

STATUS

Status: unratified, unmerged RC4 candidate. Within this candidate, Formal Specification 1.6.3 is the governing formal source; neither document is a ratified release. Formal guarantees, conditional propositions, substantive model-class constraints, constructive witnesses, empirical hypotheses, identification programs, constitutive explications, explication-adequacy questions, and contextual literature diagnoses are kept distinct. Actual contribution and local evaluative contribution are evaluated under the same fixed, prospectively selected, class-indexed constitutively adequate scheme. Unrestricted cross-scheme convergence is not assumed to follow from adequacy alone; a narrower convergence result is guaranteed only under a same-resolution, sufficiency-preserving bridge.

1. Claim discipline

A formal guarantee follows from definitions and declared bridge conditions. A substantive constraint restricts admissible model classes but is not derivable from weaker definitions. A constructive result establishes a nonempty model or countermodel. An empirical hypothesis concerns real systems. An identification program specifies how latent alternatives might be distinguished. A contextual literature diagnosis is a scoped, source-supported claim about a research field; it is neither a formal guarantee nor an empirical result about cognition. Executable coverage means only that a finite encoding computes the registered fixture. A constitutive explication states the proposed identity or necessary role-structure of a kind. A contested case can pressure that explication's adequacy, but it is not automatically an observation-level falsifier when its classification is itself disputed. Evaluation proceeds across four noninterchangeable levels: comparative explication adequacy, formal countermodel structure, operational identification, and calibrated measurement. Research feasibility concerns the feasibility and fidelity of inquiry across those levels; it is not a fifth claim status. Sustained failure on independently specified contrast or consequence tests may pressure a constitutive explication before a complete superior theory has been supplied, although a disputed lexical verdict is not automatically decisive.

2. Core assumptions

ID	Assumption
A1	Typed representations, decision models, judgment events, frames, subjects, objects, targets, routes, and profile schemes are distinguishable.
A2	Assumption: every moral token contains at least one represented moral subject and object; the subject may be explicit or implicit in surface content. This is not recorded as an established empirical finding.
A3	Assumption: an appraisal of an event represented as having no agent behind it can be normative but is not moral. This is answerable to agentless and low-recognition adverse cases rather than treated as ordinary-agency entailment.
A4	Actual contribution is evaluated under a constitutively adequate realization-profile scheme fixed prospectively for a declared comparison class; identification models remain epistemic, and no scheme or class may be tailored after a target-token verdict.
A5	Methodological admissibility and constitutive adequacy are distinct; invalidity and unresolvedness withhold Boolean application.
A6	Control-profile indicators, aggregation, intervention family, operating range, and comparison class are fixed before the target output; no indicator is defined by that same output or behavior.
A7	Only independently admitted challenges enter the skeptical defeat graph.
A8	Alignment modules beta, gamma, tau, and shared-question bridges are canonical dependencies for conflict and polyphony.

ID	Assumption
A9	Strict bending is operation-specific: the tested edge preserves the exact core relationship and changes a downstream locus. An operation that changes a core construal or builder is core-changing, not strict bending on that edge. Builder realization, input, and modulation are sufficient but nonexhaustive ways of changing the core, and one support may bear a separately evidenced strict-bender edge.
A10	CrossSchemeContributionInvariant is a substantive model-class constraint, not a theorem of constitutive adequacy alone.
A11	For every empirical hypothesis, the antecedent construct is identified independently of the outcome used to test its predicted consequence.

Retired claim ID A12 (strong semantic CET). Version 2.5 removed the universal strong-semantic-CET premise from NDT's assumptions and dependencies. The argument that every genuine semantic success condition is minimally evaluative remains an acknowledged external argument only. It neither supplies the general evaluative prerequisite nor establishes encoded-evaluativity existence or non-negligibility.

