Quantum Hyper-Aether: Semantic Geometry, Distributed Cognitive Virtual Machines, Semantic Entropy, Semantic Temperature, Goal-Relative Truth, Constraint Databases, and Human-Guided Semantic Verification

Hirofumi Inomata
Hyper-Aether Laboratory
inomata@acm.org

Abstract

This paper proposes an integrated semantic-native computational framework called Quantum Hyper-Aether. The framework unifies semantic-cognitive virtual machines, associative semantic distance, quantum-inspired semantic computation, semantic graph evolution, neural semantic communication, distributed cognition, semantic entropy, and semantic temperature.

This extended version introduces several additional mechanisms for making the theory more operational: semantic measurement by VM birth/death frequency, stability- and goal-based semantic potential, goal-relative truth, user-dialogue-based goal distance ranking, multi-VM consensus verification, LLM-assisted semantic verification, human-in-the-loop semantic audit, and a shared constraint database for security, legal, safety, privacy, and operational rules.

The central claim is that semantic correctness in AI-native systems should not be treated only as absolute truth. In many practical AI systems, correctness is defined relative to human goals, constraints, evidence, and operational approval. Therefore, Quantum Hyper-Aether is extended into a goal-aligned, auditable, and controllable semantic operating framework.

1 Introduction

Modern computing systems separate execution, communication, memory, reasoning, and semantic interpretation. Conventional AI systems primarily rely on static embeddings, symbolic computation, or isolated neural inference systems.

Quantum Hyper-Aether proposes a different architecture in which computation itself becomes semantic evolution. In this framework:

- virtual machines become semantic-cognitive entities,
- communication becomes semantic resonance,
- memory becomes semantic geometry,
- routing becomes semantic affinity propagation,
- intelligence emerges from distributed semantic interaction.

The previous formulation introduced semantic entropy and semantic temperature as measures of uncertainty, dispersion, fragmentation, and fluctuation intensity in semantic states. However, several challenges remain:

- semantic measurement,
- semantic verification,
- semantic coherence evaluation,
- semantic energy definition,
- temperature calibration,
- collapse instability,
- semantic turbulence control,
- stable but incorrect semantic states,
- operational security and legal constraints.

This paper extends the framework by introducing practical approximations and control mechanisms for these unresolved issues.

2 Quantum Semantic Foundations

Semantic computational states are represented in a semantic Hilbert space:

$$|\Psi\rangle \in \mathcal{H}.$$

A virtual machine exists as a semantic superposition:

$$|\Psi_{VM}\rangle = \sum_i c_i |VM_i\rangle,$$

where $|VM_i\rangle$ represents semantic VM states and c_i represents semantic amplitudes.

The entire computational system evolves as a semantic graph field. A semantic state is therefore not merely a symbolic label or embedding vector, but a dynamic state in a structured semantic space.

3 Semantic-Cognitive Virtual Machines

Each VM is modeled as an autonomous semantic-cognitive entity:

$$VM_i = (S_i, N_i, I_i, M_i, R_i, F_i),$$

where:

- S_i: semantic state,
- N_i: neural inference engine,
- I_i: semantic communication interface,

- M_i: semantic memory,
- R_i: routing intelligence,
- F_i: adaptive learning feedback system.

Unlike conventional virtual machines, these VMs operate as self-evolving semantic agents. They are not merely containers of computation but active nodes in a semantic field.

4 Associative Semantic Distance

For a semantic entity S_i, we define an associative semantic field:

$$A(S_i) = \{(k_n, w_n)\},$$

where:

- k_n are associative semantic keywords,
- $w_n \geq 0$ are associative strengths.

