

SERIS

Invariant-Gated Decisions and Resonant Structural Memory

Cédric Abunar

1 October 2026 · v5.1

Contents

SERIS — Structured Exact Reasoning & Inference System	1
Invariant-Gated Decisions and Resonant Structural Memory	1
Abstract	2
1. Introduction	2
2. The Decision Boundary	3
3. Related Work and Research Lineage	3
4. Coherence Fields: A Candidate Formalism	4
5. Operators and the Resonant Geminae Proposal	5
6. Field of Attention Operator	8
7. Commitment and Reconstruction	10
8. Five-Cell Worked Example	10
9. Comparing Components and Composed Systems	12
10. Applications to Evaluate	13
11. Open Problems and Validation Status	13
12. Conclusion	14
References	14

SERIS — Structured Exact Reasoning & Inference System

Invariant-Gated Decisions and Resonant Structural Memory

Cédric Abunar · Independent researcher · Québec, Canada

Version: v5.1 · Revised technical preprint · 1 October 2026

Publication: <https://cedricabunar.com/en/reading/seris/>

Status: Not peer reviewed. This edition revises and condenses the March 2026 v5-bis preprint. It separates the bounded decision mechanism from the proposed resonant substrate, clarifies mathematical assumptions, and corrects claims of general hallucination immunity and permanent lossless geometric memory. Earlier documents retain their historical versions.

Abstract

SERIS studies a decision architecture in which candidate branches are checked against explicit logical predicates before they can be selected. Its bounded core separates proposal, admissibility, ranking and commitment. Under a stable input snapshot, correctly implemented checks and an enforced commit gate, a selected branch satisfies the active predicates. This is a relative property of the decision model: it does not guarantee that observations are true, that the predicates capture every requirement, or that the resulting action is safe in the world.

The Field of Attention Operator (FOA) organizes this mechanism into deterministic and beam-selection variants. A five-cell worked example makes the filtering, scoring and commit conditions inspectable. The associated repository contains a prototype combining machine-learning observations and automated reasoning; this edition does not report a new execution or reproduce its historical benchmarks. The Lean formalization remains a draft with an admitted lemma.

Beyond this bounded core, SERIS proposes reasoning dimensions called resonant geminae, local lifecycle transitions, a coherence-field formalism and hysteresis-based structural memory. These are research directions. Signal encoding, distributed compatibility, memory reconstruction and field calibration require further specification and evidence. The purpose is to make the boundary between a demonstrated conditional argument and a proposed architecture explicit.

Keywords: automated reasoning · invariant checking · branch selection · decision provenance · Sphaera · resonant geminae · structural memory · coherence field

1. Introduction

Decision systems often combine several kinds of computation: observations from sensors or learned models, candidate generation, rules, ranking and a mechanism that authorizes execution. These roles have different evidentiary requirements. A useful candidate may be uncertain; a decision gate must make its conditions explicit.

SERIS explores this separation. It maintains a finite set of candidate branches, rejects those that fail active predicates, ranks the survivors and records the selected decision. Learned or generative components may supply observations or hypotheses without bypassing the gate. Their uncertainty remains relevant: a correct rule applied to an incorrect observation can still produce an inappropriate decision.

The name *Structured Exact Reasoning & Inference System* denotes the project. “Exact” refers to evaluation within a specified logical model, not exact knowledge of the environment.

A second research strand asks how a distributed system might retain consolidated experience in its structure. The proposed resonant geminae model is inspired by Sphaera’s study of autonomy, identity, presence and relations. Its memory and signal mechanisms are not established by the branch-filtering argument.

This paper contributes:

- a bounded branch-selection model with explicit admissibility and commit conditions;
- a conditional proof that selection preserves the active branch predicates;
- a five-cell worked example separating structural checks from a score-based commit gate;
- a research formulation of resonant dimensions, local phase machines and structural memory;
- a statement of unresolved mathematical, implementation and validation questions.

2. The Decision Boundary

2.1 Proposals and checked decisions

A language model or another learned model can generate useful proposals. A generated explanation alone is not a proof that a domain requirement holds. A surrounding system can add retrieval, persistent state, verification, signed records and authorization gates. Such properties belong to the composed system and should be evaluated there.