3. Constitutive explications

ID	Constitutive explication
CE1 Builder role structure	Conditional on NDT's explanandum and on preserving the inherent-mattering/instrumental, passive-regard/situated-involvement, and abstract-co-presence/subject-relative-position distinctions, some functional equivalent of one inherent-valence profile, perceived agency, and Relational Positioning is required. Positive and negative inherent value are independently represented components; positive-only, negative-only, mixed, and neither remain distinguishable, and no gate or aggregation may replace the pair with a net, maximum, sum, mean, or other scalar. Perceived agency requires a represented response affordance within a prospectively fixed modal-temporal horizon. Relational Positioning is valence-neutral and non-scalar. Its governed routes are identity-mediated positioning, intimacy, and direct encounter. Identity-mediated forms are shared identity, contrastive identity, and complementary-role identity. Mere category membership, role possession, or a downstream duty is insufficient. Functional access is separate. Standing-mediated is retired; an impersonal directed assignment qualifies only through an operative complementary-role bridge. The deletion argument establishes conditional nonredundancy within NDT's present decomposition, not the universal superiority of its labels or implementation. Indicators, branch criteria, population distribution, and empirical prevalence remain open.

The full model-relative moral-typing biconditional combines CE1's builder-role explication with A2 and A3, exact subject/object bindings, same-chain contribution, bounded common integration, and the ordinary output conditions. It does not convert A2 or A3 into established empirical findings.

4. Formal guarantees

ID	Guarantee
G1 Moral subset	Within an evaluable model instance, MoralJudgment implies Outputs; outside that domain the Boolean predicate is withheld rather than set to false.
G2 Subject requirement	MoralJudgment implies the existence of an explicit or implicit MoralSubjectOf witness.
G3 Agentless exclusion	Outputs and Appraise do not imply MoralJudgment when the person forming the appraisal represents no moral subject behind the appraised event.

ID	Guarantee
G4 No leakage	Frame-mismatched, inactive, access-blocked, positioning-absent, off-chain-positioning, split, profile-silent, merely reachable, coincidentally covarying, or object-detached relationship routes cannot moralize a token. The exact relationship, subject binding, and object binding must converge at an independently typed relation-bearing evaluation input on the same producing chain; meeting only at a generic guide, candidate, or output is insufficient.
G5 Provenance closure	A derived evaluative unit cannot enter a route without its construction ancestry or explicit parent import.
G6 Anchored survival	Independently live routes do not form a live chain without valid parent-child anchoring.
G7 Skeptical challenge	An admitted undefeated or undecided undercutter blocks the targeted unit under grounded semantics.
G8 Alignment discipline	Plural outputs are not polyphonic without the bridge appropriate to their target and appraisal types; unmapped or comparison-unready pairs remain unaligned.
G9 Conditional contribution convergence	If two constitutively adequate schemes share a declared resolution and are connected over a nonempty registered intervention basis by a profile isomorphism preserving intervention effects, minimal sufficiency, and the defined token-profile difference-maker verdict, they are contribution-equivalent over K.
G10 Qualifier nonmoralization	A relation-qualifier token alone implies neither BuilderQualifiedMRCondition, MoralSubjectGate, nor MoralJudgment. Its core-changing and strict-bender edges must be typed separately.
G11 Two-tier local contribution	EvalContributes, RelationshipProfileContributes_s, SubjectContributes_s, and ObjectContributes_s require route/process-specific functional access, reachability, and tier-one local actual contribution under the fixed constitutive scheme. RelationshipProfileContributes_s additionally requires, at tier two, either a controlled profile effect on the same chain or a profile-level NESS witness for the exact bound builder profile. Profile NESS can preserve redundant profile support but cannot rescue a false tier-one local difference-maker. SubjectContributes_s and ObjectContributes_s require the exact typed binding node on that chain.
G12 Evaluable-domain guard	Outputs and every moral-type predicate require a unique judging-system binding and exactly one prospectively fixed constitutively adequate scheme for the token's declared comparison class. One explicit $K = \text{ComparisonClassOf}(\kappa, c)$ is threaded through the valence, core, relationship-strength, builder-qualified, moral-relationship, and relationship-contribution conditions; criteria or configurations registered under different classes cannot be composed. Invalidity or unresolved scheme competition withholds Boolean application.
G13 Operative evaluator source	Every admitted output has at least one operative evaluator-source node on its actual contributing chain. This bounded internal lemma is not an externally exhaustive claim and does not entail strong semantic CET or universal encoded evaluativity.
G14 Determinate violation	Violates is true only when satisfaction status is determinate and false. Missing or unresolved satisfaction does not count as violation.
G15 Praise and credit	Praise is favorable appraisal that attributes responsibility or credit to its represented subject. Admiration and generic favorable appraisal do not satisfy Praise merely by sharing positive polarity.
G16 Override levels	LicensedOverride is a structural license. OperativeOverrideOn additionally requires the license's ground and warrant to occur on an actual contributing route to the token-relative evaluation; a structural license need not be operative in a given token.
G17 Anchor identity	Numerical judgment-token identity is fixed by the judgment-event anchor. Interpretation, derivation type, realization profile, and actual source are independently represented token features, not additional identity conditions.

