Given Distinction

Why the Space of Structural Reasoning Failure Is Finite and Classifiable

Brent Alan Seeley · WellFlow Labs

https://orcid.org/0009-0002-9902-5955 · bseeley@wellflowlabs.com

Version: October 2026 Preprint (Paper III) · Series: Foundations of the Convergent Semantic Architecture (Paper III of IV)

Version DOI: https://doi.org/10.5281/zenodo.23382752

Abstract

Papers I and II of this series supply the empirical content of the Convergent Semantic Architecture: CLCA recovers the constraint geometry's axes and operators from sixteen languages of thirteen independent lineages; UFA classifies every entry of a 101-fallacy inventory under five structural violation types. This paper argues why the violation space has this five-part form, and states the claim that the five are complete as a conjecture defended against the strongest candidates for a further condition. The argument does not depend on the specific inventory of AOML (which axes, which operators, which methods, which levels) but on the structural conditions that any distinction-making system must maintain to remain coherent. Any system capable of distinguishing one thing from another must maintain four conditions: what kind of distinction is being made (category), how distinctions compose (licensing), what validates a distinction (method), and at what level of abstraction the distinction operates (stratification). These four conditions yield four primitive violation types. The opposition topology that makes all four possible is a precondition whose misrepresentation yields the fifth. On the reading given in §3, philosophical traditions developed over 2,400 years without coordination have each identified one or more of these conditions; that none identified a condition outside them is the hypothesis a coded survey would test. The convergence is consistent with the structural argument. The argument establishes the form of the constraint geometry; CLCA and oAOML (operational AOML, engineering failure analysis) supply its content from independent substrates. Together, the four papers provide one connected account: what the constraints are (Paper I), what happens when they are violated (Paper II), why the violation space takes the form it does (this paper), and what truth consists of on each face of the geometry (Paper IV).

1. Introduction: The Gap Between Observation and Necessity

Paper II classified every entry of a comprehensive inventory of 101 distinct named reasoning errors under five structural violation types, with no sixth type required. That is strong evidence of coverage. But it leaves a question unanswered: is the result contingent or necessary? Could a sixth type be discovered tomorrow, the way a new element can be synthesized? Or are five types all there can be, for structural reasons that go deeper than any particular inventory of fallacies?

The difference matters. An empirical taxonomy, however comprehensive, is permanently vulnerable to the next counter-example. A structural derivation is not. If the five types can be derived from the conditions for coherent distinction-making itself, then the exhaustiveness is not an observation about the fallacies we happen to have cataloged; it is a consequence of what distinction requires. A sixth type would require a fifth condition for coherent distinction-making, not merely an uncategorized fallacy.

This paper develops that derivation. The argument proceeds from a single premise: given any system capable of treating one thing as different from another, certain conditions must be maintained for that capacity to remain coherent. These conditions are not features a system might choose to include; they are what coherence means for a distinction-making system. Violations of these conditions are not ordinary errors of fact; they are failures in how distinctions are drawn, composed, validated, or ordered. The five violation types documented in Paper II correspond to these conditions, one for one. The paper presents the structural derivation in §2 and the four-traditions convergence in §3, then sets out in §4 the convergent evidence from two independent empirical substrates: cross-linguistic analysis of human language (CLCA) and structural analysis of engineering catastrophes (oAOML).

The argument is independent of AOML's specific empirical content. It does not depend on whether there are six axes or more, whether the operator inventory is minimal or admits reduction, or whether the four levels exhaust the stratification. Those are empirical questions about the content of the constraint geometry, addressed by CLCA (Paper I) for the axes and operators, and by the continued development of the architecture for the methods and levels. This paper concerns the geometry itself: why it has the structure it does, and why the violation space has finitely many kinds.

2. The Argument from Distinction

The argument proceeds in three steps. First, the primitive: distinction itself, the capacity to treat A as different from not-A. Second, the necessary structure: any usable distinction presupposes a contrast (separation) and must be maintained along four conditions: category (kind), composition (licensing), validation (justification), and stratification (level). Third, the closure: the precondition and the four conditions form a constrained system whose violations fall into finitely many kinds. The five types are not a list but the set of ways this structure can fail; that the set is complete is the claim §2.4 states and §6 defends. The remainder of this section develops each step in turn.

2.1 The Premise

The argument begins with a single premise: there exists a system capable of distinction. A system capable of distinction can treat one thing as different from another. This is the minimal starting condition. Not "given language" or "given human cognition" or "given AOML." Just: given any system that can distinguish.

A sensor that cannot distinguish signal from non-signal is not a sensor. A language that cannot distinguish true from false is not encoding truth. A logic that cannot distinguish valid from invalid is not a logic. A software system that cannot distinguish one state from another is not maintaining state. The capacity for distinction is what makes something a system in the relevant sense, rather than undifferentiated noise.

2.2 The Four Conditions

A system capable of distinction must do more than register a momentary difference. To reason, compute, or manage state coherently, it must maintain that distinction across operations. To sustain a distinction without the boundary dissolving into noise, the system is forced to track four aspects of every distinction it holds. These are not design choices or optional features; they are the minimal bookkeeping required for distinction-making to persist. A distinction that cannot be maintained across category, composition, validation, and stratification is not a stable distinction but a transient signal indistinguishable from noise.

Condition 1: Category. The system must maintain what kind of distinction it is making. A factual distinction (corresponds to reality or not) is structurally different from a normative distinction (right or wrong), which is structurally different from an authenticity distinction (genuine or counterfeit), which is structurally different from a sincerity distinction (honest or performed). Without category tracking, the system cannot prevent the most

fundamental confusion: treating one kind of claim as if it validated a different kind. That someone is insincere does not establish that their claim is false. That something is natural does not establish that it is good. These are not merely common errors; they are the errors you get when a system composes distinctions without tracking what kind of distinction each one is.

Condition 2: Composition Licensing. The system must track how its distinctions compose, and which compositions are legitimate. "A preceded B" is a different composition from "A caused B." "A correlates with B" is a different composition from "A produces B." A system that composes distinctions without tracking which compositions are licensed will produce combinations that have the form of meaningful compositions but are not. Post hoc ergo propter hoc is exactly this: temporal precedence composed as if it were causation. The composition is performed; the licensing condition is not met.