The weighted semantic similarity between two semantic entities is:

$$Sim_w(i,j) = \frac{\sum_{k \in A_i \cap A_j} \min(w_k^{(i)}, w_k^{(j)})}{\sum_{k \in A_i \cup A_j} \max(w_k^{(i)}, w_k^{(j)})}.$$

The corresponding weighted semantic distance is:

$$D_w(i,j) = 1 - Sim_w(i,j).$$

Quantum semantic similarity is defined as:

$$Sim_Q(i,j) = |\langle \Phi(S_i) | \Phi(S_j) \rangle|^2.$$

The quantum semantic distance is:

$$D_Q(i,j) = 1 - Sim_Q(i,j).$$

The complete semantic distance metric is:

$$D_H(i,j) = \alpha D_w(i,j) + \beta D_Q(i,j) + \gamma D_G(i,j),$$

where:

$$\alpha + \beta + \gamma = 1.$$

5 Semantic Measurement by VM Birth and Death Frequency

A major unresolved issue is semantic measurement. Instead of directly measuring abstract meaning, this paper introduces a practical temporary definition:

$$M_{sem}(t) = \frac{N_{birth}(t - \Delta t, t) + N_{death}(t - \Delta t, t)}{\Delta t}.$$

Here:

- N_{birth} is the number of VM generation events,
- N_{death} is the number of VM deletion events,
- Δt is the observation time window.

This means that semantic measurement is approximated by the temporal frequency of VM birth and death. Frequent VM generation and deletion indicate semantic reorganization, instability, exploration, or turbulence.

A weighted version is:

$$M_{sem}(t) = \frac{1}{\Delta t} \sum_i w_i e_i(t),$$

where w_i is the semantic importance of VM i, and $e_i(t)$ indicates whether a birth or death event occurred.

This approximation does not measure semantic truth directly. Rather, it measures semantic structural activity. Nevertheless, it provides an observable and implementable basis for semantic temperature estimation and turbulence detection.

6 Semantic Entropy

If the semantic graph exists in a superposition:

$$|G\rangle = \sum_i \alpha_i |G_i\rangle,$$

then the probability of observing graph state G_i is:

$$p_i = |\alpha_i|^2.$$

The semantic graph entropy is defined as:

$$H_G = -\sum_i p_i \log p_i = -\sum_i |\alpha_i|^2 \log |\alpha_i|^2.$$

Distance-based semantic entropy is:

$$H_D = \sum_{i,j} p_i p_j D_H(i, j).$$

A small H_D indicates semantic concentration and coherence, while a large H_D indicates semantic dispersion, fragmentation, or turbulence.

7 Semantic Temperature

Semantic temperature is introduced by analogy with thermodynamics:

$$\frac{1}{T_{sem}} = \frac{\partial H_{sem}}{\partial E_{sem}}.$$

Equivalently:

$$T_{sem} = \left(\frac{\partial H_{sem}}{\partial E_{sem}} \right)^{-1}.$$

For practical implementation, semantic temperature may be estimated by fluctuation intensity:

$$T_{sem} = a\,Var(D_H) + b\,Var(P_{ij}) + c\,Var(S_i) + d\,R_{collapse} + e\,R_{birth/death}.$$

The semantic measurement $M_{sem}(t)$ directly contributes to $R_{birth/death}$. Thus:

$$T_{sem} \sim \text{semantic fluctuation intensity.}$$

The semantic temperature regions are:

$$T_{sem} < T_1 \Rightarrow \text{automation-dominant regime,}$$

$$T_1 \leq T_{sem} < T_2 \Rightarrow \text{adaptive reasoning regime,}$$

$$T_2 \leq T_{sem} < T_3 \Rightarrow \text{creative exploration regime,}$$

$$T_{sem} \geq T_3 \Rightarrow \text{stabilization or cooldown regime.}$$

The thresholds T_1, T_2, T_3 are not universal constants. They must be calibrated according to task type, user goals, safety requirements, and empirical behavior.

