. A *cellular sheaf* decorates a combinatorial domain with vector spaces of local data and linear maps that declare when local data agree, and its *sheaf Laplacian* has as its kernel the *global sections*, the assignments that are consistent everywhere [5, 3]. Read feature populations as the local data and coherence relations as the agreement maps, and "is this percept correctly bound?" becomes "does this assignment lie in the kernel of the sheaf Laplacian?"

There is, however, a fault line between how coherence constraints are usually modeled and what binding may demand. The sheaf constructions that scale to hypergraphs build their Laplacian from *pairwise* energies: a hyperedge on k vertices contributes the sum of its $\binom{k}{2}$ pairwise agreements, so the operator is really that of the edge's clique expansion [4]. Every constraint is thereby ground into two-variable pieces. Yet a line of thought descending from Peirce insists that some relations resist precisely this grinding. A genuinely triadic relation, taken on its own correlates, is not generated by binary relations among them; one can represent it with binary relations only by *enlarging* the carrier with an auxiliary variable and then projecting that variable away. This is the content of Peirce's *reduction thesis*, sharpened by Herzberger [7], Burch [2], and Koshkin [8], the last proving that a generic n-ary relation has ternarity $n-2$, so that a generic ternary relation is irreducibly three-way *on its carrier*. If some binding constraints are triadic in this sense, then a model that carries only pairwise cells on the same variables is structurally blind to a whole species of binding error, and no pairwise constraints on those variables can recover it, although an enlarged model with a mediating variable can. The distinction between same-carrier irreducibility and reducibility after introducing a mediator is, we will argue, the mathematically and philosophically substantive one.

This paper takes the fault line seriously and makes it precise. Section 3 introduces the *discourse complex*: a cellular sheaf on the incidence structure of a hypergraph, whose cells of arity three and higher carry a constraint through a *sum-type* coboundary rather than the pairwise clique expansion. On a dyadic cell, under the standard opposite-sign orientation, the two coboundaries coincide; on triadic cells they diverge in kind, and it is on that divergence that the same-carrier irreducible binding errors live. Within this setting we model a coherently bound percept as a global section, a misbinding as a nonzero coboundary, and the norm of the coboundary as a quantitative binding-error signal.

We outline our contributions below.

(1) *A synthesis.* We assemble discourse sheaves [6], hypergraph sheaf Laplacians [4], and the reduction thesis [8] into one linear model of higher-arity relational constraint (Section 3), and we prove that a triadic constraint is the consensus space of no dyadic complex *on the same carrier* (Section 4), while an auxiliary mediating variable can carry it.

(2) *A characterization.* We ask, in the scalar-stalk setting, when the global sections of a *general* discourse complex admit an equivalent pairwise description, and propose a purely combinatorial condition on the cell-incidence hypergraph, *binary-decomposability* (Section 6). We prove sufficiency for hypergraphs admitting an acyclic reduction order (Theorem 6.9), establish necessity up to a bipartite-matching condition (Theorem 6.10), and pin down the residual combinatorial gap as Conjecture 6.11, which we support with computational evidence. The criterion, and this combinatorial program, are to our knowledge new.

(3) *An experiment.* We construct a generative model of scenes in which a designated class of binding errors is provably invisible to the dyadic (clique-expansion) statistic T_{pair}, and show that the sum-type obstruction detects them perfectly in the noiseless limit and usefully under substantial noise, while T_{pair} sits at chance for every noise level, exactly as predicted (Section 7).

The novelty lies in the synthesis of Peircean same-carrier non-reduction, sheaf-theoretic global coherence, and an explicit distinction between same-carrier pairwise reduction and higher-arity cellular constraint, together with a combinatorial program for characterising when such constraints collapse to pairwise form. The synthesis is connective work on cited results; the non-reduction proposition specialises a known theorem; the characterization is new, its sufficiency proved in the acyclic case, its necessity proved up to an explicitly stated conjecture that we support computationally. We present the discourse complex as a linear model of one structural aspect of relational non-reduction, not as a definitive theory of feature binding or of Peircean Thirdness.

2 Background

We fix the sheaf-theoretic vocabulary and recall the threads the discourse complex unites.

Cellular sheaves and the sheaf Laplacian. Let X be a regular CW complex. A *cellular sheaf* $\mathcal{F}$ on X assigns to each cell σ a finite-dimensional real vector space $\mathcal{F}(\sigma)$, its *stalk*, and to each incidence $\tau \lhd \sigma$ a linear *restriction map* $\mathcal{F}_{\tau \lhd \sigma} \colon \mathcal{F}(\tau) \to \mathcal{F}(\sigma)$, closed under composition [3]. On a graph, the *coboundary* $\delta \colon C^0 \to C^1$ sends a vertex assignment x to $(\delta x)_e = \mathcal{F}_{u \lhd e} x_u - \mathcal{F}_{v \lhd e} x_v$ on each edge $e = \{u, v\}$; the *sheaf Laplacian* is $L = \delta^\top \delta$; and $H^0(\mathcal{F}) = \ker L = \ker \delta$ is the space of *global sections*, the assignments on which each edge's two restrictions agree [5]. Diffusion under L carries an arbitrary assignment toward H^0, the spectral gap sets the rate, and the first cohomology H^1 localises the obstructions to consistency.

Discourse sheaves. Hansen and Ghrist [6] introduced *discourse sheaves* to model opinion dynamics: agents hold opinions in stalks, edges encode the translations under which opinions count as compatible, and diffusion drives the network toward a consistent global opinion in H^0. This framing is the one we adopt, transposed from social opinion to perceptual binding.

Sheaf Laplacians on hypergraphs. To reach beyond pairwise relations, Duta et al. [4] attach a cellular sheaf to a hypergraph, with a restriction map $\mathcal{F}_{v \lhd e} \colon \mathcal{F}(v) \to \mathcal{F}(e)$ for each incident pair $v \lhd e$, and define a *linear sheaf hypergraph Laplacian* $\mathcal{L}_{\mathcal{F}}$ (their Definition 2). It is the operator of the *sheaf Dirichlet energy* (their Definition 3),

$$E^{\mathcal{F}}(x) \;=\; \tfrac{1}{2} \sum_{e} \tfrac{1}{\delta_e} \sum_{u,v \in e} \left\| \mathcal{F}_{v \lhd e} x_v - \mathcal{F}_{u \lhd e} x_u \right\|^2, \qquad \delta_e = |e|, \tag{1}$$

in which each hyperedge e contributes the sum, over pairs $u, v \in e$, of the squared disagreement of their restricted values (we display the unnormalised form; [4] additionally symmetric-normalise by node degrees). Its harmonic space $\ker \mathcal{L}_{\mathcal{F}} = \{x : E^{\mathcal{F}}(x) = 0\}$ is the locus on which, within every hyperedge, all restricted values coincide, $\mathcal{F}_{v \lhd e} x_v = \mathcal{F}_{u \lhd e} x_u$ for all $u, v \in e$. A hyperedge thus imposes the conjunction of all its pairwise agreements, so its consensus space is that of the clique expansion, pairwise by construction. This is the modelling choice we depart from.

Peirce's reduction thesis. Peirce held that triadic relations cannot in general be built from dyadic ones by composition and projection. After partial formalisations [7, 2], Koshkin [8] proved that a generic n-ary relation on an infinite domain has *ternarity* $n - 2$, so that a generic ternary relation is not the projection of any conjunction of binary relations on the same carrier. The argument is diagrammatic, in the tradition of bonding diagrams [1] and the string-diagram syntax for Peirce's beta graphs [9]. We will meet this theorem again, at $n = 3$, as the reason a single triadic cell resists pairwise expression.

What is new here. None of the four threads is by itself novel. New is their combination, a sheaf model of binding whose coboundary is chosen so that Peircean non-reduction has teeth, and the characterization of Section 6, which settles, for an arbitrary hypergraph of constraints, exactly when the pairwise default of [4] discards information.

3 The discourse complex

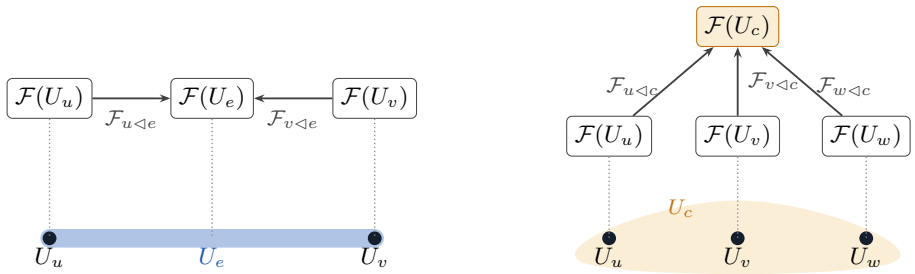
Definition 3.1 (Discourse complex)**.** A *discourse complex* is a tuple $\mathcal{D} = (A, \mathcal{C}, F, \varepsilon)$ where:

- A is a finite set of *agents*;
- F assigns a finite-dimensional real stalk $F(a)$ to each agent $a \in A$ and $F(c)$ to each cell c;
- $\mathcal{C} \subseteq 2^A$ is a set of *cells*, each a subset of A of size at least two;
- ε assigns to each incidence $a \lhd c$ (that is, each pair with $a \in c$) a sign $\varepsilon_{a \lhd c} \in \{-1, +1\}$ and a linear restriction map $F_{a \lhd c} \colon F(a) \to F(c)$.