Condition 3: Validation. The system must track what validates its distinctions, and whether the validation method is appropriate to the kind of distinction being made. A factual distinction validated by consensus rather than observation has the form of a validated distinction without the validation that would make it one. An authenticity claim validated by popularity rather than provenance has been checked, but checked with the wrong instrument. Without method tracking, the system produces something that looks verified but is not verified in the way the claim requires. Validation is not reducible to category or composition. Category determines what kind of distinction is being made; composition determines how distinctions combine. Validation determines whether the distinction has been established through an appropriate resolution mechanism. A system may correctly classify a distinction and correctly compose it, yet still fail by applying a validation method that does not match the distinction type.

Condition 4: Stratification. The system must track at what level of abstraction its distinctions operate, and must preserve level boundaries across inference. A distinction about this particular instance is structurally different from a distinction about the general pattern. A distinction about the empirical case is structurally different from a distinction about the formal definition. Without stratification, the system cannot distinguish between seeing one white swan and concluding all swans are white. It collapses instance into category, part into whole, observation into law. (Where such a crossing could have been licensed, by an adequate sampling warrant for instance, the companion Map types the error as an unlicensed composition located at a level boundary, a Type 2 root at a Type 4 locus; the stratification condition's own failure, a Type 4 root, is the case in which no legal assignment of levels exists at all, §5.)

The argument does not claim that exactly four conditions must exist as a matter of logical fiat. It claims that any system capable of sustained, coherent distinction-making must track four functional aspects of its distinctions: what kind they are, how they compose, what validates them, and at what level of abstraction they operate. These are not optional features. They are the minimal bookkeeping required to prevent the system's own distinctions from collapsing into noise. The class of systems the claim concerns should be stated: those that make distinctions of more than one kind, compose them, and carry them across levels of abstraction, which is where the four conditions have work to do. A binary memory element sustains one distinction across operations and satisfies the conditions vacuously, with a singleton category and a trivial level; that the framework can describe such a case shows representability, not necessity, and the necessity claim is made for the class just named. The opposition topology is the background precondition that makes tracking these four aspects possible. Anyone who wishes to argue for a smaller set must demonstrate a successful reduction: that one of these four functions can be fully subsumed under the others without loss of coherence. Anyone who wishes to argue for a larger set must identify a fifth function that is necessary for coherent distinction-making but is not already covered. Until such a demonstration or identification is provided, the derivation from distinction to five violation types stands as a defended conjecture: the four conditions are argued to be necessary; the exhaustiveness of the set is supported by

the failure of the candidates examined in §6 and by the companion evidence of §4; their independence has not yet been shown; and those remain the obligations a proof would discharge (§7).

2.3 The Precondition

Before any of the four conditions can operate, the system must be able to separate this from not-this. The opposition topology, the structure of what can be distinguished from what, is not a fifth condition coordinate with the other four. It is the precondition that makes all four possible. Category presupposes that categories can be distinguished from each other. Composition licensing presupposes that licensed and unlicensed compositions can be distinguished. Validation presupposes that valid and invalid methods can be distinguished. Stratification presupposes that levels can be distinguished from each other. All four conditions depend on the system maintaining a coherent space within which distinctions can be drawn.

This is why misrepresenting the opposition topology is uniquely destructive among the five, with the root form of Type 4 its only near relation (§5). The distinction between precondition and condition is a difference in logical order: the four conditions concern how distinguishable elements are managed within the distinction-space; the opposition topology concerns whether a distinction-space exists at all. A system that fails to maintain category, composition, validation, or stratification is making distinctions badly. A system that fails to maintain opposition topology cannot make distinctions at all in that respect. Failures of category, composition and validation, and the common, repairable form of a level-crossing, produce wrong answers within a functioning distinction-space. A misrepresentation of the topology damages the distinction-space itself, and so does a Type 4 root, in which no legal assignment of levels exists (§5). A false dilemma does not give you the wrong answer to a question; it gives you the wrong question. A straw man does not refute the wrong argument; it replaces the real argument with a fabrication, so the distinction between what was actually claimed and what is being attacked ceases to exist in the discourse.

2.4 The Derivation

Given the four conditions and the precondition, the derivation is immediate:

A failure of Condition 1 (category) produces a category violation: one kind of distinction is treated as if it were another kind. This is Type 1.

A failure of Condition 2 (composition licensing) produces an operator violation: a compositional rule is applied beyond the conditions under which it preserves coherent distinctions. This is Type 2.

A failure of Condition 3 (validation) produces a method violation: a distinction is verified using a method appropriate to a different kind of distinction. This is Type 3.

A failure of Condition 4 (stratification) produces a level violation: the distinctions a structure requires admit no legal assignment of levels, so that distinctions at different levels cannot be kept apart. This is Type 4. (An unlicensed crossing between levels that a license could repair is typed as a Type 2 root at a Type 4 locus, §5.)

A misrepresentation of the precondition (opposition topology) produces a topology violation: the space within which distinctions are made is itself mis-specified. This is Type 5.

Four conditions and one precondition; five violation types. Within the scheme the correspondence is one-to-one: each condition or precondition yields one type, and each type answers to one of them. A sixth type would require a further condition for coherent distinction-making beyond category, composition, validation, and stratification, or a further precondition beyond opposition topology; it need not require both. The burden of proof for a sixth type

is the burden of identifying that further condition, and the burden this paper's completeness claim carries is stated in §7.

2.5 Terms: Distinction, Integrity, Evaluability, Truth, Validation, Warrant

The argument uses several words that are easily run together, and the rest of the paper depends on keeping them apart.