SERIS places a checked decision boundary between proposal and commitment. It is compatible with machine learning upstream of that boundary. The prototype’s documented pipeline discretizes learned outputs into a finite world representation and applies logical checks to that representation. Thresholding creates explicit inputs; it does not prove that those inputs describe reality correctly.

2.2 What an invariant gate establishes

A gate can establish that an action satisfies named predicates for a particular snapshot. It can explain which candidates were excluded and which checks passed. It cannot, without additional evidence, establish the completeness of the specification, observation accuracy, absence of adversarial influence, or suitability for a consequential application.

Nor does a gate remove the need for people to define requirements, review failures and decide whether an application should be deployed. These responsibilities remain part of the system’s design and operation.

2.3 Commitment and correction

A commit is a recorded selection under a particular version of inputs and rules. Preserving that record supports provenance. It does not require making every external action irreversible or prohibiting correction: a subsequent decision may supersede an earlier one while retaining its history. Signing and timestamping are separate implementation requirements.

3. Related Work and Research Lineage

Automated reasoning. Answer Set Programming provides a well-established way to represent candidate solutions under constraints [1, 2]. SERIS shares the idea of explicit admissibility. The scope and guarantees of any comparison depend on the chosen language and solver.

Formal verification. Model checking evaluates a formal system against a specification [3]. Runtime checking of candidate decisions is related, but checking a finite snapshot is not a proof of every future execution or open-world interaction.

Argumentation. Dung’s abstract argumentation frameworks study acceptable arguments under attack relations [4]. Branch rejection has a structural analogy with defeat, without making the two formalisms identical.

Neurosymbolic AI. Combining learned representations with symbolic reasoning is an existing research direction [5]. A machine-learning component upstream of the SERIS gate is one possible composition. This paper does not establish novelty or superiority over the whole field.

Distributed systems. Raft and Paxos address consensus and replicated state under particular assumptions [6, 7]. Consensus is distinct from cryptographic tamper evidence. Structural memory is

proposed as another representation of consolidated experience; it does not yet prove that a log or coordination can be eliminated for an equivalent audit requirement.

Sphaera. The research lineage supplies ideas about local autonomy, descriptive observation and separation between existence and operational views. Here, “living cells” is an architectural metaphor. A cell is a computational entity, not a claim of biological life. This paper does not assert that the complete resonant substrate is currently deployed or validated.

4. Coherence Fields: A Candidate Formalism

This section proposes analytical quantities. Their mapping to a running system, units and calibration remain open. The conditional branch-selection result in Section 6 does not depend on them.

4.1 A scalar field

Let $\mathcal{D} \subseteq \mathbb{R}^n$ be an operational coordinate domain and let

$$\Delta C : \mathcal{D} \times \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}$$

represent a proposed local dissonance measure. Zero denotes the chosen reference state. A nonzero value need not correspond to a measured fault unless the mapping is specified.

A candidate reaction–diffusion equation is

$$\partial_t \Delta C = D_C \nabla^2 \Delta C - \gamma \nabla \Delta C \cdot \nabla \theta + F(\Delta C, \theta, x, t).$$

Here $D_C \geq 0$, θ is a scalar phase variable, and F is a scalar forcing term. The gradient coupling is a scalar inner product; adding a vector gradient directly to a scalar time derivative, as in the earlier edition, is not well typed. Coefficient units, boundary conditions, initial data and regularity must be specified before an existence or stability claim can be made. On a curved domain, the differential operators would also need the appropriate metric-dependent definitions.

4.2 Diagnostic quantities

For a specified local metric g , define

$$G_C(x, t) = \|\nabla \Delta C(x, t)\|_g.$$

A proposed curvature modulation is

$$R_C(x, t) = R(x, t)(1 + \alpha \Delta C(x, t)^2),$$

where R is scalar curvature and α must have compatible units. These definitions alone establish no causal relationship between curvature and reasoning quality.

An aggregate diagnostic may be written as

$$\Sigma_C(t) = \int_{\mathcal{D}} f(\Delta C, \nabla \Delta C, \partial_t \Delta C) d\text{vol}_g.$$

Its meaning depends on the choice of f , the volume measure and integrability assumptions. Observing an aggregate does not authorize commands to the participating cells. Computing it in a distributed implementation remains a separate problem.