Writing $C^0 = \bigoplus_{a \in A} F(a)$ and $C^1 = \bigoplus_{c \in \mathcal{C}} F(c)$, the *sum-type coboundary* $\delta \colon C^0 \to C^1$ is

$$(\delta x)_c \;=\; \sum_{a \in c} \varepsilon_{a \lhd c} \, F_{a \lhd c} \, x_a. \tag{2}$$

Cells of size two are *dyadic*, and cells of size at least three are *triadic* (or higher). The *percept space* is $P = H^0(\mathcal{D}) = \ker \delta$, and the *binding-failure statistic* is $T_{\mathrm{tri}}(x) = \|\delta x\|^2$.



(a) a dyadic cell (b) a triadic cell

Figure 1: A discourse complex is a cellular sheaf: each cell carries a stalk (a vector space of values) and each incidence a restriction map. *(a)* A dyadic cell U_e with vertices U_u, U_v is the ordinary graph-sheaf picture: the vertex stalks $\mathcal{F}(U_u), \mathcal{F}(U_v)$ map into the edge stalk $\mathcal{F}(U_e)$ through $\mathcal{F}_{u\lhd e}, \mathcal{F}_{v\lhd e}$, and consistency asks that the two images agree. *(b)* A triadic cell U_c over U_u, U_v, U_w: the three vertex stalks map into a single cell stalk $\mathcal{F}(U_c)$, and the sum-type coboundary (2) reads off one balance of the three restricted values, $(\delta x)_c = \sum_a \varepsilon_{a\lhd c}\mathcal{F}_{a\lhd c}x_a$, rather than the three pairwise agreements of the clique expansion (Remark 3.2).

The data of Definition 3.1 realise a two-level cellular sheaf on the incidence structure of the hypergraph $(A, \mathcal{C})$: agents are the 0-cells carrying degree-0 data, cells are the top cells carrying degree-1 constraints, with the incidences $a \lhd c$ and no others. Because there are no intermediate incidences, the compatibility $\delta^1\delta^0 = 0$ of a full cochain complex does not arise, and δ is simply the degree-zero *incidence* coboundary; $H^0(\mathcal{D}) = \ker\delta$ is its degree-zero cohomology, the global sections. We fix two pieces of notation used throughout:

$$V := C^0 = \bigoplus_{a\in A} F(a), \qquad R := \operatorname{row}(\delta) = (H^0)^\perp \subseteq V^*.$$

Here R, the row space of the coboundary, is spanned by the functionals $r_c := x \mapsto (\delta x)_c$, one for each cell; r_c is supported on the agents of c. The *dyadic subcomplex* $\mathcal{D}_{\text{pair}}$ retains only the dyadic cells of $\mathcal{D}$; its coboundary δ_{pair} is the ordinary graph sheaf coboundary of Section 2, and $T_{\text{pair}}(x) = \|\delta_{\text{pair}}x\|^2$ is the corresponding pairwise statistic.

Reading. Think of each agent as a feature population and its stalk as the space of values that population can take. A cell is a coherence relation among the populations it contains; its restriction maps say how each population's value is read into a common ledger $F(c)$, and its sign fixes the polarity of that reading. Equation (2) records the cell's verdict as a single quantity: the signed sum of the readings. The verdict is "consistent" precisely when that quantity vanishes. A percept is an assignment on which every cell returns the verdict "consistent"; a misbinding is one on which some cell does not, and T_{tri} totals the squared inconsistencies.

The single most consequential feature of Definition 3.1 is the form of (2): a cell contributes *one* signed sum of its incident restrictions, not the collection of its pairwise differences.

Remark 3.2 (Sum-type versus clique expansion). The coboundary (2) is the incidence, or parity-check, coboundary: it reads a hyperedge as a single higher cell carrying one balance constraint. The linear sheaf hypergraph Laplacian of [4], the operator of the energy (1), instead reads a hyperedge as a clique of pairwise agreements. On a dyadic cell, under the standard opposite-sign orientation $\varepsilon_{u\lhd e} = -\varepsilon_{v\lhd e}$, the sum-type constraint $\varepsilon_{u\lhd e}F_{u\lhd e}x_u + \varepsilon_{v\lhd e}F_{v\lhd e}x_v = 0$ is exactly the single disagreement $F_{u\lhd e}x_u = F_{v\lhd e}x_v$ of (1), so the two prescriptions coincide; with equal signs the dyadic sum-type constraint is instead $F_{u\lhd e}x_u = -F_{v\lhd e}x_v$, and orientation must be fixed for the two to agree. On a cell of size $k \geq 3$ they differ in kind (Figure 2): the clique-expansion consensus space $\ker\mathcal{L}_{\mathcal{F}}$ is the locus where *all* incident restrictions coincide, which is always pairwise-decomposable, whereas the sum-type consensus space is the single



sum-type: one balance
$\varepsilon_u x_u + \varepsilon_v x_v + \varepsilon_w x_w = 0$

clique: three agreements
$x_u = x_v = x_w$

Figure 2: Two readings of a triadic cell $\{u, v, w\}$ with scalar stalks and identity restrictions. *Left:* the sum-type coboundary imposes one balance, $x_u + x_v + x_w = 0$, a plane. *Right:* the clique expansion imposes all three pairwise agreements, forcing $x_u = x_v = x_w$, a line. The two consensus spaces differ in dimension and in kind.

hyperplane $\sum_{a \in c} \varepsilon_{a \lhd c} F_{a \lhd c} x_a = 0$, which need not be. Every result below turns on this difference: Proposition 4.2 shows the sum-type kernel can be genuinely triadic, and Section 6 characterises exactly when it is not.

The discourse complex furnishes a dictionary for binding. We *model* a coherently bound percept as a global section $x \in H^0$; the mathematical content is that a global section is a globally compatible assignment, and the perceptual reading is an additional modelling hypothesis. The *binding problem*, given detectors clamped to prescribed values u on a subset $U \subseteq A$, is then whether a global section extends the clamp. Writing $V_U = \{x \in \bigoplus_a F(a) : x_a = u_a \text{ for } a \in U\}$ for the affine set of assignments agreeing with the clamp, a bound percept consistent with U exists if and only if $V_U \cap \ker \delta \neq \varnothing$, an affine feasibility question; the diffusion L reaches such a section at a rate set by the spectral gap of L. A *misbinding* is an assignment with $\delta x \neq 0$, its severity measured by $T_{\text{tri}}(x)$. Finally, the dichotomy of Proposition 4.2 gives sheaf-theoretic content to a classical debate: a triadic constraint carried by an explicit triadic cell is *distributed* binding with no privileged locus, while the same constraint carried through an auxiliary hub is the localist convergence, or "grandmother," cell. The hub is a formal analogue of the mediating role associated with Peirce's hypostatic abstraction (Section 5); we do not claim it exhausts that notion.

4 Non-reduction

We record first that dyadic complexes express only pairwise constraints, and then that a single triad escapes them.

Lemma 4.1 (Dyadic complexes are pairwise)**.** *Let $\mathcal{D}$ be a discourse complex all of whose cells are dyadic. Then $H^0(\mathcal{D}) = \ker \delta$ is the intersection, over the edges $e = \{u, v\}$, of the consensus subspaces*

$$K_e = \{x : \varepsilon_{u \lhd e} F_{u \lhd e} x_u + \varepsilon_{v \lhd e} F_{v \lhd e} x_v = 0\},$$

each constraining only the two coordinates indexed by e. In particular H^0 is cut out by binary linear equations.

Proof. We have $\delta x = 0$ if and only if $(\delta x)_e = 0$ for every edge e. By (2), $(\delta x)_e = \varepsilon_{u \lhd e} F_{u \lhd e} x_u + \varepsilon_{v \lhd e} F_{v \lhd e} x_v$ involves only x_u and x_v; reading off components in the stalk $F(e)$ yields a finite list of equations, each in these two variables alone. □

Proposition 4.2 (Triadic non-reduction)**.** *Let $A = \{u, v, w\}$ with scalar stalks, and let $T = \{x \in \mathbb{R}^3 : x_u + x_v + x_w = 0\}$. Then:*

(i) *T is* not *the consensus space of any discourse complex on A whose cells are all dyadic;*



(ii) one triadic cell (iii) hub, dyadic only

Figure 3: The dichotomy of Proposition 4.2. The balance $x_u + x_v + x_w = 0$ is carried either by one distributed triadic cell (ii) or by an auxiliary hub h tied to u, v, w through dyadic cells (iii). The two have the same projected consensus space, but the hub is a fourth party, a posited convergence neuron, answering to nothing in the original three.

(ii) T is *the consensus space of a complex on A with a single triadic cell $\{u, v, w\}$, identity restrictions, and all signs positive;*

(iii) T is *the projection to A of the consensus space of a complex on $A \cup \{h\}$, with one auxiliary hub agent h and only dyadic cells.*

Proof. (i) Suppose $T = H^0(\mathcal{D})$ for a dyadic-only complex $\mathcal{D}$. By Lemma 4.1, H^0 is the kernel of a linear map each of whose rows is a *binary* functional on $\mathbb{R}^3$, one vanishing identically in some coordinate. As T has codimension one, this map has rank one, so $T = \ker \ell_1$ for a single binary functional ℓ_1. But $T = \ker \ell$ with $\ell(x) = x_u + x_v + x_w$, which depends on all three coordinates; equality of kernels forces ℓ_1 to be a scalar multiple of ℓ, contradicting that ℓ_1 is binary. This is Koshkin's theorem [8] at $n = 3$ in the present language.