A *distinction* is the treating of one thing as different from another: A as not-B. A claim is *structurally sound* when it satisfies the four conditions and the precondition: its kind is maintained, its compositions are licensed, its validation method fits its kind, its level is preserved, and the opposition topology it lives in is intact. Failures of structural soundness are what the five types catalog, and they are not all of a piece. A level structure no legal assignment can satisfy (a Type 4 root), or an opposition topology in which a distinction the claim needs cannot be drawn at all (the root form of Type 5: a reading whose test/production frame has merged, a measurement with no object or unit to be about), leaves nothing determinate to evaluate: such a claim has, in Paper IV's phrase, exited the domain where truth and falsity are possible. Type 5 has a second form, and it should be kept apart from the first. A false dilemma ("the die can show only 1 or 2") and a straw man are misrepresentations of the frame by claims that are themselves evaluable and false; what they destroy is not the truth-aptness of a proposition but the adequacy of the frame to the question at issue, so that the question answered is not the question asked. A failure of category, composition or method (Types 1 to 3), or a level-crossing whose license could be repaired, leaves a claim that is still truth-apt but is offered without the support or license its kind requires. *Structural coherence* is soundness taken over a whole structure of distinctions rather than one act (§6). *Evaluability* is weaker than soundness and is not to be confused with it: a claim is evaluable when it is about something and raises a determinate issue, so that it could be true or false; a claim can be about the right thing and raise the right issue while failing a condition. *Distinction integrity*, in the sense Paper IV gives the term, is the condition in which the distinctions a claim's kind requires are intact, so that a verdict can be given on it at all. Paper IV treats it as what admits a claim to evaluation, not as what makes it true, which places it with evaluability: both are conditions on a claim's being assessable, not on its being correct, and soundness asks more, that the claim's kind be maintained and its composition and support licensed as well. What makes an admitted claim true is, on that account, the realizer proper to its axis, and a structurally sound claim that fails its realizer is *false* rather than defective. (Distinction integrity is not truth: Paper IV §2.1 keeps the two apart, and the usage here follows that account.) *Validation* is the act of establishing a claim by a method; *warrant* is what a subject holds when validation appropriate to the claim's kind has succeeded. A step may be compositionally licensed subject to a condition: admitted while a warrant it depends on remains pending, as when a conclusion about a population is drawn from a sample with the question of the sample's representativeness left open and marked as open. Such a step is not a violation. It becomes one only when the pending warrant is treated as already in hand, so that the conclusion is asserted as though the condition had been met. A classification that read every conditional step as a structural failure would collapse the distinction between an incomplete argument and a fundamentally flawed one, and the five types would become nothing more than a catalog of unfinished proofs.

Two consequences follow, and both matter for reading this paper. First, an ordinary false claim is not a violation of any condition. "The door is open," asserted when the door is closed, identifies the door and retains the open/closed issue: it is evaluable, its kind is veridical, its composition and method are unexceptionable, and it is false. It has distinction integrity in Paper IV's sense and fails that axis's realizer, and nothing in its structure is at fault; the five types are not a taxonomy of falsehood. The line the paper draws is between structural failures, which the types catalog, and ordinary falsehood, which is not a structural failure at all. That line does not run between claims and arguments. An inference whose composition is unlicensed (Type 2), or a claim whose validation method is inadmissible for its kind (Type 3), is a structural failure and is cataloged; an inference that is merely unsound because a premise is false, or a belief that is merely unwarranted because a validation of the right

kind has not yet been performed, is not. Second, warranted access to a truth is a different relation from the truth itself. A stopped clock reads the right time twice a day: the claim it makes is then true, and the subject reading it has no warrant. A frozen camera feed that shows the door open when it is open has selected a method of the right type, observation, while the instance's coupling to the door has failed: the claim is true, the access is defective, and if the door then closes the claim becomes false without anyone's warrant changing. The relation that constitutes a truth and the relation through which a subject establishes or maintains access to it come apart in both directions, true but unwarranted and warranted by a method of the right kind but false. This paper concerns the first relation, the constitution of distinctions and their integrity; the second is the province of validation and warrant, which the framework addresses through its method condition and, downstream, through provenance. *Tether integrity* is the program's term for that second relation: the maintained live coupling between a validation instance and what it validates, which the stopped clock and the frozen feed show can be severed while a claim stays true, and which can equally be intact while what it delivers is false. When §7 says that a system violating the conditions produces outputs that merely resemble distinctions, it means structural failures, the failures the types catalog: not false claims, not the unsoundness of an argument with a false premise, not beliefs awaiting a validation of the right kind, each of which is a distinct thing with a distinct analysis.

This also fixes what this paper does and does not establish relative to Paper IV. This paper argues that any system that preserves distinctions under composition must satisfy the four conditions and their precondition, and that the failure space they define has five classes. It does not, by itself, say what truth consists in on any axis. Paper IV adds the premise that each axis carries its own realizer, the condition under which an admitted claim comes out true, defends the assignments, and reads the four truth traditions against them. The bridge from here to there is short and is stated at the end of §6; the defense of the axis-specific truth-conditions belongs to Paper IV and is not borrowed back here as evidence for the closure claim.

3. Independent Convergence: Every Condition Independently Discovered

The structural argument would be interesting but inconclusive if it stood alone. It gains substantial weight from an independent observation: across roughly 2,400 years of philosophy, developed without shared vocabulary or coordination, every one of the conditions has been independently identified at least once, several of them more than once. The convergence is not "one tradition per condition" (a symmetry that would itself be suspicious) but that the conditions keep being rediscovered from unrelated starting points. That no sustained tradition has identified a condition outside the set is the second half of the claim, and it is a universal negative that selected correspondences cannot establish; it is stated here as the hypothesis the coded survey would test. (The claim as made here is the author's reading of the cited passages; a coded survey of the traditions, with a stated selection and coding procedure, is future work.)

Condition	**Tradition**	**Core Concern**	**UFA Type**
Category	Category error analysis (developed in different forms by Aristotle, 1984b, and Ryle, 1949)	Certain predications are not false but nonsensical	Type 1: Axis Transfer Error
Composition	Formal logic (Aristotle's syllogistic, 1984a, through modern predicate logic)	Which derivation rules are valid	Type 2: Operator Overextension
Validation	Irrelevant appeals tradition (informal logic, argumentation theory)	Whether evidence matches the claim type	Type 3: Verification Habitat Mismatch

Stratification	Part-whole and abstraction (scholastic logic, Buddhist two-truths, Kantian distinctions, Tarski's language hierarchy)	Level-crossing without justification	Type 4: Stratification Failure (root: no legal level assignment; repairable crossings are Type 2 roots at a Type 4 locus)