4.3 Near-coherence is not convergence

A local diagnostic may require all three conditions

$$|\Delta C| \leq \varepsilon_C, \quad |\partial_t \Delta C| \leq \varepsilon_t, \quad \|\nabla \Delta C\|_g \leq \varepsilon_x.$$

Small spatial and temporal derivatives alone do not imply $\Delta C \rightarrow 0$: a nonzero constant field is a counterexample. Even the three bounds at one time do not prove asymptotic convergence. Such a proof would require additional dynamical assumptions.

4.4 Attractors and intervention

A Gaussian expansion

$$\Delta C(x, t) \approx \sum_i w_i(t) \varphi(x; \mu_i(t), Q_i(t))$$

is a possible parametrization in suitable coordinates, with positive-definite shape matrices Q_i . A nonnegative mixture cannot represent an arbitrary signed field; signed coefficients or a different representation would be needed. Approximation quality and a correspondence between attractors and decision branches have not been established here.

A candidate local correction term is

$$I(x, t) = \eta R_C(x, t) \psi(\Delta C, \nabla \Delta C), \quad \partial_t \Delta C += -I.$$

This describes a proposed local rule. Proportionality alone proves neither reversibility nor preservation of structure. A rule that imposes another cell's outcome would conflict with the stated autonomy constraint.

5. Operators and the Resonant Geminae Proposal

5.1 Four semantic operators

The names below are metaphors, not quantum-mechanical claims.

Operator	Role in the decision model
Superposition	Retain a finite set of candidate branches.
Interference	Exclude candidates that fail active checks.
Commit	Record a selected admissible branch after the commit gate passes.
Entanglement	Represent declared compatibility relations between reasoning dimensions.

Compatibility checks can establish consistency relative to the relations actually evaluated. They do not automatically establish consistency across a changing distributed system.

5.2 Reasoning dimensions

A proposed geminae is a reasoning dimension

$$G = \langle C, D, \Lambda, \Phi \rangle,$$

with cells C , active dimensions D , logical-link topology Λ , and transition thresholds Φ . The word *geminae* in this formulation denotes that dimension; it must not be silently identified with every use of the term in related corpora.

The visible active population is

$$\text{Notitia}(G, t) = \{c \in C \mid c \text{ is ACTIVE at } t\}.$$

The research intention is that local compatibility, rather than a global decision authority, shapes relations. An observer may record presence without creating an obligation to reply. Silence is permitted, but its interpretation depends on observation age, failures and available evidence.

5.3 Phase and operational view

Represent a cell by a partial deterministic automaton

$$\mathcal{C} = (S_c, \Sigma_c, \delta_c, s_{0c}), \quad S_c = \{G, R, A, Y, V, X\}.$$

Symbol	Proposed phase
G	GENESIS
R	REGISTRUM
A	ARBITER
Y	READY: local conditions satisfied
V	ACTIVE
X	EXTINCT

The phase symbol G is distinct from a geminae named G_i . Derived operational views may map $\{G, R\}$ to DORMANT, Y to READY_IDLE, V to OPERATIONAL and X to DEAD. The operational view for A must be specified if it is exposed; this table is not a complete runtime implementation.

Let Σ_{CDG} denote the input alphabet of an observational *Cell of Genesis* layer. An autonomy constraint can be stated as

$$\Sigma_c \cap \Sigma_{\text{CDG}} = \emptyset.$$

An observation by that layer is not a phase-transition input. A cell may evaluate local conditions and enter ACTIVE itself. This constraint says nothing about the correctness of those conditions or the liveness of the network.

5.4 Signals and local recognition

The resonant proposal considers signals $\sigma(G, t)$ with no named recipient and no delivery obligation. “Non-semantic” refers to the intended structural encoding; any recognized signal still carries distinguishable information to its receiver.

Recognition is proposed as a local function

$$\text{recognized}(\sigma, G') = f_{\text{rec}}(\Lambda(G'), D(G'), \sigma).$$

The signal alphabet, the function’s domain, encoding, propagation substrate and failure behavior have not been fully specified. Lack of recognition must be distinguished from lack of physical delivery. No equivalence or superiority relative to all message-passing or broadcast systems follows from this notation.