ID	Guarantee
G18 Prospect-order boundary	The represented preorder over uncertain prospects agrees with the outcome order on degenerate prospects. NDT imposes no universal expected-utility functional or complete risk theory for nondegenerate prospects.
G19 Output-language boundary	The current formal output is OughtChoose rather than an ordinary-language all-things-considered ought. Singleton and all-tied action sets need not yield OughtChoose, and Permission is not a primitive in the current six-output grammar.
G20 Valence-profile preservation	InherentValenceProfile retains independent positive and negative components through every builder gate and relationship contribution test. The four profile states remain distinguishable after ValenceGate; only the two branch-level Boolean results may be disjoined at that decision boundary.
G21 Positioning taxonomy	RelationalPositioning is realized exactly by an identity-mediated, intimacy, or direct-encounter route in Specification 1.6.3. Identity-mediated positioning is realized exactly by a shared, contrastive, or complementary-role identity bridge. Access, distance, valence, generic association, and downstream duties do not supply those routes.
G22 Strict-bender boundary	StrictBender implies CoreProfilePreservedOn and a downstream bender locus for the same tested operation. It therefore excludes a core change on that operation, but it does not exclude the unindexed CoreChangingFactor predicate: the same support may have a distinct core-changing operation. Builder realization, input, or modulation implies CoreChangingFactor but does not exhaust it.
G23 Independent edge credit	ValenceEdge, RelationalPositioningEdge, StrictBenderEdge, and AttractorEdge are nonentailing functional credits. One representation may bear several only when each edge independently passes its binding, intervention, and contribution test.

5. Constructive results

ID	Constructive result
C1 Content-provenance nonidentity	Same prescription; care route moral, policy route nonmoral.
C2 Crossed subject gate	Human and nonhuman referents receive the same verdict under the same attributed-recognition profile.
C3 Positive third-personal judgment	A competent represented stranger-child relation yields moral wrongness and blame under the stated conditions.
C4 Agent-implicit appraisal	An institution or responsible party can be implicit in surface content but present in the actual derivation.
C5 Representationally agentless event	“The earthquake was terrible” remains nonmoral when the person forming the judgment represents no person, collective, institution, or other agent as subject of the event.
C6 Independent builder gates	Admissible finite assignments separately witness ValenceGate, the agency floor, and RelationalPositioningGate. Valence plus agency can pass while positioning fails; positioning can pass while the valence profile is neither; all can pass together. Functional access remains separately variable.
C7 Redundant support	Moral and nonmoral routes may both actually contribute; inactive or profile-silent moral backup does not moralize.
C8 Violation-sensitive tragedy	Formal: a lower-priority violation can be guided over an outcome-better higher-priority violation. Plain language: when every option breaks a rule, the model can steer toward breaking the less important one even when breaking the more important one would produce a better outcome.

ID	Constructive result
C9 Indexed polyphony	Distinct frame events with surviving comparison bridges and conflict coexist without local inconsistency.
C10 Unaligned diagnostic subtypes	Plurality may lack an action bridge or may share an execution while lacking appraisal comparison-readiness.
C11 Moral self-relation	Role-distinct self subject and self object bindings can support a first-personal moral token.
C12 Relation-qualifier bending	An ownership qualifier can alter priority or override without directly conferring moral type.
C13 Invalid model boundary	A post hoc or non-preserving profile scheme invalidates the application rather than producing Outputs=false.
C14 Empirical unresolvedness	Divergent live candidate schemes can withhold classification without collapsing into one negative verdict.
C15 Positioning deletion and route substitution	With the valence profile, perceived agency, access, subjecthood, object binding, and output held fixed, removing a sole qualifying positioning witness, or all witnesses when several exist, makes RelationalPositioningGate and moral type fail. Deleting one witness need not do so when another qualifying route remains.
C16 Valence quadrant separation	Positive-only, negative-only, mixed, and neither profiles are pairwise distinguishable before and after the Boolean ValenceGate decision. Positive-only reproduces the 1.5.3 boundary; a maximum, net, sum, mean, or one-scalar replacement is rejected.
C17 Identity-form separation	Shared, contrastive, and complementary-role identity-mediated positions can be independently realized. An operative impersonal directed assignment may satisfy the complementary-role form; a role label or downstream duty without the bridge cannot.
C18 Strict-bender lesion	Holding the exact core profile and its actual relationship contribution fixed across the registered source and counterpart states, an urgency or qualifier edge can change a downstream requirement or priority. If the same tested operation changes a builder or core construal, the StrictBender verdict fails while a core-changing type may apply.
C19 Orthogonal multi-edge representation	One representation can carry distinct positioning and strict-bender edges, or valence and attractor edges, without collapsing their tests; lesioning one credited edge need not delete another.
C20 Crossed comparison classes	A valence configuration registered under one comparison class cannot be combined with a relationship-strength configuration registered under another to satisfy the moral-relationship condition for one judgment token.