Each tradition discovered its condition independently, gave it its own vocabulary, and built its own analytical apparatus around it, with no reason to suspect it was finding one face of a larger structure. Ryle (1949; category) was not working on the same problem as Aristotle's syllogistic (composition licensing), which was not working on the same problem as the informal logicians analyzing irrelevant appeals (validation), which was not working on the same problem as Nāgārjuna's analysis of two truths (*Mūlamadhyamakakārikā* 24.8–10; stratification). Stratification in particular has been found repeatedly: the Madhyamaka two truths, the scholastic part/whole analyses, Kant's distinction between phenomena and noumena (Kant, 1781/1787/1998), and Tarski's object-language/metalanguage hierarchy (Tarski, 1933/1983) are four discoveries of a single condition, made independently of one another. The two truths were also the source of the companion specification's level dimension, which was first drawn from them. That does not touch the tradition's independence from the other three; it means only that the architecture's agreement with the two truths is inherited rather than convergent, while its agreement with the scholastic, Kantian and Tarskian discoveries is not. The relation between traditions and conditions is many-to-the-set, not one-to-one: the stronger pattern, since multiple independent rediscoveries of the same condition are harder to dismiss as coincidence than a tidy one-to-one map.

This convergence is what the structural argument predicts. If the conditions are genuinely required for coherent distinction-making, then any sustained investigation of reasoning failure will eventually identify them (and will not identify anything outside them), because, if the argument is right, they are the only places where coherent distinction structurally can fail. That the conditions keep being independently rediscovered while nothing outside the set is, is strong evidence that both the argument and the traditions are tracking the same structural reality. The honest test of the claim is a candidate condition that resists the set; §6 takes up the two strongest (relevance and identity-persistence) and shows they resolve into the maintenance of an existing condition rather than a new one.

The opposition topology, the FRAME, occupies a different position from the other four conditions, but not because it has gone unstudied. It is the precondition the other four presuppose, and it has been the subject of sustained independent programs: Spencer-Brown's *Laws of Form* (1969) takes "draw a distinction" as its primitive act (nearly the premise of §2.1) and builds a formal calculus on it; Luhmann (1995) makes distinction and observation the foundation of an entire systems theory; Saussure's differential theory of meaning (1916/1959) (a term's value constituted by its opposition to others) is the same insight at the base of structural linguistics; and the dialectical tradition is a long inquiry into opposition itself. What none of these did was connect the precondition to a theory of reasoning error. Type 5 phenomena (false dilemma, straw man, false equivalence, and, as §6 argues, relevance failures) are recognized across traditions but were never gathered into a single account of how mis-specifying the distinction-frame is a structural failure mode in its own right, alongside the four. That connection is this framework's contribution: the precondition has been studied as deeply as any condition, but as the foundation of meaning, not as a source of classifiable error.

4. Form and Content: The Two Evidence Streams

The argument from distinction establishes, given the conditions of §2, the form of the constraint geometry: there must be conditions corresponding to category, composition, validation, and stratification, plus a precondition corresponding to opposition topology. It does not specify the content: which categories, which composition rules, which validation methods, which stratification levels. If the argument is correct, these conditions govern any

distinction-making system, not just human language. They must appear in engineered state-machines as rigorously as in natural language, because both are systems that maintain distinctions.

The content comes from two independent empirical evidence streams, each operating on a different substrate:

CLCA (Cross-Linguistic Compositional Analysis): Paper I reports the composition rules (eight operators with licensing conditions) on which sixteen languages from thirteen genetically independent lineages (twelve families and one isolate) converge, replicated across two language sets and two protocols; native-speaker validation of the elicited data is pending (Paper I, Appendix B). The six primary axes organize those rules, and the unseeded protocol, given no axis vocabulary, recovers five of them on its own (VER, AUTH, SINC, NORM and VERIF; only ESS has no unseeded counterpart; the G7b comparisons in the public repository). This is the linguistic instantiation of the constraint geometry. The validation methods (nine, with axis-specific habitats) and the stratification levels (four) are stated in the AOML v2.2 specification. All four dimensions began in the exploratory comparison of truth vocabularies that preceded the CLCA protocol (Paper I, §1), the levels in Sanskrit's two truths (*paramārtha-* and *saṃvṛti-satya*) and the methods in the languages' validation vocabularies (*pramāṇa*, 证/验, הוֹכָחָה). The validation habitats are keyed to the axes: what can validate a claim of a kind follows from what the kind is about (Paper IV §2.2), the unseeded run elicits evidential marking in both language sets, and proposed habitats have been tested against the corpus, with one rejected. What awaits direct elicitation is the inventory of method names and the level inventory; a run eliciting validation vocabulary and level marking would test both, and the axis-keying of the habitats is the prediction it would test.

oAOML (Operational AOML, engineering failure analysis): The companion oAOML Engineering Failure Analyses ledger documents the same violation types in purely operational, non-linguistic systems, classified against primary sources (accident boards, regulators, canonical post-mortems) at the level of the individual engineering judgment rather than the incident: an incident is a bundle of judgments, and the bundle usually mixes types. Four of its cases show the form. In the Therac-25 radiation accidents, the judgment that committed the beam mode from the represented rather than the actual machine state was an unlicensed operation inside a frame that could represent the difference (a Type 2 root at a Type 5 locus); the judgment that treated a bounded counter as a faithful flag was a Type 2 root at a Type 4 locus; the removal of the hardware interlocks was contributing and is left unclassified by the record. The Mars Climate Orbiter was lost when a thrust value in pound-force seconds was consumed by navigation software specified for newton-seconds: one kind of quantity substituted for another at an interface whose specification named the kind, the operational form of a Type 1 axis substitution. Ariane 5 Flight 501 was destroyed when a 64-bit to 16-bit conversion licensed for the Ariane 4 flight envelope was reused outside that range; the conversion raised an operand error, the exception handling shut the unit down, and its diagnostic output was read as flight data: a Type 2 root at a Type 4 locus, with a separate upstream judgment, the reuse of Ariane 4's qualification evidence without re-establishing its warrant, classified Type 3. Knight Capital lost more than $460 million in about forty-five minutes when an incomplete deployment left one server running a repurposed flag that activated defective legacy code: the deployment judgment treated the rebind as uniform across the fleet when it was not, a Type 2 overextension of the rebind operation. None of these systems uses language. The engineering findings are the investigators'; the classification of the judgments, read at the level of the judgment rather than the incident, is the ledger's.