8 Stability- and Goal-Based Semantic Potential

The previous formulation introduced semantic energy as:

$$E_{sem} = \sum_i p_i V(S_i) + \sum_{i,j} w_{ij} D_H(i,j) + \lambda C_{comp}.$$

A key unresolved issue was the definition of semantic potential $V(S_i)$. This paper introduces the temporary definition:

$$V(S_i) = Instability(S_i).$$

Equivalently:

$$V(S_i) = 1 - Stability(S_i).$$

However, stability alone is insufficient because a semantic state can be stable but incorrect, misleading, or misaligned with human goals. Therefore, the definition is extended as:

$$V(S_i|G) = Instability(S_i) + GoalMismatch(S_i, G),$$

where G is the human-defined goal.

A more complete formulation is:

$$V(S_i|G,U,C) = \alpha Instability(S_i) + \beta D_{goal}(S_i,G,U) + \gamma Risk(S_i) + \kappa ConstraintViolation(S_i,C) - \delta HumanApprova$$

Here:

- G is the human-defined goal,
- U is the user dialogue history,
- C is the shared constraint database,
- $D_{goal}(S_i, G, U)$ is the goal distance learned from user interaction,

- $ConstraintViolation(S_i, C)$ is the degree of violation against security, legal, safety, privacy, or operational constraints.

Usually:

$$\kappa \gg \alpha, \beta, \gamma,$$

because security, legal, and safety violations should dominate ordinary goal optimization.

Thus, semantic energy becomes:

$$E_{sem} = \sum_i p_i V(S_i|G, U, C) + \sum_{i,j} w_{ij} D_H(i, j) + \lambda C_{comp}.$$

This means:

$$E_{sem} = \text{semantic instability} + \text{goal mismatch} + \text{constraint violation} + \text{semantic dispersion} + \text{computational cost}.$$

9 Semantic Free Energy

The semantic free energy is defined as:

$$F_{sem} = E_{sem} - T_{sem} H_{sem}.$$

Minimizing F_{sem} does not simply mean minimizing entropy. Instead, it means maintaining a useful balance between semantic stability and semantic diversity.

At low temperature, the system favors stable automation. At medium temperature, it supports adaptive reasoning. At high temperature, it permits creative exploration. At excessive temperature, stabilization mechanisms must intervene.

10 Goal-Relative Semantic Truth

In many AI systems, absolute truth exists only in limited domains such as mathematics, formal logic, and certain empirically verifiable physical laws. However, when an AI system is developed for human-defined purposes, correctness is often determined by whether a semantic state contributes to the achievement of those purposes.

Therefore, semantic truth is defined as goal-relative truth:

$$Truth_{sem}(S_i|G) = Fitness(S_i, G).$$

Thus:

$$Correct(S_i) \neq AbsoluteTruth(S_i),$$

but rather:

$$Correct(S_i) = Useful(S_i, G).$$

Semantic verification is therefore reformulated as:

$$V_{sem}(S_i, G) = \text{Does semantic state } S_i \text{ help achieve goal } G?$$

A practical verification score is:

$$V_{sem}(S_i, G, U, C) = w_1 GoalFit(S_i, G, U) + w_2 ConstraintFit(S_i, C) + w_3 HumanApproval(S_i) + w_4 VMConsensus(S$$

This definition shifts semantic verification from absolute truth judgment to goal-aligned validity judgment.

11 Goal Distance Learning from User Dialogue

The goal distance $D_{goal}(S_i, G, U)$ does not need to be defined as an absolute metric from the beginning. Instead, it can be learned from end-user dialogue as a ranking problem.

If the user prefers semantic state S_a over S_b, we write:

$$S_a \succ S_b.$$

This implies:

$$D_{goal}(S_a, G, U) < D_{goal}(S_b, G, U).$$

Thus, goal distance can be learned from pairwise or ranked preferences:

$$S_1 \succ S_3 \succ S_2 \Rightarrow D_{goal}(S_1, G, U) < D_{goal}(S_3, G, U) < D_{goal}(S_2, G, U).$$

A practical formulation is:

$$GoalFit(S_i, G, U) = f_\theta(S_i, G, U),$$

and:

$$D_{goal}(S_i, G, U) = 1 - GoalFit(S_i, G, U).$$

The user dialogue history U provides signals such as:

- preference between alternatives,
- rejection of unsuitable interpretations,
- requests for safer, more practical, or more creative outputs,
- approval of goal-aligned results,
- correction of semantic misunderstanding.