6. Empirical hypotheses, identification programs, and contextual diagnoses

ID	Claim or program
H1 Semantic-control association	Association claim only: within an independently fixed operating regime, semantic performance is associated with a pre-specified control-sensitive order beyond an evaluatively empty baseline. H1 does not by itself establish encodedness, action guidance, pervasiveness, adaptation, or a particular evolutionary or learning history.
H1b Practical control bridge	Open empirical generalization: when an independently specified representation is correct or relevant to an object, functionally active, and independently shown to be consumed by standing control, there exists a realized stake map to an outcome-evaluative structure. Control recruitment alone entails neither encodedness, stable order preservation, nor action guidance. The antecedent indicators, intervention, operating range, and outcome measures must be specified separately.

ID	Claim or program
H2a Relationship-model identification	Given independently validated indicators and fixed parameters, measured positive and negative inherent-value components and perceived agency, together with independently identified non-scalar Relational Positioning, classify which real episodes instantiate NDT's three-builder relationship core. Independently measured functional access and profile-sensitive actual contribution distinguish an unavailable or merely available relationship from one that helped form the token. The valence pair must preserve the four profile states. Identity-mediated, intimacy, and direct encounter are the governed positioning routes; shared, contrastive, and complementary-role are the governed identity forms. This is an operational application of the constitutive roles, not their source.
H2b Relationship-model consequences	Once the three builders have been identified independently, registered interventions and held-out transfer tests distinguish their work. Relational Positioning should add preregistered information beyond the inherent-valence profile plus agency, access-only, salience, content-only, and downstream-duty partitions. Removing a positioning route should selectively alter its registered transfer or process signature without necessarily eliminating positioning or moral type when another route remains. Repeated failure of a mature positioning measure to add such information pressures H2b and CE1.
H3 Represented subject and third-person gate	Given an independently specified subject-representation model, third-person episodes that instantiate NDT's subject condition track attributed integrative recognition rather than privileged entity kind. A robust episode for which validated measures provide sufficiently strong negative evidence for explicit and implicit subject representation under a declared detection model is adverse evidence against the subject condition when independent process, transfer, and contrast measures place it with paradigm moral judgments rather than tragic or axiological appraisal. The detection model states its temporal window, unit, domain, sensitivity, false-negative risk, and evidential threshold. A superior rival strengthens but is not required for the evidence to be adverse; no single lexical verdict is decisive.
H4 Provenance	Relationship-grounded and provenance-free same-content tokens exhibit discriminating intervention, transfer, or process signatures.
H5 Admission	Registered, credible, target-relevant challenges are admitted independently of output change and condition skeptical defeat.
H6 Frame variation	Holding referents fixed while changing construal changes relationship or judgment signatures.
H7 Polyphony	Increasing accessible framings increases divergent outputs when decision resources and comparison bridges are present.
H8 Collective subjecthood	Holding policy, harm, and member conduct fixed, attributed collective recognition predicts institution- versus member-directed blame.
H9 Relation qualifiers	Holding core relationship and content fixed, independently represented qualifiers selectively alter downstream action descriptions, duties, permissions, priorities, appraisals, or responsibility attributions. This is a strict-bender test. A manipulation that changes a core builder or construal tests a core-changing edge instead.
IP1 Provenance identification	Triangulate intervention, transfer, process measures, lesion/rescue, and held-out model comparison. In third-person and first-person reconstructive cases independently measure whether the focal appraisal targets conduct, character, moral worth, or judgment provenance; include first-person prospective cases that the character-appraisal interpretation does not explain.
IP2 Actual contribution	Use active-process evidence, minimal sufficient active sets, organization-preserving profile schemes, and pre-specified indicators. Before any target-token verdict, preregister the judging system, declared comparison class or task index, admissible scheme set, resolution, intervention coverage, equivalence margin, reliability threshold, and evidence needed to retire a candidate scheme. Different genuine classes may use different schemes; post-hoc token-specific class or scheme tailoring invalidates the application.