The four conditions manifest differently across substrates but are structurally identical. In human language, an "axis" is a kind of truth claim: veridical, normative, authenticity. In an engineering system, under the operational reading the companion ledger uses, an "axis" is a kind of state variable that the system's own specification distinguishes: voltage level versus data signal, one kind of quantity versus another. These are declared categories of the system's state, playing in the projection the role the truth-kinds play in language; they are not new kinds of truth, and the axes of the companion specification classify kinds of claim, not the dimensions a subject matter can

vary along. That reading is what licenses treating a unit substitution as the operational form of a Type 1 violation. The structural role is the same: the system must track what kind of distinction it is making, and violations of that tracking produce the same structural failure regardless of substrate. Similarly, "composition licensing" in language governs which operators may combine with which axes; in engineering, it governs which state transitions are permitted from which states. "Validation" in language concerns which evidence types match which claim types; in engineering, it concerns which verification procedures match which system properties. "Stratification" in language concerns levels of abstraction; in engineering, it concerns data type boundaries and representation levels. The content differs because the substrates differ. The form is the same because both are distinction-making systems. The engineering diagnoses were made by the investigators on their own terms, before any AOML framework existed; the ledger's classification reads those diagnoses in AOML's vocabulary and records where the reading changed.

The convergence across substrates is significant. CLCA reports the geometry in human language. oAOML reads it in engineered systems. Neither substrate's facts borrow from the other's, although both are read in the same vocabulary, the ledger by the method the companion Map applies to fallacies. On the paper's argument, the geometry reappears because both substrates are distinction-making systems, and any distinction-making system must instantiate it to remain coherent.

The structural argument for the form, the two empirical streams filling in its content, and the convergence of §3 are together the warrant for the five types. It does not rest on any single leg:

The structural argument alone would be a philosophical claim without empirical grounding. CLCA alone would be a linguistic finding that might not generalize. oAOML alone would be a set of engineering case studies. The four-traditions convergence alone would be a historical observation. Together, they provide four reasons, independent in their sources though not in the vocabulary that reads them, to think the five types are tracking something real: a priori reasoning, cross-linguistic empirical data, cross-substrate engineering data, and 2,400 years of independent philosophical investigation.

5. Type 5: The Precondition, Not the Fifth Coordinate

Type 5 (Contrast Frame Collapse) requires separate treatment because it occupies a different structural position from Types 1 through 4.

Types 1 through 3, and the common form of a level-crossing, are violations of conditions within a functioning distinction-space; the root form of Type 4 is the exception taken up below. They presuppose that the space is correctly configured and then fail to maintain one of the four conditions. Type 5 is the failure that occurs when the space itself is misconfigured. The other four conditions have something to operate on only if the opposition topology is intact; Type 5 attacks the topology itself. It does so in two forms (§2.5): in the root form a distinction the claim needs cannot be drawn, and nothing is left to evaluate; in the misrepresentation form, the false dilemma and the straw man, the frame is misdescribed by claims that remain evaluable and false, and what is lost is the frame's adequacy to the question. The consequences below hold of both, because both leave the system answering a question other than the one at issue.

This asymmetry has several consequences. First, Type 5 violations are uniquely corrosive. A Type 1 error (axis substitution) gives you the wrong kind of answer; a Type 5 error (false dilemma) gives you the wrong space of possible answers. A Type 2 error (operator overextension) produces an unlicensed composition; a Type 5 error (straw man) produces a fabricated opposition that makes the question of licensing meaningless. Types 1 through 3 damage individual moves within the game; Type 5 damages the board.

Second, Type 5 is irreducible to Type 2 (operator overextension), despite superficial similarity. Both involve structural misapplication, but they operate at different logical orders. Type 2 assumes a correctly specified domain and misapplies a rule inside it. Type 5 mis-specifies the domain. The irreducibility has a clean test: for a Type 1, 2 or 3 error, including a repairable level-crossing (a Type 2 root at a Type 4 locus), you can fix the error by supplying what is missing within the existing frame: the right axis, the right licensing, the right evidence, the right level bridge or warrant. (The companion Map's Second Edition records most level-crossing fallacies in this form, Type 2 roots at a Type 4 locus, repaired by supplying a license; it also records method and contrast roots at a level locus, #94 and #74.) The root form of Type 4, which the Second Edition reserves for cases where no legal level assignment exists at all, the Liar among them, is repaired only by revising the stratification, the level structure the claim lives in. That is also a revision of the space rather than a move within it, and it should be said plainly; what separates it from a Type 5 repair is which structure is revised. A stratification revision redraws the levels while leaving the opposition topology, what is being distinguished from what, as it was; a frame revision redraws the oppositions themselves. For a Type 5 error, supplying warrant within the frame never suffices, and redrawing the levels does not touch it. If someone presents a false dilemma, you cannot resolve it by providing better evidence for one of the two options; you must show that the binary itself is wrong, that the space has dimensions the dilemma suppressed. No Type 5 error can be resolved purely by supplying warrant within the existing frame; the frame itself must be revised. This is what makes Type 5 structurally irreducible: it is the only violation type whose correction requires redescribing the opposition topology itself, rather than correcting a move within it or redrawing the levels on which moves are made.

Third, Type 5 explains why an attack on contrast frames is the most economical move available to propaganda. The priority is logical, not temporal: propaganda need not destroy frames before it enforces false categories (Type 1), manufactures false causes (Type 2), substitutes evidence (Type 3), or collapses levels (Type 4), and historically it does all of these at once. But "you are either with us or against us" does something the other tools do not. Because the opposition topology is the precondition for all four conditions, damaging it disables the detection of the other four violation types together. That is why the frame attack recurs across otherwise different propaganda systems and why it deserves separate treatment; the claim is about what the attack accomplishes, not about the order in which techniques are deployed.