The goal representation may be updated gradually:

$$G_t = (1 - \eta)G_{t-1} + \eta \Delta G_t,$$

where η controls how quickly the system adapts to new user preferences.

This prevents the system from overreacting to short-term dialogue while still allowing it to adapt to evolving user goals.

12 Shared Constraint Database

Goal learning from user dialogue should not override security, legal, safety, privacy, or operational constraints. Therefore, Quantum Hyper-Aether introduces a shared constraint database:

$$ConstraintDB = (SecurityRules, LegalRules, SafetyRules, PrivacyRules, OperationalLimits).$$

The constraint fitness is:

$$ConstraintFit(S_i, C) = Check(S_i, ConstraintDB).$$

The database contains:

- security policies,
- access-control rules,
- legal and regulatory constraints,
- privacy rules,
- safety rules,
- operational limits,
- audit requirements,
- forbidden semantic or action states,
- evidence requirements.

The constraint database should not be treated as ordinary self-learning memory:

$$ConstraintDB \neq SelfLearningMemory.$$

VMs may query the constraint database, but they should not freely rewrite it. Updates require versioning, authorization, audit logs, and human or organizational approval.

If a candidate semantic state violates critical constraints:

$$ConstraintViolation(S_i, C) \geq \theta_{critical},$$

then:

$$Reject(S_i).$$

Thus, even if:

$$GoalFit(S_i, G, U)$$

is high, the candidate state must be rejected if it violates mandatory constraints.

13 LLM-Based Semantic Verification

Existing large language models can be used as semantic teachers or semantic verifiers. They can assist in:

- assigning semantic labels,
- generating associative keywords,
- estimating semantic similarity,
- detecting contextual inconsistency,
- explaining semantic collapse,
- proposing alternative interpretations,
- standardizing semantic verification procedures.

An LLM-based semantic verification function may be defined as:

$$V_{LLM}(S_i, Ctx, R, P),$$

where:

- S_i is the semantic state,
- Ctx is the context,
- R is reference knowledge,
- P is system policy.

However, LLM verification is not absolute. LLMs may contain biases, outdated knowledge, hallucinations, and domain-specific weaknesses. Therefore, LLM-based verification must be combined with VM consensus, external evidence, constraint checking, and human audit.

14 Multi-VM Consensus Verification

A stable semantic state is not necessarily correct:

$$Stable(S_i) \neq Correct(S_i).$$

To reduce the risk of stable but incorrect semantic states, Quantum Hyper-Aether uses multi-VM consensus.

For candidate semantic state S, each VM produces a vote:

$$Vote_i(S) \in \{0, 1\}.$$

A simple majority consensus is:

$$Consensus(S) = \frac{1}{N} \sum_{i=1}^{N} Vote_i(S).$$

However, simple majority voting is insufficient when VMs share the same bias, training source, or semantic attractor. Therefore, weighted consensus is used:

$$Score(S) = \sum_i w_i Vote_i(S).$$

The weight of VM i may be:

$$w_i = a\, Reliability_i + b\, Diversity_i + c\, Expertise_i + d\, PastAccuracy_i - e\, Instability_i.$$

This gives higher influence to VMs that are reliable, diverse, specialized, historically accurate, and stable.

Semantic collapse may be permitted only when:

$$Consensus(S) > \theta_{collapse}.$$

If the consensus score is below the threshold, the system may:

- postpone collapse,
- request more VM evaluations,
- invoke LLM verification,
- consult external evidence,
- trigger human review.

15 Human-in-the-Loop Semantic Audit

Even VM consensus and LLM verification do not guarantee correctness. Therefore, stable semantic states should be periodically sampled and presented to humans for approval.