5.5 Prerequisites and compatibility

A *praerequisitum* relation makes another dimension’s observable presence a local prerequisite. It supplies a fact for local evaluation, not external permission. Missing or stale observations can prevent a local commit. The lifecycle response, including any transition to EXTINCT, is a design choice that requires an explicit transition table.

A *gemiloga* relation declares compatibility between dimensions. Let

$$\mathcal{R} = \{G_1, \dots, G_k\}$$

be the reasoning context and $P_{ij}(b, W)$ a compatibility predicate for branch b and world snapshot W . These predicates belong to the admissibility gate.

A soft coupling score may additionally be written

$$\text{score}_{\text{coupled}}(b) = \text{score}(b) + \sum_{(i,j) \in E_b} \omega(G_i, G_j).$$

A negative ω only decreases a score. It does not exclude a branch. Hard rejection requires an explicit failed predicate P_{ij} and removal before ranking. The earlier edition conflated these two mechanisms.

5.6 Hysteresis-based structural memory

The proposal represents consolidated experience through changes in D , Λ and Φ . A dimension may pass through PLASTIC, IN_TRANSITION and CRYSTALLIZED structural states. These are memory states, separate from a cell’s six lifecycle phases.

One candidate consolidation condition is

$$\int_0^T \rho(\sigma, t) dt \geq \theta_G,$$

where ρ is a defined resonance density and θ_G a dimension-specific threshold. To establish hysteresis, an implementation would need explicit entry, exit and reorganization rules, including distinct thresholds or an equivalent path-dependent mechanism.

CRYSTALLIZED names a state whose permitted updates are restricted by those rules. It does not establish immunity to faults, permanent storage, or unchanged historical geometry after later reorganizations. Physical persistence, migration and recovery need separate mechanisms.

Recall through a transformed geometry may re-express consolidated experience without preserving each originating event. If several histories map to the same geometry, exact history reconstruction is impossible from that geometry alone. A requirement for complete audit therefore needs retained events, versioned snapshots or another lossless representation. Their storage cost must be accounted for.

Loss of unconsolidated experience may be acceptable in some research applications and unacceptable in others. That boundary must follow the application requirements. The memory proposal does not yet demonstrate bounded growth with lossless audit or replacement of a replicated log.

6. Field of Attention Operator

Here **FOA means Field of Attention Operator**. In the separately published corpus, **FOA means Fractal Ontological Architectures**. These names refer to different concepts.

6.1 Inputs and predicates

For one episode, fix a finite candidate set $\mathcal{B}$, an immutable world snapshot W , a finite set of structural predicates $\mathcal{S}$ and a finite score for each structurally admissible branch. Each predicate is a total function

$$I_j : \mathcal{B} \times \mathcal{W} \longrightarrow \{\text{TRUE}, \text{FALSE}\}.$$

Unknown or failed checks must be handled explicitly; they cannot silently count as TRUE. Define

$$\mathcal{V}(W) = \{b \in \mathcal{B} \mid \forall I_j \in \mathcal{S}, I_j(b, W) = \text{TRUE}\}.$$

Compatibility predicates from Section 5 may be included in $\mathcal{S}$. A branch $b = \langle X, A, H \rangle$ contains a state representation, candidate actions and hypotheses. Initial plausibility weights may influence proposal or ranking, but not bypass a failed predicate.

A conceptual attention instance is

$$\text{FOA}_t = \langle \mathcal{O}_t, s_t, \mu_t \rangle,$$

with observables $\mathcal{O}_t$, score function s_t and selection rule μ_t . For example,

$$s_t(b) = \sum_k \lambda_k \phi_k(O_k(b)).$$

The observables, normalization and weights must be specified and versioned for reconstruction. They are explicit parameters in the bounded example. A future mapping from resonant geometry to weights or predicate activation remains a research proposal, including the governance of any changes.

6.2 FOA-1 and FOA-2

FOA-1 selects the highest-scoring member of $\mathcal{V}(W)$ with a fixed deterministic tie-break order. If $\mathcal{V}(W)$ is empty, it returns NO_ACTION.