6. The Open Question

The argument from distinction establishes that conditions of these four kinds, with their precondition, must hold for any distinction-making system. The CLCA evidence shows that human languages converge on a specific inventory of categories and composition rules (six axes, eight operators), which the AOML specification completes with nine methods and four levels. Two distinct open questions remain. The first is completeness: are there additional axes, operators, methods, or levels that CLCA has not yet documented, discoverable through expansion to other semantic domains? For the methods and levels, the run described in §4 is the direct test. The second is minimality: are the documented primitives all independent, or could some be derived from others, reducing the inventory to a smaller set? Both questions are empirical rather than structural; the argument from distinction constrains the form (there must be categories, licensing, validation, stratification) but does not determine the exact number of primitives within each dimension.

A sharper challenge than completeness or minimality is whether the set of conditions itself is complete: whether there is some fifth condition the four-plus-one misses. Two candidates have the best claim. Relevance (the pertinence of a contribution to the matter at hand) is real, and an independent tradition (Grice's maxims, 1975; Sperber and Wilson's relevance theory, 1986) has studied it as a condition of coherent exchange. A red herring can divert from one veridical claim to another, same axis, same level, licensed compositions and valid methods throughout: none of the four conditions is touched, yet the move fails. Identity-persistence (re-identifying a

distinction as the same across uses) has the same shape: equivocation is a term's distinction drifting mid-argument, and drift over occurrences is not obviously category, composition, validation, or stratification.

Both resolve into the FRAME, and into the same aspect of it. The FRAME is not only the static space of distinctions but the active question within it: what is being distinguished from what, here, now. A relevance failure presents an off-question claim as occupying the position of the question at issue; an equivocation lets the position of a single distinction drift across the discourse. Both are failures to maintain the FRAME coherently over time: the temporal persistence of the precondition, which the companion Map records as temporal frame failures. The test is Type 5's correctability criterion, together with §5's distinction between revising the frame and revising the levels: a relevance failure cannot be repaired by supplying more warrant within the frame; more evidence for the irrelevant point only deepens the diversion; it is repaired only by restoring the question, what is being distinguished from what, which is a revision of the frame and not of the levels. It is therefore a Type 5 violation manifested over a discourse rather than in a single step. Relevance and identity-persistence are not a fifth condition; they are the maintenance-over-time of the precondition, and Grice's relevance is the tradition that studied that maintenance. The cooperative-pragmatic residue of Grice (quantity, manner, the cooperative principle itself) governs exchange between agents rather than distinction-making per se, and is outside this framework's scope, as syntax is.

A third candidate has a different shape: groundedness. Yablo's paradox (Yablo, 1993) exhibits an infinite sequence of sentences, each saying only that every subsequent sentence is untrue. No sentence refers to itself; each is locally well-formed; each draws its distinction within a fixed category, composes it by licensed means, validates it appropriately, and respects stratification at every step, yet no consistent assignment of truth-values to the sequence exists. The classical Liar at least wears its defect on its surface, and the traditional diagnosis blames circularity; Yablo's sequence deletes the circle and keeps the collapse (whether a hidden circularity survives is disputed: Priest, 1997; Sorensen, 1998). If this failure is a violation of none of the four conditions, then groundedness is a fifth condition, and the closure argument fails by its own criterion.

The resolution is that the conditions have been stated act-by-act but hold structure-wide. The four conditions must hold not only of each distinction act but of the total structure of distinctions: the dependency graph of distinctions, where distinction A depends on distinction B when A's coherence presupposes B's evaluation, must be well-founded. Every evaluative chain must terminate in base distinctions whose satisfaction does not depend on further evaluation. Local satisfaction of all four conditions does not entail global coherence; a system of individually legal distinctions with an ungrounded dependency structure cannot sustain determinate evaluation as a whole. Four properties should be kept apart here, because the paradoxes fail them differently: consistency (some valuation satisfies all the constraints), uniqueness (exactly one does), groundedness (the valuation is reached from base distinctions along a well-founded dependency), and termination (a procedure that computes it halts). The Liar and Yablo's sequence have no consistent valuation; the truth-teller ("this sentence is true") has two and no way to choose, a failure of uniqueness with consistency intact; a benign cyclic constraint such as $x = y$ together with $y = x$ has consistent assignments and no unique one, and is no paradox at all. The condition this paper needs is groundedness. Cyclicity alone does not entail inconsistency, and consistency alone does not supply a grounded evaluation. On this paper's proposal it is not a fifth condition but the global reading of the four: each condition quantifies over the distinction structure, not merely over its members. That the structure-wide reading is a reading of the stratification condition rather than an added requirement is asserted here and not derived; it is listed among the obligations of §7. Yablo's sequence fails the stratification condition read globally, since the sequence admits no well-founded assignment of levels, each sentence evaluating all its successors.

That the global reading is principled rather than protective is shown by a case the framework classified long before this challenge arose: circular reasoning. Begging the question is a chain of individually licensed inferential steps arranged in a dependency loop, every local move legal, the whole ungrounded, and the companion Map has

always typed it within the five (as identity offered where derivation is required, #28). What was ever defective in circular justification was not the loop's shape but its ungroundedness, and Yablo's sequence proves the point by exhibiting the same defect with the loop deleted: an infinite descent in place of a cycle. It can be read as begging the question at infinity; the analogy holds for the ungroundedness and is not offered as an equivalence between the two. The rival diagnosis, which blames self-reference or circularity, tracks a surface symptom; the conditions framework tracks the structural failure and locates the defect of the Liar, Yablo's sequence, and circular justification alike in ungroundedness, whatever type their local form receives. The major formal treatments of the paradoxes each answer a different one of the four questions just distinguished, and the convergence claimed here is this paper's reading of them, not a result any of them states. Kripke's fixed-point construction (1975) characterizes groundedness, the sentences that receive a value at the least fixed point, and leaves the Liar and the truth-teller ungrounded together; Leitgeb (2005) makes the dependency analysis exact and Yablo (1982) anticipated it. Revision theory (Gupta and Belnap, 1993) classifies sentences by their behavior under revision, so that the Liar is paradoxical, unstable in every revision sequence, and the truth-teller merely pathological, stable in each sequence but neither categorically true nor categorically false. Non-well-founded semantics (Barwise and Etchemendy, 1987) traces the paradox to an unsituated (Russellian) conception of propositions rather than to circularity as such. What they share, on this reading, is that none finds the defect in self-reference itself. The framework does not legislate against ungrounded structures. It observes, with these traditions, that groundedness is what secures evaluation from base distinctions, and that the paradoxical cases receive no value in any fixed point. The boundary should be stated: Kripke's construction also gives some ungrounded sentences determinate values, those true in every fixed point (the largest intrinsic fixed point), so ungroundedness as such is not semantic failure; what fails is the dependency structure that admits no consistent valuation at all.