(ii) With $c = \{u, v, w\}$, all signs positive, and identity restrictions, (2) gives $\delta = (\,1 \;\; 1 \;\; 1\,)$, whence $\ker \delta = T$.

(iii) Let $F(h) = \mathbb{R}^2$ with coordinates (a, b), and adjoin three dyadic cells $e_u = \{u, h\}$, $e_v = \{v, h\}$, $e_w = \{w, h\}$ with scalar cell stalks, identity agent-side restrictions, and hub-side restrictions

$$F_{h \lhd e_u}(a, b) = a, \qquad F_{h \lhd e_v}(a, b) = b, \qquad F_{h \lhd e_w}(a, b) = -a - b,$$

signs positive on agents and negative on the hub. The three constraints read $x_u = a$, $x_v = b$, $x_w = -a - b$, so $x_u + x_v + x_w = 0$ on the consensus space, and conversely every point of T is realised by $x_h = (x_u, x_v)$. Projecting away x_h recovers exactly T. □

The contrast between (ii) and (iii) is the substantive content, and it is worth dwelling on. A triadic constraint can be borne by one triadic cell, distributed across the three agents with no privileged party; or it can be borne by an auxiliary hub whose dyadic ties to the three agents jointly encode the same balance (Figure 3). The two models share a projected consensus space, but the pairwise model pays for it with an extra agent that corresponds to nothing in the original problem, a posited coordinator in a social network, a posited convergence neuron in a brain. The empirical question "is there a grandmother cell?" thereby sharpens into the structural question "is the binding constraint carried triadically, or hypostatically through a hub?"

The construction isolates a distinction we shall use throughout. The two systems in Proposition 4.2(ii) and (iii) are *extensionally equivalent*: they define the same ternary relation after projection. They are not *structurally* equivalent: one encodes the relation as a genuine triadic cell, the other only after enlarging the carrier with a mediator. Accordingly we call a relation *intrinsically triadic* when its constraint space is not generated by binary-supported functionals on the given carrier, as in Proposition 4.2(i), and *mediately triadic* when it becomes pairwise representable after the carrier is enlarged with an auxiliary variable, as in Proposition 4.2(iii). Every intrinsically triadic relation is mediately triadic, by the hub; the reduction thesis is that the converse fails, that a genuine triad cannot be made pairwise without the mediator. This

same-carrier-versus-mediated distinction, rather than the bare count of three variables, is the content of "non-reduction" here; we return to it in Section 5.

Example 4.3 (The smallest irreducible complex)**.** Take $A = \{u, v, w\}$, one triadic cell, scalar stalks, identity restrictions, positive signs. Then $\delta = (\,1\ 1\ 1\,)$, so $R = \mathrm{Span}\{(1,1,1)\}$ and $H^0 = T$ is the plane $x_u + x_v + x_w = 0$, of dimension two. The only vector in R supported on two or fewer agents is 0: any nonzero multiple of $(1,1,1)$ touches all three coordinates. We shall see in Section 6 that this single observation, no nonzero binary functional lives in R, is exactly what obstructs pairwise expression.

5 From Peircean relational graphs to discourse complexes

Peirce's existential graphs, in their Beta part, are a diagrammatic logic of relations. An individual is a *line of identity*, a heavy line on the sheet of assertion; a predicate is a *spot*, or rheme, of definite valency, whose hooks are filled by the lines of the individuals it relates; conjunction is juxtaposition, negation is enclosure by a cut, and existential quantification is carried by the lines themselves [11]. The reduction thesis lives in this syntax: teridentity, the three-way branching of a line, is the least genuinely triadic construction, and the negative half of the thesis is the claim that it cannot be assembled from dyadic identities alone [7, 2, 8]. The discourse complex of Section 3 is, in one phrase, a Beta graph with values: it keeps the relational skeleton, who is related to whom and with what valency, and adds at each place a space of admissible values and at each relation a linear law. This section makes the correspondence explicit, term by term, and shows that the hub of Proposition 4.2(iii) is exactly the graphical move of adjoining a mediating spot and then erasing its line by projection, the operation whose impossibility for genuine triads is Peirce's thesis. Figures 4 and 5 carry the two halves of the dictionary.

A feature population is a locus of possible interpretation. In the Beta graph a line of identity denotes an individual, and a hook is a place awaiting one. The discourse complex thickens the place into a space: a feature population a carries a stalk $F(a)$, and to interpret the locus is to choose a point of $F(a)$ (Figure 4). Where Peirce's line commits to a single individual, the stalk is the whole range of values the locus admits, the space of interpretations available there. A percept assigns one interpretation per locus; the stalk is what makes the assignment possible.

A cell is a relational predicate. A spot of valency k asserts a k-ary predicate of the individuals on its hooks. A cell c, with its restriction maps $\{F_{a \lhd c}\}$ and signs $\{\varepsilon_{a \lhd c}\}$, is its valued counterpart: it asserts of the interpretations at its correlates the single linear predicate $\sum_{a \in c} \varepsilon_{a \lhd c} F_{a \lhd c} x_a = 0$ of (2). Dyadic cells are dyadic predicates, the reactions of Secondness; cells of valency three and higher are the genuinely relational predicates of Thirdness. The Betagraphic calculus of [9, 1] supplies the bonding syntax; the sum-type coboundary supplies the law each bond enforces.

A global section is a coherent assignment. A model of a Beta graph assigns individuals to lines so that every scribed predicate holds. A global section $x \in H^0(\mathcal{D})$ assigns an interpretation to every locus so that every cell's predicate is satisfied, $\delta x = 0$. It is the coherent reading of the whole graph, the state of affairs in which all the asserted relations obtain at once (Figure 4). Perceptually it is a bound percept; logically it is a satisfying valuation; and the binding-failure statistic $T_{\mathrm{tri}}(x) = \|\delta x\|^2$ measures how far a candidate assignment falls short of coherence.

A triadic cell is a non-factorizable relation. Teridentity is Peirce's paradigm of a relation that no composition of dyads can build *on its own correlates*. Proposition 4.2(i) is its valued

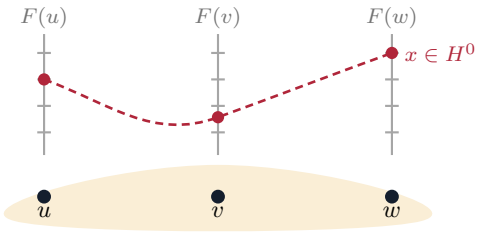


Figure 4: A feature population as a *locus of possible interpretation*, and a global section as a *coherent assignment*. Over each correlate a stands its stalk $F(a)$, the space of values (interpretations) the locus admits; the shaded region is a relation binding u, v, w. A global section $x \in H^0$ chooses one interpretation at each locus (dots) so that every relational predicate is satisfied, $\delta x = 0$.

form: the balance $x_u + x_v + x_w = 0$ is the consensus space of no dyadic complex on $\{u, v, w\}$, since a nonzero functional touching all three coordinates is not a scalar multiple of any binary one. A triadic cell is thus *intrinsically* triadic, an irreducible Third on its carrier, and not the conjunction of its pairwise shadows there (Figure 5, left; contrast the clique reading of Remark 3.2). It may still be *mediately* triadic, representable pairwise once the carrier is enlarged, which is the next item.

An auxiliary hub is an introduced mediating object. To carry a triadic predicate with dyadic ones, Proposition 4.2(iii) adjoins a new correlate h, tied to the three by dyadic cells, whose two-dimensional stalk holds the mediation. Graphically this is scribing a fresh spot on the sheet and running a new line of identity to it. This introduction of a mediating object is a formal analogue of the role played by Peirce's *hypostatic abstraction*, the operation that turns a predicate into a subject in its own right [10, 4.235]; we take the hub to be a mathematical counterpart of that mediating move, and do not claim that it exhausts Peirce's richer notion. The hub answers to nothing in the original three: it is the logician's introduced mediator and the neuroscientist's convergence, or "grandmother," cell (Figure 5, right).

Projecting the hub away is relational reduction. The original triad is recovered from the hub construction by projecting away the hub's stalk, that is, by existentially quantifying the introduced object out,

$$R(u, v, w) \;\equiv\; \exists h \left[\, D_1(u, h) \wedge D_2(v, h) \wedge D_3(w, h) \,\right],$$

with D_1, D_2, D_3 the three dyadic cells (Figure 5, arrow). This manoeuvre, introduce a mediator and then eliminate it by projection, is precisely how relational logic attempts to reduce a relation to lower-valency ones, and the reduction thesis, from Herzberger through Burch to Koshkin [7, 2, 8], is the theorem that for a genuine triad it cannot be completed without the posited h. Proposition 4.2 gives the statement in the valued setting: the hub is necessary, and its elimination is a projection, never a factorization. The reduction is real, but only as a reduction *to a triad plus an existential*, which is Peirce's positive thesis, and not a reduction of the triad away, which his negative thesis forbids.