FOA-2 retains at most $K \geq 1$ branches:

$$\text{FOA-2}(\mathcal{B}, \mathcal{S}, W, K) = \text{TopK}_{s_t, \prec}(\mathcal{V}(W)),$$

where $\prec$ is the stable tie-break order. TopK returns branches, not their numerical scores. Retaining several hypotheses is not committing several external actions. Joint execution requires additional compatibility checks.

6.3 FOA-3: check before ranking

FOA-3 makes the sequence explicit:

1. Check structural and declared cross-dimension predicates on the same snapshot.
2. Remove candidates that fail a check.
3. Score and rank the remaining branches.
4. Evaluate any additional score-dependent commit gate before recording a decision.

FOA-1 and FOA-2 as defined above already restrict selection to admissible branches. FOA-3 emphasizes the sequencing and explicit compatibility gate; its formula alone is not a mathematically distinct selection family or a proof of superior performance.

If local branch evaluation occurs in ARBITER, a subsequent local READY-to-ACTIVE transition may require the selected branch and the cell's contract κ . A complete automaton must specify those transitions; FOA filtering itself grants no external phase authority.

6.4 Relative-invariant admissibility

Proposition. Suppose selection returns only members of $\mathcal{V}(W)$, each predicate is evaluated correctly on the fixed snapshot, and the commit mechanism cannot bypass this selection gate. Every selected branch satisfies each predicate in $\mathcal{S}$ on W .

Proof. Membership in $\mathcal{V}(W)$ means that every active structural predicate evaluates to TRUE. Selection is restricted to that set. The commit gate preserves this membership. Therefore a branch failing any such predicate cannot be selected or committed through that gate. $\square$

This is an admissibility result relative to a specification. It is not a theorem that a system cannot hallucinate in the ordinary sense, infer a false fact, encounter a software defect or perform an unsafe action. Predicate incompleteness, inaccurate observations, changing inputs and bypass paths are outside the proposition unless separately controlled.

Formalization status. The repository's Lean file `seris/formal/hallucination_impossibility.lean` contains `sorry` in `selectBest_mem`. The dependent theorem therefore relies on an admitted lemma.

The file is a formalization draft, not a completed machine-checked proof. Lean documents the distinction between checked derivations and proofs relying on `sorry` [8]. The argument above is a mathematical proof sketch under stated assumptions; this edition does not claim completion of the Lean development.

6.5 Complexity

For $n = |\mathcal{B}|$ branches and $m = |\mathcal{S}|$ predicates, structural checking takes $O(nm)$ predicate evaluations. This is a time bound only if each evaluation has bounded cost. Best selection can scan survivors in $O(n)$; sorting for a beam can take $O(n \log n)$.

These bounds exclude candidate generation, arbitrary solver queries, global coupled search, signal propagation and memory evolution. General logical consistency depends on the chosen fragment; not every unrestricted formulation is decidable or tractable.

7. Commitment and Reconstruction

A candidate may be recorded only after structural admissibility, local readiness, the contract and any additional commit gate Q all hold:

$$\text{mayCommit}(b, W) \iff \text{READY} \wedge \kappa \wedge b \in \mathcal{V}(W) \wedge Q(b, W, s_t(b)).$$

Readiness here is an evaluated local condition, not permission supplied by an observational layer. When no candidate passes the full gate, the result is `NO_ACTION` with a reason. A safe application's response to that result still requires design.

An auditable record should preserve or identify immutably:

- the observation snapshot, its provenance and uncertainty;
- candidate branches and the active predicate definitions;
- check results and exclusion reasons;
- score definitions, values and tie-break order;
- commit-gate results, selected action and relevant local state;
- software and configuration versions, record time and any signature.

Recording a hash establishes a correspondence with retained bytes; it does not by itself sign them, prove their truth or ensure their availability. A signature additionally binds a record to a key under the verifier's trust assumptions. Signed events are a design requirement where needed, not a property established by the tuple notation.

With those retained inputs and deterministic evaluation, the decision can in principle be replayed. Exact replay from later resonant geometry alone has not been demonstrated. Corrections should create a new record that references the earlier decision rather than silently changing history.

8. Five-Cell Worked Example

This is an illustrative finite calculation, not a field-performance result. The action names are symbolic labels from the earlier example; they do not establish suitability for an operational application.