Groundedness in this evaluative sense, well-foundedness of the dependency structure of distinctions, is distinct from grounding as external correspondence, addressed below. The former concerns whether the structure of evaluations terminates; the latter concerns whether certain axes require tethering to something beyond the structure. Both are enforced by the five types; neither is a sixth.

The division of labor between the argument and the empirical program is deliberate. The argument gives the structural reason for the architecture, given its conditions; the empirical program determines what it contains. The two are complementary, not competing. Future CLCA work may add axes, refine operators, or discover new structural rules. None of these refinements would affect the structural argument, which concerns the form of the geometry rather than its specific content. The five violation types are derived from the form, and so long as the conditions stand, they stand regardless of how the content evolves.

A clarification on the relationship between structural integrity and grounding. The five types concern structural soundness, not correspondence. On the dispositional axis the realizer is internal (the fit, below), and the types catch failures in how such claims are composed and validated, R1's blocked operation among them. On the veridical axis soundness neither secures truth nor excludes it: a claim can be sound and false, or unsound and true (the stopped clock of §2.5), and the types do not catch falsehood. What they catch is the structural part of a grounding failure. A method violation (Type 3) on a veridical claim offers support that cannot connect the claim to what it is about; a Type 4 root, where the level structure connecting a representation to what it represents admits no legal assignment, leaves nothing determinate to connect; a frame collapse (Type 5) in which the distinction between simulated and real breaks corrupts the distinction grounding presupposes. Whether a sound, admitted claim in fact corresponds is the realizer's question, and the maintained coupling between a validation instance and what it validates is tether integrity (§2.5); neither is a sixth type, and neither is settled by the five. Runtime implementation of correspondence tracking (provenance chains, evidential audit trails) operationalizes that coupling, alongside what the five types structurally enforce.

The relationship between structural integrity and truth is not uniform across all axes. Truth-conditions on every axis share one schema with two steps: the distinctions a claim requires must be intact, which admits the claim to evaluation, and an axis-specific realizer then determines its verdict. What differs between axes is the realizer, and in particular whether it is a relation to something outside the claim-making system. On the sincerity axis the realizer is internal: it is the fit between what is internally held and what is outwardly performed, a relation the agent maintains, which is R1's structural source, because an external cause can install either state but not, directly, the fit between them. On the essence axis, at the formal level, the realizer is likewise internal: a theorem or a definitional truth holds by following under the system's own constitutive rules, with nothing external to correspond to. On the veridical and verification axes the realizer is external, observable reality in the one case and an appropriate verification process actually carried out in the other; the five types enforce this through method constraints (Type 3) and level constraints (Type 4) on those axes. Authenticity is mixed, by aspect. What makes an object authentic is a historical relation, its actual origin, and the distinction between genuine origin and counterfeit appearance must be drawable for the question to arise at all; a claim *that* the object is authentic is true by correspondence with that provenance, which is the axis's realizer; and the evidence that warrants the attribution is a third thing again, which can move without the origin changing (newly found documents establish an authorship that was always the case, and lost documents do not unmake it). (Verisimilitude is a meta-axis modulation of the veridical axis rather than an axis of its own; its realizer is the correspondence of a representation to what it represents.) The two-ingredient decomposition is offered here as a structural reading of the axis classification, not as a position on the philosophy of truth. The axes discussed here differ in which ingredients their truth-conditions require (the normative axis, whose realizer Paper IV leaves dependent on its pole, is set aside here); whether one then identifies truth with correspondence, coherence, constitutive rules, or some combination is a separate question this paper does not engage. Systematic treatment is developed in Paper IV (Truth on an Axis). That is the bridge to Paper IV, and it is all this paper supplies: the argument here concerns the conditions and the failure space; the premise that each axis carries its own realizer, and the defense of the assignments just sketched, are Paper IV's, and nothing in Paper IV's reading of the truth traditions is used here as evidence that the failure space is complete.

7. Conclusion

The argument has three parts. A primitive: distinction. A necessary structure: a fixed set of conditions required to maintain it. A closure claim: those conditions admit a finite set of failure modes, and the five types are argued to exhaust it, a conjecture defended in §6 and stated with its remaining obligations below. The claim is not a taxonomy but a closure property: given distinction, both coherence and its failure are structurally constrained. Any system that appears to make distinctions while violating these conditions is producing outputs that merely resemble licensed distinctions (in the sense fixed in §2.5: a structural failure, not a false or unwarranted claim).

The following paragraph restates the core of the argument. Any system capable of treating one thing as different from another must maintain four conditions (category, composition licensing, validation, stratification) plus the opposition topology that makes all four possible. Violations of these conditions are not ordinary errors of fact; they are failures in how distinctions are drawn, composed, validated, or ordered. The five types documented in Paper II correspond to these conditions, one for one. On the reading of §3, philosophical traditions developed over 2,400 years have each identified one or more of the four conditions; that none identified a condition outside them is a hypothesis the coded survey of §3 would test. Two empirical evidence streams (CLCA in human language, oAOML in engineered systems) fill in the specific content of the geometry from independent substrates. The convergence across a priori reasoning, empirical data, engineering analysis, and philosophical history provides four reasons to think the five types are tracking something real.

The argument does not close the question of what specific constraints the geometry contains; that is the ongoing work of the empirical program. It gives the structural reason the violation space has the form it does: the conditions for coherent distinction are four plus one, and each way of failing them is a type. That the five exhaust the space is stated as a conjecture and defended: the four conditions are argued necessary for any system that sustains distinctions; the two strongest candidates for a further condition, relevance and identity-persistence, resolve into the precondition maintained over time, and the third, groundedness, into the conditions read structure-wide (§6); and the companion evidence of §4 has so far found nothing outside the set. The obligations a proof would discharge are named: a necessity argument for each condition within the stated class of systems, independence witnesses (one condition failing while the others hold, under fixed definitions), an exhaustiveness argument that does not rest on the absence of a proposed sixth, and a derivation of the structure-wide reading of §6 from the stratification condition, in place of the identification asserted there.