Two payoffs follow from reading the discourse complex as a Beta graph with values. First, the relational ontology is inherited: loci, predicates, coherent assignments, and the same-carrier irreducibility of the triad are read from Peirce's own framework. Second, the added metric layer lets us do what the bare graph cannot, namely *measure* the failure of a relation through T_{tri}, and, for a whole graph at once, ask combinatorially which of its triads are intrinsically triadic and

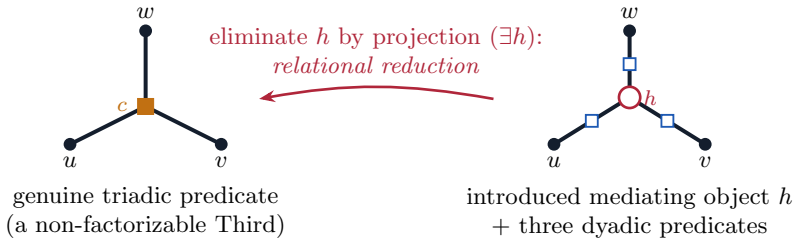


Figure 5: A triadic cell as a *non-factorizable relation*, and the hub as an *introduced mediating object* whose *projection* is relational reduction. *Left:* a genuine triadic predicate c binds u, v, w in one act. *Right:* the dyadic surrogate posits a new object h, tied to u, v, w by three dyadic predicates. Eliminating h by projection (existential quantification) recovers the triad; Peirce's reduction thesis (Proposition 4.2) is that for a genuine Third this positing of h is unavoidable.

which are only mediately so. That question is the subject of Section 6: binary-decomposability (Definition 6.5) is offered as a *conjectured* combinatorial test for when every triad in a fabric of relations is mediately triadic (pairwise after enlarging the carrier), leaving no triad that is intrinsically triadic on its own carrier. A mediator always dissolves a single triad; the content is whether a mediation is intrinsic to the given carrier or must be introduced from outside it.

6 When are pairwise constraints enough?

Proposition 4.2 shows a lone triad is irreducible on its carrier. But a discourse complex is a whole hypergraph of interacting cells, and the interactions can conspire either to preserve same-carrier triadic content or to dissolve it. The central question of this paper is when they dissolve it: for which hypergraphs is the global-section space expressible, on the same carrier, by pairwise constraints? We propose a condition on the bare cell-incidence structure and conjecture that it answers the question for generic scalar-stalk data (Conjecture 6.11), proving one direction outright and the other up to a matching condition. We work throughout this section with *scalar* stalks, so that each restriction $F_{a \lhd c}$ is a nonzero scalar and the generic-rank machinery below applies. The extension to vector stalks with matrix-valued restrictions, where a restriction may mix coordinates within a stalk and so cannot be split into independent scalar agents, requires a block-supported analogue of the binary-support and matroid arguments used here, and is left to future work (Section 9).

6.1 Genericity, made precise

Fix the hypergraph $(A, \mathcal{C})$ and the signs ε. With scalar stalks the remaining data are the restriction scalars $r = (r_{a \lhd c})$ ranging over $(\mathbb{R}^\times)^E$, where $E = \{(a, c) : a \in c\}$ is the incidence set. A property holds for *Zariski-generic* r if it holds on the complement of the zero set of some nonzero polynomial in the $r_{a \lhd c}$, in particular on a dense set of full Lebesgue measure. Genericity here refers only to the continuous restriction scalars: the incidence pattern $(A, \mathcal{C})$ and the sign pattern ε are fixed and discrete, and are never varied, so a phrase such as "generic signs" would be meaningless. All statements below are of the form "for each fixed incidence and sign pattern, the asserted property holds for Zariski-generic nonzero restriction scalars." We use one standard fact about such genericity, which converts every rank we need into a combinatorial quantity.

Lemma 6.1 (Generic rank equals term rank)**.** *Let $M(r)$ be a matrix whose entries are either zero or one of the algebraically independent scalars $r_{a \lhd c}$ (up to the fixed signs), each variable occurring in a single prescribed position. Then for Zariski-generic r,*

$$\operatorname{rank} M(r) \;=\; \textit{term rank of the support pattern of } M,$$

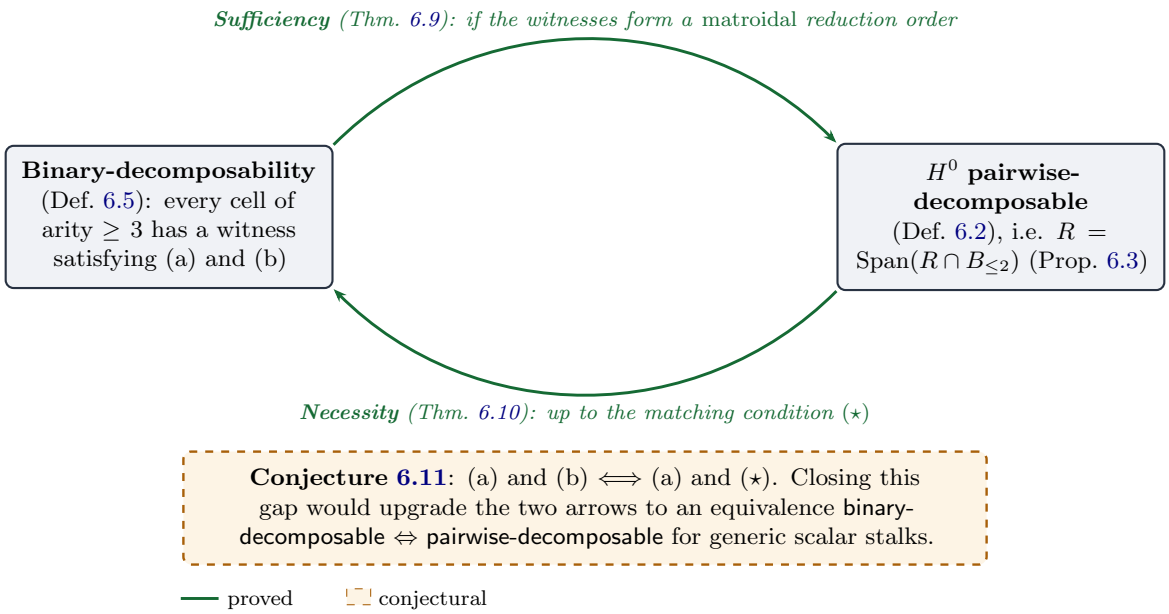


Figure 6: Map of this section. The two boxes are the combinatorial criterion and the cohomological property it is meant to capture; the proved implications run between them, and the one remaining gap is isolated in the conjecture below.

the maximum number of nonzero positions no two sharing a row or a column, equivalently, the maximum matching in the bipartite graph of the pattern.

Proof sketch. A matching of size t selects t nonzero positions in distinct rows and columns; the corresponding $t \times t$ minor has, among the terms of its Leibniz expansion, the monomial $\pm \prod r_{a \lhd c}$ over the matched positions, which no other term can cancel because each variable occurs once. So that minor is a nonzero polynomial and is nonzero for generic r; hence $\operatorname{rank} M \ge t$. Conversely every nonzero minor selects positions forming a matching, so $\operatorname{rank} M \le$ term rank. Both bounds are attained generically. □

Because distinct incidences carry algebraically independent scalars, Lemma 6.1 applies to δ and to every submatrix we form from its rows. In particular the rows $\{r_c\}$ are supported exactly on their cells, and the generic dimension of any space we build from them is a matching number.

6.2 The right notion of pairwise expressibility

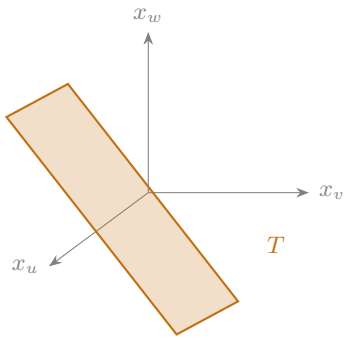
For $S \subseteq A$ let $V_S^* \subseteq V^*$ be the functionals supported on S, those factoring through the projection $\pi_S \colon V \to \bigoplus_{a \in S} F(a)$, and let $B_{\le 2} = \bigcup_{|S| \le 2} V_S^*$ be the functionals supported on at most two agents. Note that $B_{\le 2}$ is a union of subspaces, not itself a subspace.

Definition 6.2 (Pairwise-decomposable)**.** The global-section space $H^0(\mathcal{D})$ is *pairwise-decomposable*, PD for short, if

$$H^0 \;=\; \bigcap_{|S|=2} \pi_S^{-1}\big(\pi_S(H^0)\big);$$

that is, H^0 is determined by, and recoverable from, its projections to pairs of agents.

Reading. The right-hand side is the smallest set containing H^0 that is "visible pairwise": an assignment x lies in it exactly when, for every pair of agents, the two coordinates x shows there could belong to some genuine global section. Pairwise-decomposability says this pairwise shadow reconstructs H^0 exactly: nothing is lost by looking two agents at a time. Proposition 4.2 already



each coordinate shadow of T fills its plane,
so the pairwise closure of T is all of $\mathbb{R}^3 \neq T$

Figure 7: The plane $T = \ker(x_u + x_v + x_w)$ projects onto every coordinate plane surjectively. Its pairwise closure is therefore all of $\mathbb{R}^3$, strictly larger than T: the triad is not pairwise-decomposable.

shows the triad T fails this: its shadow on every coordinate pair is the whole plane $\mathbb{R}^2$ (Figure 7), so the reconstruction is all of $\mathbb{R}^3$, strictly larger than T.

The following reformulation is the working form of the definition, and the conceptual anchor of this section: it translates the geometric question about H^0 into a clean linear-algebraic one about the constraint space R, and every combinatorial result below is an analysis of when R is spanned by its binary-supported functionals. It also disposes of a tempting but vacuous variant.