ID	Claim or program
IP3 Unresolvedness	Return EmpiricallyUnresolved when live candidate schemes remain observationally underdetermined and verdict-divergent. A program-level maturity and futility rule must state when the registered interventions and discriminators have been exhausted, when the unresolved verdict remains informative, and when persistent competition instead counts against identification viability or the larger hypothesis being tested.
IP4 Convergence testing	Test whether candidate adequate schemes are connected by a same-resolution, sufficiency-preserving bridge before asserting scheme-invariant contribution.
IP5 Bender/core-change identification	For each candidate source, hold the exact subject/object tuple, inherent-valence profile, perceived agency, and Relational Positioning fixed while testing a preregistered downstream locus. Classify a successful core-preserving downstream effect as strict bending. If a builder or core construal changes, classify the operation as core-changing. Builder realization, input, and modulation are sufficient but nonexhaustive ways of changing the core; each role requires its own evidence. Urgency, vulnerability, power, sentience, ownership, and other source labels do not determine the result.
H10 Violation-sensitive fallback	Holding outcome scores fixed, changing the represented priority and violation profile changes the selected fallback when every option violates at least one requirement.
H11 Target-sort appraisal	Holding valence and broad event content fixed, appraisal grammar varies systematically with whether the represented target is an act, agent, institution, outcome, relationship, or event.
H12 Alignment-bridge manipulation	Holding local outputs fixed, adding or removing a represented act, target, shared-question, or dimension bridge moves a pair between unaligned plurality and comparison-ready conflict.
H13 Frame individuation	Pre-specified frame-boundary indicators predict when construal, action space, object binding, or operative question changes enough to constitute a distinct frame activation.
IP6 Control-recruitment identification	Use pre-output control-state indicators, intervention, held-out transfer, and rival-model comparison so the predicted behavior does not define ControlRecruited. Independently identify consumption by standing control and policy-state sensitivity to intervention. A positive result tests H1b's antecedent but does not by itself establish encodedness, stable semantic-control order preservation, or action guidance.
IP7 Broad NDT empirical program	Coordinate separate programs for CET/control recruitment, normative detachment, builder measurement, subjecthood, challenge admission, fallback, appraisal targets, frame individuation, collective agency, relation qualifiers, and alignment/polyphony. Provenance measurement addresses one central MDT thesis, not the whole theory.
D-LIT-001 Explanandum divergence	Scoped contextual literature diagnosis, not a formal guarantee or empirical result. Across the representative programs audited in Moral Cognition Explanandum Audit 1.2, operational targets differ and no shared token-level criterion separates moral from nonmoral normativity. This supports stating that NDT begins from a classification diagnosis; it does not prove that every existing or possible theory lacks such a criterion.
H14 Weak semantic CET	Across a declared population and operating range, greater semantic differentiation predicts selective or control-sensitive organization beyond an evaluatively empty baseline. This is an association claim. Evolution or learning may motivate it but is not a governed causal, adaptation, or pervasiveness claim.
H15a Encoded-evaluativity existence	Within an independently delimited population of representation tokens, at least one token instantiates EncOutcomeEval. This existential claim does not establish prevalence or person-level pervasiveness.
H15b Encoded-evaluativity non-negligibility	Within the same representation-token population and sampling frame used for H15a, encoded evaluativity occurs at a rate governed in advance as non-negligible. NDT fixes no universal numerical threshold; sampling, measurement, uncertainty, and the non-negligibility rule require future empirical governance.