The conjecture can fail, and it is worth saying how, because a classification of this kind is always open to the charge that its categories are drawn so widely that any failure can be fitted to one of them. Three findings would count against it. The first would be a structural failure, one that is not simply a false claim or an argument with a false premise (§2.5), that none of the five repairs fixes: not supplying a license, not changing the method, not restoring the level, not revising the stratification, not correcting the category, not redrawing the opposition topology. The second would be a Type 2 assignment whose missing license cannot be named before the case is classified, only after it. Composition licensing is not defined by the failures it is used to explain; the licenses are fixed in advance by the companion specification's legality matrix and structural rules, drawn largely from cross-linguistic evidence gathered without reference to any fallacy (the temporal substitution rule, R6, also draws on the fallacy tradition, and some licenses, such as the inductive license, are stated in the companion Map rather than the matrix). An assignment that could be justified only by reasoning backward from the failure would show the category doing no work. The third would be systematic disagreement between independent coders assigning types to the same arguments, which is what the inter-rater study reported in Paper II is designed to detect; its first phase, with language-model raters, found moderate agreement (Fleiss's kappa of 0.51), and the phase with human raters is to come. Of the three, the most demanding test is the one closest to practice: whether independent readers, typing the same argument without seeing each other's work, arrive at the same assignment. Where the typing is done by someone with a stake in the result, the classification can pass every check above while the fitting happens in the reading. Blind assignment and rules fixed before the cases are examined are the safeguard, and the classification is only as strong as the tests it has passed under those conditions.

Paper I recovered the routines that unrelated languages share and the limits on what they will say. Paper II cataloged the illegal operations. This paper argues why there are five kinds of illegal operation and states, as a defended conjecture, that there are no more. Paper IV argues for an account of what truth is on each axis when the conditions are satisfied, completing the theoretical arc. Together, these four papers constitute the theoretical foundation of the Convergent Semantic Architecture: what the constraints are, what happens when they are violated, why the violation space takes the form it does, and what truth consists of on each face of the geometry. The implementation of these constraints in artificial and engineered systems, including runtime enforcement architectures and protection during system self-modification, builds on this foundation as a subsequent body of work.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the author used Claude (Anthropic) in order to assist with drafting, revising, and checking the text. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

Aristotle. (1984a). *Prior Analytics* (A. J. Jenkinson, Trans.). In J. Barnes (Ed.), *The Complete Works of Aristotle* (Vol. 1). Princeton University Press.

Aristotle. (1984b). *On Sophistical Refutations* (W. A. Pickard-Cambridge, Trans.). In J. Barnes (Ed.), *The Complete Works of Aristotle* (Vol. 1, pp. 278–314). Princeton University Press.

Barwise, J., & Etchemendy, J. (1987). *The Liar: An Essay on Truth and Circularity.* Oxford University Press.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics, Vol. 3: Speech Acts* (pp. 41–58). Academic Press.

Gupta, A., & Belnap, N. (1993). *The Revision Theory of Truth.* MIT Press.

Kant, I. (1998). *Critique of Pure Reason* (P. Guyer & A. W. Wood, Eds. & Trans.). Cambridge University Press. (Original work published 1781/1787)

Kripke, S. (1975). Outline of a theory of truth. *Journal of Philosophy*, 72(19), 690–716.

Leitgeb, H. (2005). What truth depends on. *Journal of Philosophical Logic*, 34(2), 155–192.

Luhmann, N. (1995). *Social Systems* (J. Bednarz, Jr., with D. Baecker, Trans.). Stanford University Press. (Original work published 1984)

Nāgārjuna. (c. 150 CE). *Mūlamadhyamakakārikā.*

Priest, G. (1997). Yablo's paradox. *Analysis*, 57(4), 236–242.

Ryle, G. (1949). *The Concept of Mind.* Hutchinson.

Saussure, F. de. (1959). *Course in General Linguistics* (C. Bally & A. Sechehaye, Eds., with A. Riedlinger; W. Baskin, Trans.). Philosophical Library. (Original work published 1916)

Seeley, B. A. (2026). AOML v2.2: Semantic Constraint Architecture. Companion document to Paper II. Revision 2. https://doi.org/10.5281/zenodo.23243616

Seeley, B. A. (2026). Cross-Linguistic Compositional Analysis (CLCA): A Reproducible Protocol for Recovering Shared Conceptual Structure from Cross-Linguistic Evidence. Paper I, Foundations of the Convergent Semantic Architecture. Revision 2. https://doi.org/10.5281/zenodo.23242522

Seeley, B. A. (2026). oAOML: Engineering Failure Analyses. Canonical Incidents Mapped to the Five Structural Violation Types. Companion document to Papers II and III. https://doi.org/10.5281/zenodo.21960566

Seeley, B. A. (2026). The Universal Fallacy Architecture: Structural Error Prediction from AOML Constraint Geometry. Paper II, Foundations of the Convergent Semantic Architecture. Revision 2. https://doi.org/10.5281/zenodo.23249789

Seeley, B. A. (2026). The Universal Fallacy Map: Classical and Contemporary Fallacies Mapped to Five Structural Error Types (Second Edition). Companion document to Paper II. https://doi.org/10.5281/zenodo.21880635

Seeley, B. A. (forthcoming). Truth on an Axis: How the Philosophy of Truth Maps to the Semantic Constraint Geometry. Paper IV, Foundations of the Convergent Semantic Architecture.

Sorensen, R. (1998). Yablo's paradox and kindred infinite liars. *Mind*, 107(425), 137–155.

Spencer-Brown, G. (1969). *Laws of Form.* Allen & Unwin.

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition.* Blackwell.

Tarski, A. (1983). The concept of truth in formalized languages. In J. Corcoran (Ed.), *Logic, Semantics, Metamathematics* (J. H. Woodger, Trans.; 2nd ed., pp. 152–278). Hackett. (Original work published 1933)

Yablo, S. (1982). Grounding, dependence, and paradox. *Journal of Philosophical Logic*, 11(1), 117–137.

Yablo, S. (1993). Paradox without self-reference. *Analysis*, 53(4), 251–252.