The process is:

$$Stable(S_i) \to Sample(S_i) \to HumanReview \to Approved/Rejected/Modified.$$

The sampling probability may be:

$$P_{sample}(S_i) = \alpha Risk_i + \beta Importance_i + \gamma Change_i + \delta LowConfidence_i.$$

Thus, the system prioritizes semantic states that are:

- high risk,
- high impact,
- recently collapsed,
- low in LLM confidence,
- low in VM consensus,
- generated in a high-temperature region.

Human feedback is used to update:

$$HumanApproval(S_i) \to Update(V_{sem}, D_H, V(S_i), w_i, \theta_{collapse}).$$

Thus, human audit becomes a learning signal for semantic verification, semantic distance calibration, VM reliability estimation, collapse threshold adjustment, and semantic potential correction.

16 Hierarchy of Semantic Truth

Quantum Hyper-Aether distinguishes several levels of truth:

Truth Type	**Meaning**	**Verification Method**
Absolute Truth	Mathematics, formal logic, some physical laws	Formal proof or experiment
Factual Truth	Real-world facts and observations	External evidence
Goal-Relative Truth	Usefulness for human-defined goals	Goal fitness evaluation
Operational Truth	Accepted in system operation	Human audit and policy approval
Semantic Truth	Contextual coherence and meaning validity	LLM and VM consensus
Constraint Truth	Compliance with mandatory rules	Constraint database check

Table 1: Hierarchy of truth in Quantum Hyper-Aether.

The most important notion for an AI-native operating system is not always absolute truth, but goal-relative truth under operational constraints.

17 Entropy–Temperature–Verification Control

A semantic temperature controller may be defined as:

$$u_T(t) = K_p(T_{target} - T_{sem}(t)) + K_i \int_0^t (T_{target} - T_{sem}(\tau))d\tau + K_d \frac{d}{dt}(T_{target} - T_{sem}(t)).$$

In the extended framework, this controller must also consider verification, constraints, and human approval:

$$u(t) = f(T_{sem}, H_{sem}, E_{sem}, V_{sem}, Consensus, ConstraintFit, HumanApproval).$$

The controller adjusts:

- VM birth rate,
- VM deletion rate,
- routing randomness,
- collapse threshold,
- exploration probability,
- semantic memory activation range,
- LLM verification frequency,
- human audit sampling rate,
- constraint checking intensity.

Thus, the system does not merely stabilize semantic dynamics. It stabilizes semantic dynamics toward human-defined goals under mandatory constraints.

18 Collapse Reversibility and Auditability

Semantic collapse should not be treated as an irreversible final decision. Instead, it should be treated as an auditable and reversible provisional commitment.

For each collapse event:

$$|\Psi\rangle \to |G_k\rangle,$$

the system stores:

- candidate semantic states before collapse,
- VM consensus score,
- LLM verification result,
- constraint database decision,
- external evidence references,
- human audit state,

- reason for collapse,
- rollback possibility.

If later evidence contradicts the collapsed state, the system may perform:

$$Rollback(G_k) \rightarrow Reopen(\{G_1, G_2, \ldots, G_n\}).$$

Thus, collapse becomes:

$$Collapse = ProvisionalCommitment + AuditTrail + Reversibility.$$

19 Overall Workflow

The integrated workflow is:

$$Input \rightarrow UserDialogue \rightarrow GoalExtraction \rightarrow ConstraintCheck \rightarrow CandidateSemanticStates \rightarrow GoalRanking$$

This means that Quantum Hyper-Aether operates as a closed-loop semantic control system.

20 Remaining Challenges

The proposed extensions significantly reduce several unresolved issues, but they do not eliminate all challenges.