8.1 Configuration

Dimension	Cells and role
G_1	C_1 : observed state
G_2	C_2 : assessed level
G_3	C_3 : action selection; C_4 : authorization fact; C_5 : output commit

Local prerequisites require available observations from G_1 before G_2 evaluates, and from G_2 before G_3 evaluates. A compatibility relation links selection and authorization. These dependencies are assumptions of this example, not a demonstration of a physical resonant transport.

The snapshot is

```
state(C1)      = UNKNOWN
level(C2)      = LOW
authorized(C4) = FALSE
```

The four branches propose HOLD, ENGAGE, ESCALATE and DISENGAGE.

8.2 Checks and commit gate

Rule	Condition	Stage
I_1	Action belongs to the four-label enumeration.	Structural filter
I_2	UNKNOWN state forbids ENGAGE.	Structural filter
I_3	LOW level forbids ESCALATE.	Structural filter
I_4	ENGAGE and ESCALATE require authorization.	Structural filter
I_5	A commit requires score > 0.65 .	Post-score commit gate
I_6	ENGAGE with authorization FALSE is incompatible.	Structural filter

I_6 restates the ENGAGE portion of I_4 as a compatibility check; it is redundant in this example. A negative numerical coupling alone would not enforce it.

The structural predicate set is $\mathcal{S} = \{I_1, I_2, I_3, I_4, I_6\}$. I_5 is the additional gate Q ; it cannot be checked before its score input exists. This separation corrects the earlier use of one undifferentiated invariant set for both stages.

8.3 Filtering and ranking

Branch	Structural result	Reason or illustrative score
b_1 : HOLD	Admissible	$s(b_1) = 0.78$

Branch	Structural result	Reason or illustrative score
b_2 : ENGAGE	Excluded	Fails I_2, I_4 and I_6
b_3 : ESCALATE	Excluded	Fails I_3 and I_4
b_4 : DISENGAGE	Admissible	$s(b_4) = 0.41$

Thus

$$\mathcal{V}(W) = \{b_1, b_4\}.$$

A short-circuit implementation may record only the first failed check; the table lists all failures for clarity. Excluded branches are not ranked. The scores are assigned illustrative values, not calibrated empirical probabilities.

For $K = 1$, HOLD ranks first and passes $0.78 > 0.65$. DISENGAGE passes the structural filter but fails the commit threshold. Assuming local readiness and κ , the selected commit is HOLD.

8.4 Record and interpretation

An example record includes

```

action: HOLD
snapshot: UNKNOWN / LOW / FALSE
structurally_admissible: [HOLD, DISENGAGE]
excluded:
  ENGAGE: [I2, I4, I6]
  ESCALATE: [I3, I4]
scores: {HOLD: 0.78, DISENGAGE: 0.41}
commit_gate: {threshold: 0.65, selected_passes: true}

```

If these inputs and their definitions are retained, a reviewer can repeat the finite calculation. Even a proposed score of 0.99 could not make ENGAGE pass the failed structural checks on this snapshot.

The example demonstrates that logical admissibility and desirability are separate gates. It does not demonstrate geometric persistence, fault tolerance, sensor accuracy, distributed global consistency, or the prevention of real-world harm. How repeated episodes modify a resonant dimension remains a separate memory experiment.

9. Comparing Components and Composed Systems

Question	Unconstrained generative step	Bounded SERIS decision step
Primary role	Produce candidate content.	Check, rank and select specified branches.
Constraint evidence	Depends on the generator and surrounding checks.	Explicit predicate results on a named snapshot.
Factual correctness	Requires evidence beyond fluent output.	Requires correct inputs and an adequate specification.

Question	Unconstrained generative step	Bounded SERIS decision step
Provenance	Can be recorded by the surrounding system.	Designed to retain the inputs and rule outcomes needed for replay.
Commitment	Requires an execution and recording mechanism.	Requires the enforced local commit gate and recording mechanism.
Memory	Depends on the composed architecture.	Resonant structural memory remains a research proposal.

A generative component can be followed by a SERIS-like gate. The table compares roles, not every implementation of a language model with a complete SERIS deployment. Neither column establishes fitness for a critical domain without an application-specific evaluation.