Proposition 6.3 (Dual form)**.** *$H^0(\mathcal{D})$ is pairwise-decomposable if and only if $R = \mathrm{Span}(R \cap B_{\leq 2})$; that is, if and only if the row space of the coboundary is spanned by the functionals it contains that are supported on at most two agents. Equivalently, a global-section constraint system is pairwise expressible exactly when its constraint space is generated by binary-supported functionals.*

Proof. Suppose $H^0 = \bigcap_i K_i$ with $K_i = \pi_{S_i}^{-1}(L_i)$ and $|S_i| \leq 2$. Taking orthogonal complements in the finite-dimensional V, $R = (H^0)^\perp = \sum_i K_i^\perp$, and each $K_i^\perp$ lies in $V_{S_i}^* \cap R \subseteq B_{\leq 2} \cap R$; hence $R = \mathrm{Span}(R \cap B_{\leq 2})$. Conversely, if R has a basis $b_1, \ldots, b_m$ with each $b_i \in R \cap V_{S_i}^*$ and $|S_i| \leq 2$, put $K_i = \ker b_i = \pi_{S_i}^{-1}\big(\ker b_i|_{V_{S_i}}\big)$; then $H^0 = R^\perp = \bigcap_i \ker b_i = \bigcap_i K_i$, which is Definition 6.2. □

Remark 6.4 (A definition that looks correct but is empty). One might instead ask that $R \subseteq \mathrm{Span}(B_{\leq 2})$. This holds for *every* subspace, because the binary functionals already span all of V^*: a functional supported on a single agent is binary in the degenerate case $|S| = 1$, and these span V^*. The content of pairwise-decomposability is that the binary functionals doing the spanning must themselves lie in R, precisely the condition $R = \mathrm{Span}(R \cap B_{\leq 2})$. Example 4.3 is the check: there $R = \mathrm{Span}\{(1,1,1)\}$ meets $B_{\leq 2}$ only in 0, so $\mathrm{Span}(R \cap B_{\leq 2}) = \{0\} \neq R$, and T is not PD, as it must be.

6.3 Two plausible criteria that fail

Before the criterion that works, two that do not; each rules out a natural guess and sharpens what the correct condition must accommodate.

The first guess is that pairwise-decomposability holds exactly when every cell is dyadic. It fails. On four agents the three triadic cells $\{0,1,2\}, \{1,2,3\}, \{0,2,3\}$ leave $\dim H^0 = 1$ generically (three independent constraints in $\mathbb{R}^4$), and a one-dimensional subspace is always pairwise-decomposable, since a line is recovered from its coordinate projections. Triadic cells

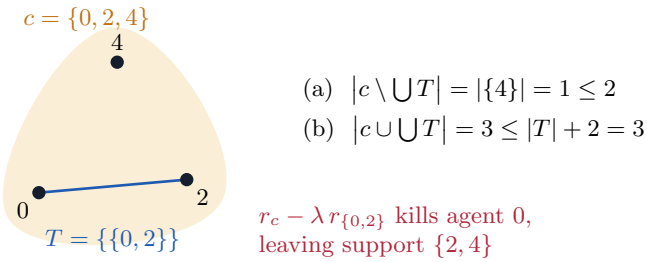


Figure 8: Binary-reducibility of the triadic cell $c = \{0, 2, 4\}$ via the single dyadic witness $T = \{\{0, 2\}\}$. Conditions (a) and (b) both hold, and concretely a multiple of the dyadic row cancels agent 0 from the triadic row, leaving a binary functional on $\{2, 4\}$.

can therefore coexist with pairwise-decomposability when there are enough of them to shrink the section space.

The second guess is that each triadic cell must contain at least $|c| - 2$ proper subcells covering it. It too fails. The complex $\{0, 2\}, \{0, 2, 4\}$ has a single dyadic subcell of its triadic cell, yet is pairwise-decomposable: subtracting a suitable multiple of the dyadic row from the triadic row cancels agent 0 and leaves a functional supported on $\{2, 4\}$. A lone dyadic subcell suffices. The correct invariant must be subtler than a subcell count.

6.4 Binary-decomposability

Definition 6.5 (Binary-reducible cell; binary-decomposable hypergraph)**.** A cell $c \in \mathcal{C}$ is *binary-reducible* if $|c| \leq 2$, or if there is a set $T \subseteq \mathcal{C} \setminus \{c\}$ of other cells satisfying both

$$\text{(a)} \quad \Big|c \setminus \bigcup T\Big| \leq 2 \qquad \text{and} \qquad \text{(b)} \quad \Big|c \cup \bigcup T\Big| \leq |T| + 2.$$

The hypergraph $(A, \mathcal{C})$ is *binary-decomposable* if every cell of arity at least three is binary-reducible.

Both conditions read off the bare cell supports; neither mentions stalks, signs, or restriction values. Condition (a) is a *covering* requirement, all but at most two agents of c must be touched by the other cells in T. Condition (b) is a *deficiency* bound, the sub-hypergraph $T \cup \{c\}$ must span at most two more agents than it has cells. The two are independent (Example 6.14). The intuition, made exact below, is that (a) leaves at most two "free" agents of c to survive the cancellation, while (b) guarantees enough cells to cancel the rest.

Example 6.6 (Reducible: the dyadic-in-triadic complex)**.** Consider the complex on $A = \{0, 2, 4\}$ with the two cells $\{0, 2\}$ and $\{0, 2, 4\}$, scalar stalks, and identity restrictions, giving rows $r_{\{0,2\}} = (1, -1, 0)$ and $r_{\{0,2,4\}} = (1, 1, -1)$ in coordinates (x_0, x_2, x_4). The triadic cell is binary-reducible with $T = \{\{0, 2\}\}$: condition (a) gives $|c \setminus \{0, 2\}| = |\{4\}| = 1 \leq 2$, and (b) gives $|c \cup \{0, 2\}| = 3 \leq |T| + 2 = 3$ (Figure 8). Concretely,

$$r_{\{0,2,4\}} - r_{\{0,2\}} = (0, 2, -1),$$

supported on $\{2, 4\}$, is a binary functional in R. Together with $r_{\{0,2\}}$ it spans the two-dimensional R, so $R = \mathrm{Span}(R \cap B_{\leq 2})$ and H^0 is pairwise-decomposable by Proposition 6.3. The whole hypergraph is binary-decomposable.

Example 6.7 (Irreducible: the lone triad)**.** The single triadic cell of Example 4.3 has no other cells to serve as a witness, so the only candidate is $T = \varnothing$, for which (a) reads $|c \setminus \varnothing| = 3 > 2$ and fails. The cell is not binary-reducible, the hypergraph is not binary-decomposable, and $H^0 = T$ is not PD, consistent with Proposition 4.2(i).

6.5 The characterization

Our main results bracket the equivalence "binary-decomposable $\iff$ pairwise-decomposable." We prove sufficiency under an acyclicity hypothesis together with a matching condition, and necessity up to the same matching condition; we then isolate the exact residual gap. Two notions make the hypotheses precise.

Definition 6.8 (Reduction order; matroidal witness)**.** A *reduction order* for $(A, \mathcal{C})$ is an ordering $c^{(1)}, \dots, c^{(k)}$ of the cells of arity ≥ 3 such that each $c^{(j)}$ has a witness $T_j \subseteq \{\text{dyadic cells}\} \cup \{c^{(1)}, \dots, c^{(j-1)}\}$ satisfying (a) and (b) of Definition 6.5. Such a witness is *matroidal* if there is a set $K_j \subseteq \bigcup T_j$ with $|K_j| = |c^{(j)} \cup \bigcup T_j| - 2$, containing no agent of $c^{(j)} \setminus \bigcup T_j$, for which the matching identity $(\star)$ below holds with $(c, T, K) = (c^{(j)}, T_j, K_j)$.

Informally, a reduction order lets every higher cell be reduced using only dyadic cells and cells already reduced, with no circular dependence; matroidality is the extra incidence condition ensuring the reduction can actually be carried out over the values, not merely counted. Most binary-decomposable hypergraphs admit a reduction order; the smallest that does not is the triangle of triads (Example 6.13).

Theorem 6.9 (Sufficiency)**.** *Fix the incidence and sign patterns. If $(A, \mathcal{C})$ admits a reduction order all of whose witnesses are matroidal, then for Zariski-generic restriction scalars $H^0(\mathcal{D})$ is pairwise-decomposable.*

Proof. Fix generic r; by Lemma 6.1 all ranks below are matching numbers. Let $W = \mathrm{Span}(R \cap B_{\leq 2})$. We show $r_c \in W$ for every cell c; then $W = R$, which is pairwise-decomposability by Proposition 6.3. Dyadic rows are supported on two agents, hence lie in $R \cap B_{\leq 2} \subseteq W$. Treat the higher cells in the reduction order $c^{(1)}, \dots, c^{(k)}$ and induct on j that $r_{c^{(j)}} \in W$. Write $c = c^{(j)}$, let $T = T_j$ be its witness and $K = K_j$ the set witnessing matroidality, and put $U = c \cup \bigcup T$, $m = |T|$.