The remaining challenges are:

1. **Prototype validation.** The system must be tested with a small prototype to verify whether VM birth/death frequency correlates with real semantic activity and temperature.
2. **Goal ranking stability.** The goal distance learned from user dialogue must distinguish short-term preference from long-term goal alignment.
3. **Constraint database governance.** The shared constraint database requires version control, authorization, auditability, regional rule selection, and protection against unauthorized modification.
4. **Consensus bias.** Multi-VM consensus can fail if VMs share the same LLM bias, data bias, or semantic attractor. Diversity and independence of VMs must be measured.
5. **Human audit load.** Human-in-the-loop verification must be sampled carefully to avoid excessive burden while still catching high-risk stable-but-wrong semantic states.
6. **Temperature threshold calibration.** The thresholds T_1, T_2, T_3 must be empirically calibrated for automation, adaptive reasoning, creative exploration, and stabilization.
7. **Collapse reversibility.** Collapse must be made auditable, reversible, and explainable to prevent incorrect semantic attractors from becoming permanent.
8. **External evidence integration.** Factual and operational truth require integration with databases, logs, sensors, documents, legal sources, standards, and benchmark results.
9. **Scalability.** Semantic distance computation, VM consensus, LLM verification, constraint checking, and human audit queues must scale through hierarchical clustering.

These challenges indicate that the next stage of Quantum Hyper-Aether research is not only theoretical refinement but also prototype implementation and empirical evaluation.

21 Advantages of the Extended Framework

The extended Quantum Hyper-Aether framework provides:

- observable semantic measurement through VM birth/death frequency,
- implementable semantic temperature estimation,
- stability- and goal-based semantic potential,
- goal-relative semantic truth,
- user-dialogue-based goal distance learning,
- shared constraint database for mandatory rules,
- LLM-assisted semantic verification,
- multi-VM consensus verification,
- human-in-the-loop semantic audit,
- reduced risk of stable but incorrect semantic states,
- semantic collapse control and reversibility,
- improved alignment with human-defined objectives.

22 Conclusion

This paper extended Quantum Hyper-Aether by introducing practical mechanisms for semantic measurement, semantic verification, semantic potential definition, goal-distance learning, constraint management, and goal-aligned control.

Semantic measurement is temporarily approximated by the temporal frequency of VM birth and death. Semantic potential is defined as semantic instability plus goal mismatch, risk, and constraint violation. Semantic truth is reinterpreted as goal-relative truth rather than universal absolute truth. Goal distance is learned from end-user dialogue as a ranking problem. Stable semantic states are verified through multi-VM consensus, LLM-based verification, external evidence, constraint database checks, and human-in-the-loop audit.

The resulting framework transforms Quantum Hyper-Aether into a semantic-native, self-stabilizing, auditable, constraint-aware, and goal-aligned AI computing architecture. It does not assume that stability alone implies correctness. Instead, correctness emerges from the interaction of semantic stability, goal fitness, mandatory constraints, VM consensus, LLM verification, human approval, and external evidence.

The central shift is from semantic theory alone to goal-aligned, auditable, controllable AI-native operation.

References

[1] John von Neumann, *Mathematical Foundations of Quantum Mechanics*, Princeton University Press, 1955.

[2] Richard P. Feynman, "Space-Time Approach to Non-Relativistic Quantum Mechanics," *Reviews of Modern Physics*, 1948.

[3] Kurt Gödel, "On Formally Undecidable Propositions of Principia Mathematica and Related Systems," 1931.

[4] Saunders Mac Lane, *Categories for the Working Mathematician*, Springer, 1971.

[5] Ilya Prigogine and Isabelle Stengers, *Order Out of Chaos*, Bantam Books, 1984.

[6] Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, 1948.

[7] E. T. Jaynes, "Information Theory and Statistical Mechanics," *Physical Review*, 1957.

[8] Geoffrey Hinton, "Learning Distributed Representations of Concepts," Proceedings of the Eighth Annual Conference of the Cognitive Science Society, 1986.

[9] Ashish Vaswani et al., "Attention Is All You Need," NeurIPS, 2017.

[10] Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control*, Viking, 2019.

[11] Paul Christiano et al., "Deep Reinforcement Learning from Human Preferences," NeurIPS, 2017.