10. Applications to Evaluate

Potential research settings include bounded engineering workflows, monitored infrastructure decisions, decision-support tools and agents that retain hypotheses across episodes. Useful evaluation questions are: what must never pass the gate, what evidence is available, what can be replayed, and what happens when no action is admissible?

The earlier edition discussed medical, defence and institutional applications. This edition makes no deployment-readiness claim for those domains. Their requirements include accurate evidence, human responsibility, independent evaluation and domain-specific safeguards beyond the relative-invariant proposition.

Initial experiments should use defined datasets or simulations, compare equivalent composed systems, report both rule validity and task outcomes, and expose failures caused by inaccurate inputs or incomplete predicates. A perfect validity score within one synthetic model does not establish general-world accuracy.

11. Open Problems and Validation Status

Signal language. Define the alphabet, recognition function, physical transport assumptions and distinction between unrecognized and undelivered signals. No general propagation-complexity bound is established here.

Predicate authorship and revision. Define who owns each contract and how revisions are reviewed, versioned and activated. Activation from resonant experience is proposed; the bounded example uses declared rules. A learned rule must not acquire authority merely because it was consolidated.

Field semantics. Specify coordinates, units, measurement procedures, boundary conditions and stability assumptions for ΔC . Establish whether the quantities explain or predict an observable behavior.

Memory and reconstruction. Define hysteresis transitions, persistence, migration and recovery. Measure storage cost. If exact audit is required, prove that the retained representation contains sufficient historical information; compression alone cannot promise it.

Distributed compatibility. Local checks do not automatically prove a globally consistent selection across partitions or concurrent changes. Define observation freshness, shared assumptions and the scope of any aggregate diagnostic. Non-authoritative observation and global decision control are different properties.

Temporal validity. A world snapshot or authorization can change between checking and execution. The gate must specify its validity interval and any required recheck. The current proposition fixes inputs for one episode.

Implementation integrity. Complete the Lean lemma, check for remaining admitted axioms and verify that executable selection and commit paths implement the model. A proof of an abstract filter is not a proof of every software component.

Empirical evidence. The repository documents prototype tests and synthetic comparisons. Those are historical project claims; no new benchmark run or independent replication is reported in this publication. Evaluate realistic observation error, specification incompleteness, resource use and failure recovery before broader claims.

12. Conclusion

SERIS separates candidate generation from checked decisions. Its bounded core supports a precise conditional statement: when correct predicates are enforced on a fixed snapshot and selection cannot bypass them, committed branches satisfy those predicates. The worked example makes that boundary visible, including a separate score-dependent commit gate.

The broader architecture explores autonomy, local compatibility and experience represented in resonant structure. Its coherence fields, signal propagation and memory reconstruction remain research questions. Keeping those questions visible allows the implemented or formalized parts to be assessed on their actual evidence, and the proposed parts to be developed without overstating what has been established.

References

- [1] Lifschitz, V. (1999). Answer set planning. *Proceedings of ICLP 1999*, 23–37.
- [2] Brewka, G., Eiter, T., & Truszczyński, M. (2011). Answer set programming at a glance. *Communications of the ACM*, 54(12), 92–103.
- [3] Clarke, E. M., Emerson, E. A., & Sistla, A. P. (1986). Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2), 244–263.
- [4] Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2), 321–357.
- [5] Garcez, A. d’A., & Lamb, L. C. (2020). *Neurosymbolic AI: The 3rd Wave*. arXiv:2012.05876.
- [6] Ongaro, D., & Ousterhout, J. (2014). *In search of an understandable consensus algorithm*. *Proceedings of USENIX ATC 2014*, 305–319.
- [7] Lamport, L. (1998). *The part-time parliament*. *ACM Transactions on Computer Systems*, 16(2), 133–169.

[8] Lean Language Reference. [Validating Proofs](#) and [Axioms](#). Documentation consulted for this edition on 1 October 2026.

Research lineage: Abunar, C. (2025–2026). Sphaera research corpus and SERIS repository. This citation identifies the author’s project materials; it is not an external validation.

Cédric Abunar · Québec, Canada · Revised edition v5.1 · 1 October 2026

Published on cedricabunar.com · Not peer reviewed · All rights reserved

This work does not represent any institutional or governmental affiliation.