A binary functional through c. Requiring a combination $b = \beta r_c + \sum_{c_i \in T} \alpha_i r_{c_i}$ to vanish on the $|U| - 2$ agents of K is a homogeneous linear system $M\mu = 0$ in the $m + 1$ unknowns $\mu = (\beta, \alpha)$, where M has rows indexed by K and columns by $T \cup \{c\}$, with entry $\varepsilon r_{a \lhd c'}$ when $a \in c'$ and 0 otherwise. A solution with $\beta \neq 0$ exists precisely when the column of c lies in the span of the columns of T; by Lemma 6.1 this is the matching identity

$$\nu(K, T) = \nu\big(K, T \cup \{c\}\big), \tag{$\star$}$$

where $\nu(K, X)$ is the maximum matching in the incidence graph on $K \times X$, and $(\star)$ holds because the witness is matroidal. Since K avoids $c \setminus \bigcup T$, the two free agents $U \setminus K$ contain $c \setminus \bigcup T$, so b is supported on at most two agents; thus $b \in R \cap B_{\leq 2} \subseteq W$ with $\beta \neq 0$.

Descent. From $\beta \neq 0$,

$$r_c = \tfrac{1}{\beta}\, b - \tfrac{1}{\beta} \sum_{c_i \in T} \alpha_i\, r_{c_i}.$$

Each $c_i \in T$ is either dyadic, whence $r_{c_i} \in W$, or an earlier higher cell $c^{(j')}$ with $j' < j$, whence $r_{c_i} \in W$ by the inductive hypothesis. Therefore $r_c \in W$, completing the induction. □

Theorem 6.9 assumes matroidal witnesses. That conditions (a) and (b) always furnish one, so that a plain reduction order is automatically matroidal, is one half of Conjecture 6.11 below; we verify it on several hundred instances (Section 6.6) but do not prove it. Under that conjecture, Theorem 6.9 yields pairwise-decomposability for every binary-decomposable hypergraph admitting a reduction order.

Theorem 6.10 (Necessity, up to matching)**.** *Suppose $H^0(\mathcal{D})$ is pairwise-decomposable for all restrictions in a Zariski-open dense set. Then for every cell c with $|c| \geq 3$ there exist $T \subseteq \mathcal{C} \setminus \{c\}$ and $K \subseteq \bigcup T$ with $|K| = |c \cup \bigcup T| - 2$ satisfying condition* (a) *of Definition 6.5 and the matching identity $(\star)$.*

Proof. Fix generic r for which H^0 is PD and a cell c with $|c| \geq 3$. By Proposition 6.3, $r_c = \sum_i \lambda_i b_i$ with each $b_i \in R \cap B_{\leq 2}$ and $\lambda_i \neq 0$. The cell-rows need not be linearly independent, so a representation of a vector in the row space in terms of the rows is not unique; to extract a coefficient cleanly we pass to a basis. Since $r_c \neq 0$ for generic r, choose a basis $\mathcal{B}$ of R consisting of cell-rows and containing r_c (any spanning set of vectors contains a basis, and one may be chosen through the nonzero vector r_c). Expand each b_i in this basis, $b_i = \sum_{c' \in \mathcal{B}} \mu_{i,c'} r_{c'}$, where the coefficients are now unique. As r_c is itself a basis vector, matching its coefficient in $r_c = \sum_i \lambda_i b_i$ gives $\sum_i \lambda_i \mu_{i,c} = 1$, so some term has $\mu_{i,c} \neq 0$. Fix that i, drop the index, and set $T = \{c' \in \mathcal{B} \setminus \{c\} : \mu_{c'} \neq 0\}$, $U = c \cup \bigcup T$, and $K = U \setminus \mathrm{supp}(b)$, so $|K| = |U| - 2$. By construction the rows $\{r_{c'} : c' \in T \cup \{c\}\}$ are linearly independent, so the term-rank computations below are unambiguous.

For condition (a): if an agent $a^* \in K$ lay in $c \setminus \bigcup T$, the equation $b(a^*) = 0$ would read $\mu_c \, r_{a^* \lhd c} = 0$, forcing $\mu_c = 0$ for generic $r_{a^* \lhd c} \neq 0$, a contradiction. Hence $c \setminus \bigcup T \subseteq \mathrm{supp}(b)$, whence $|c \setminus \bigcup T| \leq 2$ and, in passing, $K \subseteq \bigcup T$.

For ($\star$): the requirement $b(a) = 0$ for all $a \in K$ is the system $M\mu = 0$ of the previous proof, and the existing solution has $\mu_c \neq 0$, so the column of c lies in the span of the columns of T. By Lemma 6.1 this is exactly $\nu(K, T) = \nu(K, T \cup \{c\})$. □

Theorem 6.10 extracts from pairwise-decomposability the covering condition (a) and the matching identity ($\star$). The identity is condition (b) in disguise except in one bipartite-matroid corner; closing that corner, in both directions, is the sole obstruction to a clean two-sided equivalence, and is the content of Conjecture 6.11.

Conjecture 6.11 (Combinatorial closure)**.** *Fix a cell c and a set $T \subseteq \mathcal{C} \setminus \{c\}$. There is a set $K \subseteq \bigcup T$ with $|K| = |c \cup \bigcup T| - 2$ avoiding $c \setminus \bigcup T$ and satisfying* (a) *and* ($\star$) if and only if *some $T' \subseteq \mathcal{C} \setminus \{c\}$ satisfies* (a) *and* (b)*. Consequently, together with Theorems 6.9–6.10: for generic scalar-stalk data on a hypergraph admitting a reduction order, $H^0(\mathcal{D})$ is pairwise-decomposable if and only if $(A, \mathcal{C})$ is binary-decomposable.*

Remark 6.12 (Why ($\star$) and (a) appear to force (b)). Condition (b) is equivalent to $|K| \leq |T|$. Suppose instead $|K| > |T|$. Then the incidence graph $G(K, T)$ has a Hall violator: some $a^* \in K$ is left unmatched by every maximum T-matching. For ($\star$) to hold, no augmenting path may run from a^* to c; in particular $a^* \notin c$, and every cell of T meeting a^* must also meet a matched vertex. But condition (a) places all but at most two agents of c inside $\bigcup T$, so c is adjacent to much of the alternating structure rooted at a^*, which generically produces exactly the forbidden path. Turning this heuristic into a proof, a Kőnig–Egerváry argument on the incidence graph of $(K, T \cup \{c\})$, is the content of the conjecture.

Example 6.13 (The circular case: a triangle of triads)**.** On $A = \{0, 1, 2, 3\}$ take $\mathcal{C} = \{\{0, 1, 2\}, \{1, 2, 3\}, \{0, 2, 3\}\}$. Each cell is binary-reducible: for $c = \{0, 1, 2\}$ the witness $T = \{\{1, 2, 3\}, \{0, 2, 3\}\}$ has $\bigcup T = \{0, 1, 2, 3\} \supseteq c$, so (a) gives $|c \setminus \bigcup T| = 0$ and (b) gives $|c \cup \bigcup T| = 4 \leq |T| + 2 = 4$; the other two cells are symmetric. So $(A, \mathcal{C})$ is binary-decomposable. But no reduction order exists: reducing any one cell needs the other two, which have not yet been reduced; the dependence is genuinely circular, and Theorem 6.9 does not apply. Nonetheless H^0 *is* pairwise-decomposable, for a different reason: here $\dim R = 3$ and $\dim H^0 = 1$, and R, being the orthogonal complement of a generic line, is spanned by the six one-dimensional binary subspaces $R \cap V_S^*$ (one per pair S), whose span is generically all of R. This is the mechanism Theorem 6.9 does not capture, and it is why the general sufficiency statement, every binary-decomposable hypergraph is PD, remains part of Conjecture 6.11 rather than a theorem.

6.6 Verification

We tested the whole circle of statements computationally with independent tools; source and seeds accompany the manuscript. A first tool computes $\dim R$ and the *binary rank* $\dim \mathrm{Span}(R \cap$

instance	$\lvert A\rvert$	$\lvert\mathcal{C}\rvert$	binary-decomp.	empirical PD	agree
dyadic path	4	3	yes	40/40	✓
single triad	3	1	no	0/40	✓
two disjoint triads	6	2	no	0/40	✓
triangle of triads	4	3	yes	40/40	✓
dyadic-in-triadic	5	2	yes	40/40	✓
four-ary with triadic subcell	5	2	no	0/40	✓
adversarial disconnect	8	5	no	0/40	✓
120 random hypergraphs	3–8	1–9	mixed	120/120	✓

Table 1: Combinatorial verdict of Definition 6.5 against the numerical test for pairwise-decomposability (PD). "Empirical PD" reports, over random draws of restrictions, how often H^0 was numerically pairwise-decomposable.

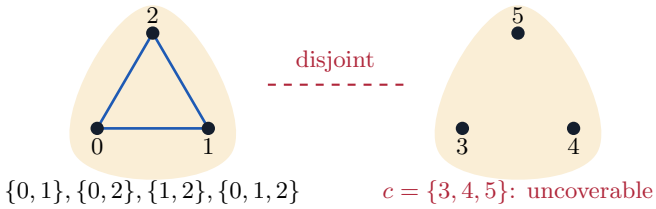


Figure 9: The adversarial disconnect of Example 6.14. The triadic cell $\{3,4,5\}$ shares no agent with the rest of the complex, so condition (a) fails for every witness. The hypergraph is not binary-decomposable and H^0 is not PD, yet condition (b) alone is attainable, which is why both conditions are needed.