ID	Claim or program
H16 Robustness ordering	Under matched perturbation and measurement quality, relationship builders are expected to be more stable determinants of the represented relationship than downstream strict benders, and strict benders more stable than surface attractors. This ordering is a working hypothesis, not a constitutive theorem. Evidence status. Beal and Rottman (2024) provide preliminary consilience: issue-targeted candidate-bender items retained associations in models where broad value and generic attractor measures did not. The design does not identify H16's matched robustness ordering. The published labels do not establish the new core-preservation test. Specificity and semantic proximity, asymmetric and tuned item construction, backward stepwise selection without held-out validation, and unresolved causal direction keep the result exploratory for H16.
H17 Full-architecture instantiation	At least some real judgment episodes instantiate NDT's complete framed-detachment architecture: independently identified frame, evaluative structure, decision or appraisal model, candidate token, and actually contributing live chain jointly produce the output; moral episodes additionally satisfy the relationship, subject, object, and same-chain provenance conditions. Fit to isolated components does not establish this hypothesis.
H18 Provenance-based kind unity	Independently identified provenance-based token partitions exhibit greater held-out intervention, transfer, process, or explanatory stability than partitions based only on content, lexical labels, overt behavior, character appraisal, or free-standing policy source. Repeated failure to achieve the preregistered substantively meaningful advantage across feasible pre-specified tests counts against NDT's proposed provenance-based unifier; consistent rival outperformance strengthens that adverse result but is not required. Heterogeneity alone is not decisive because a functional kind may be multiply realized.
IP8 Explication adequacy	Compare NDT's proposed kind with independently specified rivals using paradigm and contrast coverage, formal and conceptual noncircularity, integration with nonmoral normativity, causal or explanatory stability, and discriminating consequences. The definitions are formally non-self-referential, and the architecture proposes a conceptually noncircular distinction; validated operationally independent identification remains open. Avoid target leakage by validating antecedent measures independently and evaluating consequences in held-out cases. Treat disputed lexical extension as evidence to analyze, not as a pre-labeled outcome.
IP9 Relational-Positioning identification	Within a declared domain, validate the three governed route families and three identity forms using exact subject/object bindings, positive and negative controls, held-out target substitution, and transfer tests fixed before the target classification. Compare positioning with inherent valence, agency, access, salience, generic association, content, and downstream-duty rivals. Identity-mediated positioning requires an independently specified operative schema and bridge contribution; category membership, role possession, or duty association is insufficient. Under an adequate scheme, absence of every governed route is positioning absence rather than distance or an unnamed rescue route. Repeated mature failure to obtain incremental information, or systematic absence of the registered signature in paradigm-patterning episodes, pressures H2b or CE1.

7. Executable coverage 2.2

Executable Partial Reference Model 2.2 is the unratified finite RC4 companion. Its registered controls must preserve the positive and negative inherent-value components through gates and contribution tests; keep positive-only, negative-only, mixed, and neither distinguishable after the ValenceGate decision; reject net, maximum, sum, mean, and other one-scalar substitutions; exercise identity-mediated, intimacy, and direct-encounter positioning plus the shared, contrastive, and complementary-role identity forms; keep access separate; bind both finite gate configurations to the token's single prospective comparison class and reject crossed-class composition; and test the strict-bender/core-changing split. Those fixtures are additional to the Gate 1 system-binding, contribution, typed integration, definedness, praise, override, uncertainty, and nontriviality controls. The program and its mutation audit establish only finite conformance to registered fixtures. They do not discover latent routes from raw data, prove the general theory, establish empirical identifiability, or prove prose-code equivalence.

8. Formal companion dependency order

- 1. Types, anchors, and construals
- 2. Semantic, outcome, control, and action evaluation
- 3. Decision models, standards, and fallback
- 4. Judgment grammar and target formation
- 5. Moral relationships and represented subjects
- 6. Actual support, provenance, and import/export
- 7. Challenge admission and defeat
- 8. Action, target, dimension, and question alignment
- 9. Formal guarantees and constructive witnesses
- 10. Executable fixtures and empirical identification protocols

9. Article and visualization crosswalk

Article section	Primary destinations
1-2	A1-A11; retired A12; CE1; G1-G3
3	H1, H1b, H14, H15a, H15b; control and detachment modules
4-5	CE1; G1-G7, G10; C1-C8, C11-C14; H2-H6, H9; IP1-IP4
6	G8; C9-C10; H6-H7
7-8	H1-H14, H15a-H15b, H16-H18; IP1-IP9; D-LIT-001; executable coverage and empirical limits
Visual start module	C1-C3; G1-G4; designated inert referent control
Worked cases	C1-C19 and finite executable boundaries
Full architecture	All assumptions, constitutive explications, guarantees, hypotheses, constraints, and crosswalks