$B_{\leq 2}$) numerically from a coboundary matrix, with agent-aware supports, declaring pairwise-decomposability exactly when the two agree. A second evaluates the combinatorial criterion of Definition 6.5 on cell supports. Across a curated suite of eighteen cases, single and disjoint triads, the triangle of triads, the dyadic-in-triadic complex, four-ary cells, and the adversarial disconnect of Example 6.14, and one hundred twenty random hypergraphs with $|A| \in [3,8]$ and arities in $\{2,3,4\}$, the combinatorial verdict agreed with the numerical one in all 138 instances (Table 1).

Two targeted audits address the gaps in the proofs. For local reducibility (that (a) and (b) yield a matroidal witness, one half of Conjecture 6.11), over 731 randomly generated cells satisfying (a) and (b), an admissible K with $(\star)$ was found in every case. For the sufficiency mechanism, of 510 binary-decomposable hypergraphs, 458 admitted a reduction order and the remaining 52 were only-circular; all 52, the triangle of triads among them, were nonetheless numerically pairwise-decomposable, consistent with Conjecture 6.11. A third tool searched, over 234 random PD hypergraphs and all admissible (T,K), for a witness satisfying $(\star)$ but not (b), and found none. These computational experiments provide supporting evidence for Conjecture 6.11 but do not constitute a proof; in total the characterization is consistent with every one of several hundred computationally tested instances, and the two places where a proof is missing are exactly the two halves of the conjecture.

Example 6.14 (Both conditions are needed: the adversarial disconnect)**.** Let $A = \{0,\dots,7\}$ and

$$\mathcal{C} = \{\{0,1\},\ \{0,2\},\ \{1,2\},\ \{0,1,2\},\ \{3,4,5\}\}.$$

For the cell $c = \{3,4,5\}$ every other cell lies within $\{0,1,2\}$, so $\bigcup T \subseteq \{0,1,2\}$ for any candidate T and $|c \setminus \bigcup T| = 3 > 2$: condition (a) fails for all T (Figure 9). The hypergraph is not binary-decomposable, and indeed H^0 is not PD (0/40 empirically). What makes the example instructive is that condition (b) *is* attainable here: taking T to be the four cells on $\{0,1,2\}$

gives $|c \cup \bigcup T| = 6 = |T| + 2$, so (b) alone would wrongly certify reducibility. Both conditions of Definition 6.5 are genuinely independent, and both are needed.

7 An experiment: detecting binding failure

The theorems of Section 6 concern which subspaces a discourse complex can express. We turn to a concrete consequence: whether the sum-type statistic T_{tri} can *detect* binding errors that the dyadic statistic T_{pair} must miss. Here $T_{\text{pair}}(x) = \|\delta_{\text{pair}} x\|^2$ is the statistic a model built from the dyadic sub-complex would use, the value-theoretic counterpart of the clique-expansion Laplacian of [4]. We do *not* claim invisibility to every function of pairwise data; the scope of the claim is exactly T_{pair}, and we make this precise in Remark 7.2 below. To keep the claim falsifiable rather than rhetorical, we build a generative model in which a designated class of errors is provably invisible to T_{pair}, and measure how well T_{tri} recovers them.

7.1 The generative model

We use twelve scalar feature channels. Four dyadic cells $\{0,1\}, \{2,3\}, \{4,5\}, \{6,7\}$ impose $x_a - x_b = 0$; four triadic cells $\{0,2,8\}, \{4,6,9\}, \{1,5,10\}, \{3,7,11\}$ impose $x_a + x_b - x_c = 0$. Channels 8 through 11 are *integrator* channels: they occur only in triadic cells, so the pairwise coboundary δ_{pair} has identically zero columns there. The full coboundary δ is an 8×12 matrix of rank 8, so the percept space $P = \ker \delta$ has dimension four.

7.2 Scene classes

At each noise level σ we draw two thousand scenes from each of three classes. A *bound* scene is $x = p + \sigma\varepsilon$ for a percept $p \in P$ and standard Gaussian ε. A T_{pair}-*visible* misbinding swaps one channel in $\{0, \ldots, 7\}$ between two independent percepts, breaking a dyadic cell and so raising T_{pair}. A T_{pair}-*blind* misbinding swaps one channel in $\{8, \ldots, 11\}$, breaking only a triadic cell; because δ_{pair} never touches these channels, T_{pair} cannot distinguish such a scene from a bound one.

7.3 A prediction and its proof

The design makes one prediction exactly, not merely statistically.

Proposition 7.1 (Pairwise blindness)**.** *For the model above, at every noise level σ the pairwise statistic T_{pair} has area under the ROC curve equal to $\frac{1}{2}$ in expectation for separating bound scenes from T_{pair}-blind misbindings.*

Proof. A T_{pair}-blind misbinding alters only channels 8 through 11. Since δ_{pair} has zero columns on those channels, $T_{\text{pair}}(x) = \|\delta_{\text{pair}} x\|^2$ depends only on $(x_0, \ldots, x_7)$, whose joint law is identical under the bound and T_{pair}-blind constructions. The two distributions of T_{pair} therefore coincide, and no threshold on T_{pair} separates them better than chance. □

Remark 7.2 (Scope of the blindness). The blindness is a property of T_{pair}, not of pairwise information in general. A T_{pair}-blind swap does change some *pairwise* marginals of the scene: in a bound percept a triadic cell $\{a, b, c\}$ makes x_a and x_c dependent, and the swap on the integrator channel c replaces x_c by an independent draw, so the joint law of the pair (x_a, x_c) differs between the two classes. A statistic that inspected such a pair could therefore detect the swap. What T_{pair} cannot do is form that pair: the integrator channels 8 through 11 appear in no dyadic cell, so T_{pair}, built from the dyadic sub-complex, never observes them. The experiment thus shows that a model carrying only the dyadic (clique-expansion) constraints is blind to violations of the triadic ones, which is exactly the sum-type-versus-clique distinction of Remark 3.2 made observable; it is not a claim that no function of pairwise data can separate the classes.

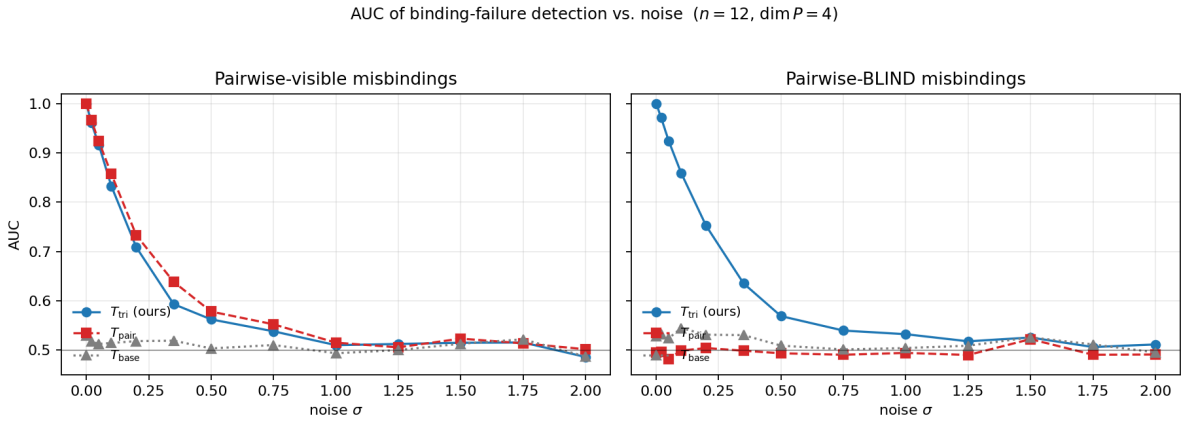


Figure 10: Discrimination (AUC) versus noise. *Left:* T_{pair}-visible misbindings, where the triadic (T_{tri}) and pairwise (T_{pair}) statistics track one another. *Right:* T_{pair}-blind misbindings, where T_{pair} remains at chance for every noise level while T_{tri} is perfect at $\sigma = 0$ and stays above 0.85 through $\sigma = 0.10$. The variance baseline T_{base} is uninformative in both panels.

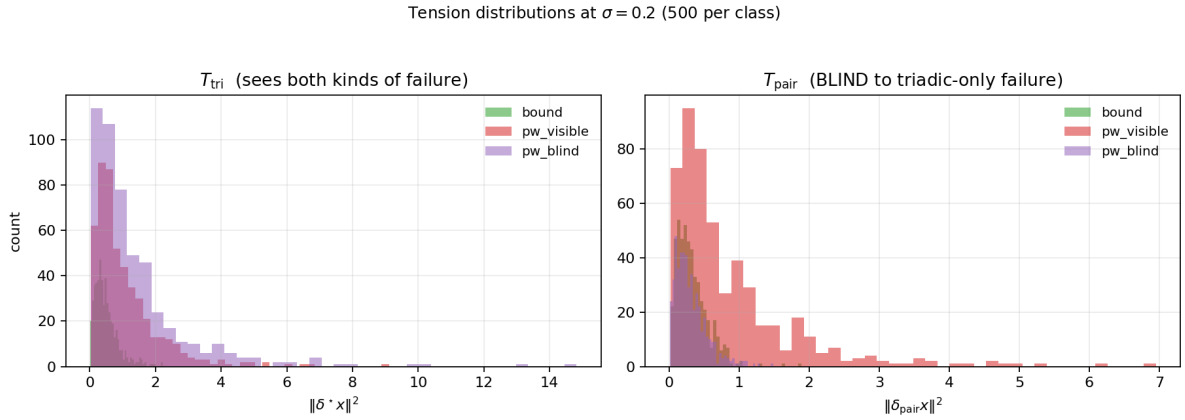


Figure 11: Tension distributions at $\sigma = 0.20$. *Left:* T_{tri} separates both misbinding classes from bound scenes. *Right:* T_{pair} separates only the T_{pair}-visible class; the bound and T_{pair}-blind distributions are essentially identical, illustrating Proposition 7.1.

7.4 Results

Figure 10 plots the area under the ROC curve (AUC) for each statistic against noise, and Figure 11 shows the underlying tension distributions at a representative level; Table 2 gives the numbers. On T_{pair}-visible errors the triadic and dyadic statistics behave almost identically, both decaying gracefully from perfect discrimination toward chance as noise grows. On T_{pair}-blind errors they part company completely: T_{pair} sits at 0.50 for every noise level, as Proposition 7.1 requires (mean 0.496, standard deviation 0.010, range $[0.483, 0.522]$ across thirteen levels), while T_{tri} achieves perfect discrimination in the noiseless limit and remains above 0.85 through $\sigma = 0.10$. A non-relational variance baseline is uninformative throughout. The separation on the blind class is not a matter of degree but of kind: it reflects a structural blindness of T_{pair}, which carries only the dyadic constraints and never forms the pairs that would reveal the swap (Remark 7.2), a blindness no amount of data removes from T_{pair} itself.

σ	tri / visible	tri / blind	pair / visible	pair / blind	base / blind
0.00	**1.000**	**1.000**	**1.000**	0.495	0.529
0.02	0.961	0.971	0.966	0.496	0.532
0.05	0.916	0.924	0.924	0.483	0.524
0.10	0.833	0.859	0.857	0.499	0.545
0.20	0.709	0.752	0.733	0.504	0.532
0.50	0.563	0.569	0.578	0.494	0.509
1.00	0.510	0.533	0.515	0.495	0.504
2.00	0.486	0.512	0.502	0.491	0.496

Table 2: AUC for separating each misbinding class from bound scenes, by statistic and noise level; two thousand scenes per class per level, random seed 20260523. The pair/blind column is 0.50 to within sampling noise at every level, confirming Proposition 7.1.

8 Discussion

What we have shown. The discourse complex gives binding a sheaf-theoretic home in which the operative modelling choice, sum-type coboundary rather than clique expansion, is made explicit and shown to matter. Within it, a triadic constraint is provably not the consensus space of any dyadic complex *on the same carrier* (Proposition 4.2), while an auxiliary mediator can carry it, so that "non-reduction" here is the same-carrier distinction between intrinsic and mediated triadicity. In the scalar-stalk setting we proposed a combinatorial invariant of the cell incidences, binary-decomposability, and *conjectured* that it characterises when the global-section space collapses to pairwise form. We proved sufficiency for hypergraphs admitting a matroidal reduction order (Theorem 6.9) and necessity up to a single matching identity (Theorem 6.10), and located the residual gap precisely in Conjecture 6.11, whose two halves we support on several hundred computationally tested instances. The accompanying experiment shows the distinction is not merely formal: the sum-type obstruction detects binding errors that the dyadic (clique-expansion) statistic T_{pair} is structurally unable to see.

Limitations Four limits deserve emphasis. First, the experiment's sharpest prediction is a consequence of the model's construction (channels 8 through 11 were designed to be triadic-only), so it establishes that the sum-type obstruction has the discriminative power the theory ascribes to it, but not that any real perceptual system decomposes its features this way; whether cortical binding has irreducibly triadic structure is an empirical question we do not settle, and the blindness is of T_{pair} specifically, not of pairwise information in general (Remark 7.2). Second, the model is synthetic, and a benchmark on real relational data is the natural next test, the obvious route being to substitute the sum-type coboundary for the clique-expansion Laplacian of [4] within their network. Third, the combinatorial characterization is developed for scalar stalks; the vector-stalk case, where a matrix restriction can mix coordinates within a stalk, needs a block-supported analogue and is left open. Fourth, even in the scalar case the characterization is complete only modulo Conjecture 6.11: the general sufficiency (dropping acyclicity) and local reducibility (that (a) and (b) yield a matroidal witness) both reduce to one statement of bipartite-matroid theory.

Relation to prior work and novelty. The four ingredients we draw on, cellular sheaves and their Laplacians, discourse sheaves, hypergraph sheaf networks, and the Peircean reduction thesis, are each established, and our non-reduction proposition specialises Koshkin's theorem. The novelty lies in the synthesis of Peircean same-carrier non-reduction, sheaf-theoretic global coherence, and an explicit distinction between same-carrier pairwise reduction and higher-arity cellular constraint, together with a combinatorial program, largely conjectural, for characterising

when such constraints collapse to pairwise form. We are careful to separate the mathematical content from the philosophical interpretation: the discourse complex is a linear model realising higher-arity relations of the form $\ker\big(\bigoplus_{a\in c} F(a) \to F(c)\big)$, and we do not claim that all Peircean or perceptual triadic relations are linear, nor that the model is a complete theory of feature binding or of Peircean Thirdness.

9 Open problems

Problem 1 (Close the characterization). Prove Conjecture 6.11. Its two halves, local reducibility (that (a) and (b) yield a matroidal witness, which would make every reduction order matroidal and remove the matroidality hypothesis from Theorem 6.9) and the removal of the acyclicity hypothesis, are both statements about matchings in the incidence bipartite graph of $(K, T \cup \{c\})$; a Kőnig–Egerváry argument is the natural approach. Resolving them upgrades Theorems 6.9–6.10 to a clean equivalence: for generic scalar-stalk data, H^0 is pairwise-decomposable if and only if $(A, \mathcal{C})$ is binary-decomposable.

Problem 2 (Complexity of the criterion). Is binary-decomposability decidable in time polynomial in $|A| + |\mathcal{C}|$? The present verifier enumerates candidate sets T. A polynomial algorithm, via linear-programming relaxation or matroid intersection, would make the criterion practical at the scale of learned models and enable a sum-type sheaf-hypergraph layer.

Problem 3 (A relational benchmark). Replace the clique-expansion Laplacian of [4] with the sum-type coboundary in their hypergraph network and evaluate on standard node-classification benchmarks, seeking a task or error regime where the triadic inductive bias measurably helps.

Problem 4 (Vector-valued stalks). Extend the characterization to stalks $F(a)$ of dimension greater than one with matrix-valued restrictions $F_{a\lhd c}$. Here binary support must become *block*-support (support on at most two agents' blocks, not two scalar coordinates), and Lemma 6.1 must be replaced by a block/rank analogue; a block-supported matroid formulation of Definition 6.5 is the natural target.

Problem 5 (Nonlinear relational cells). The present cells are linear: a cell c constrains its correlates to the kernel of $\bigoplus_{a\in c} F(a) \to F(c)$. More generally one could attach to each cell an arbitrary relation $\mathcal{R}_c \subseteq \prod_{a\in c} F(a)$ and call an assignment coherent when $(x_a)_{a\in c} \in \mathcal{R}_c$ for every c; the discourse complex is the linear case $\mathcal{R}_c = \ker L_c$. Which of the present results, in particular the intrinsic-versus-mediated dichotomy, survive for nonlinear $\mathcal{R}_c$, and what replaces binary-decomposability there?

Problem 6 (A behavioural test). Design a psychophysical paradigm with both T_{pair}-visible and triadic-only conjunctions, for instance agent/patient/action role bindings, where correct binding requires a joint three-way commitment, and ask whether human binding errors track the triadic obstruction T_{tri} rather than the dyadic T_{pair}.

References

[1] G. Brady and T. H. Trimble. A categorical interpretation of C. S. Peirce's propositional logic Alpha. *Journal of Pure and Applied Algebra*, 124(1–3):1–33, 2000.

[2] R. W. Burch. *A Peircean Reduction Thesis: The Foundations of Topological Logic*. Texas Tech University Press, 1991.

[3] J. M. Curry. *Sheaves, Cosheaves and Applications*. PhD thesis, University of Pennsylvania, 2014.

[4] I. Duta, G. Cassarà, F. Silvestri, and P. Liò. Sheaf hypergraph networks. In *Advances in Neural Information Processing Systems 36*, 2023. arXiv:2309.17116.

[5] J. Hansen and R. Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3(4):315–358, 2019.

[6] J. Hansen and R. Ghrist. Opinion dynamics on discourse sheaves. arXiv:2005.12798, 2020.

[7] H. G. Herzberger. Peirce's remarkable theorem. In *Pragmatism and Purpose: Essays Presented to Thomas A. Goudge*, pages 41–58. University of Toronto Press, 1981.

[8] S. Koshkin. Logical reduction of relations: from relational databases to Peirce's reduction thesis. arXiv:2406.14094, 2024.

[9] J. McLaughlin, A.-V. Pietarinen, and F. Bellucci. Betagraphic: a diagrammatic syntax for Peirce's beta graphs. *Transactions of the Charles S. Peirce Society*, 2015.

[10] C. S. Peirce. *Collected Papers of Charles Sanders Peirce*, 8 vols., ed. C. Hartshorne, P. Weiss, and A. W. Burks. Harvard University Press, 1931–1958. Cited by volume and paragraph.

[11] D. D. Roberts. *The Existential Graphs of Charles S. Peirce*. Mouton, The Hague, 1973.

[12] A. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136, 1980.

[13] C. von der Malsburg. The correlation theory of brain function. Internal Report 81-2, Max-Planck-Institute for Biophysical Chemistry, 1981.