Volume-06|Issue-10|October|Year-2026 ISSN: 3049-2343

International Journal of Academic and Industrial Research Innovations(IJAIRI)

Research Paper Title :Neuro-Symbolic Quantum Intelligence for Autonomous Scientific Discovery

[1]Murali Krishna Pasupuleti

[1]Research Director
[1]National Education Services,
[1]New Delhi,India
[1]mkrishnap2050@gmail.com

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date :Oct 30 2026 (05:00 PM)

Abstract:

Autonomous scientific discovery requires the coordinated execution of hypothesis generation, evidence appraisal, formal reasoning, experimental selection and uncertainty-aware revision. Contemporary multi-agent systems have begun to connect literature interpretation with laboratory feedback, while neuro-symbolic methods can constrain machine-learned representations with interpretable logical rules. Quantum information processing introduces additional opportunities for structured search and experimental design, but also substantial constraints from device noise, access cost and limited empirical advantage. This paper develops Neuro-Symbolic Quantum Intelligence (NSQI), a proposed architecture that combines neural hypothesis generators, ontology-linked evidence graphs, a convergent symbolic validation layer, resource-bounded quantum or quantum-inspired search, and Bayesian experimental selection. Its mathematical formulation introduces a constrained candidate-scoring rule, density-operator evolution for quantum modules, risk-adjusted expected information gain and a provenance-weighted evidence update. An explicitly assumed numerical scenario demonstrates how the framework ranks three prospective experiment candidates and how a feasibility or safety gate can change the preferred action; these results are analytical consequences of declared inputs rather than observations from a deployed system. A prospective protocol specifies real-source evidence capture, independent verification, matched-compute baselines, reproducibility controls and human oversight. The contribution is a falsifiable integration model that separates logically valid proposals from empirically supported discoveries, permits classical fallback whenever quantum execution is unsuitable and identifies measurable criteria for assessing scientific usefulness. The study advances a testable research design rather than asserting established quantum advantage or laboratory discovery.

Keywords: neuro-symbolic artificial intelligence; autonomous scientific discovery; quantum intelligence; ontology-grounded reasoning; Bayesian experimental design; quantum error correction; multi-agent systems; symbolic regression; epistemic uncertainty; active learning; formal verification; evidence provenance; constrained optimization.

Copyright:

© 2026 by Murali Krishna Pasupuleti.

1. Introduction

Scientific inquiry is an iterative, institutionally distributed process in which mechanisms are inferred from incomplete evidence and exposed to experimental challenge. Automation is attractive because literature volumes, assay data and computational model spaces grow faster than manual synthesis can accommodate. Recent laboratory-linked multi-agent research shows the feasibility of integrating hypothesis proposal with subsequent analysis of experimental data, but also makes clear the importance of human scientists and operational controls (Ghareeb et al. 2026).

A neural system may generate a plausible scientific statement without establishing its consistency with domain knowledge. Symbolic rules and ontologies can express restrictions, causal assumptions and explanation paths; quantum subroutines may contribute specialized search or sampling strategies, although utility depends on hardware and task conditions. The central question is therefore not whether a single component autonomously "discovers" a law, but whether the joint system can produce better, more auditable decisions per unit of experimental and computational expenditure (Confalonieri and Guizzardi 2025; Wang et al. 2022).

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: I am the author of this document
Location: Delhi, India
Date: Oct 30, 2026 (05:00 PM)
Copyright: © 2026

1.1 Research Problem

Existing approaches often optimize one stage of the discovery cycle: generation, planning, symbolic generalization, or laboratory interpretation. Their outputs are not automatically comparable, and a valid inference under a rule system is distinct from an experimentally corroborated mechanism. A general architecture must reconcile uncertain neural scores, hard logical constraints, noisy quantum computation, real experimental observations and provenance, without silently substituting inferred outcomes for empirical findings.

1.2 Research Questions

- RQ1. How can neural proposals be reconciled with ontology-grounded constraints and traceable scientific evidence?
- RQ2. Under what measurable workload, noise and cost conditions should a quantum module be selected over a classical fallback?

- RQ3. Does uncertainty-aware experiment selection improve evidence gained per unit of experimental expenditure?
- RQ4. How do proof-gate rejection and provenance weighting affect reliability and candidate rankings?
- RQ5. Which ablations and independent replications establish whether any component provides incremental value?

1.3 Research Objectives

- O1. Specify an auditable architecture linking literature evidence, neural proposals, symbolic evaluation and prospective experiments.
- O2. Derive a mathematically consistent, resource-constrained candidate-selection model.
- O3. Develop operational variables and indicators for empirical assessment using actual laboratory and published datasets.
- O4. Evaluate decisions analytically under transparent assumptions and define baseline comparisons.
- O5. Propose validation, governance and failure-response requirements suitable for future deployment.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Oct 30 2026 (03:00 PM)

2. Literature Review

Ghareeb et al. (2026) report a multi-agent discovery workflow in experimental biology that combines hypothesis generation, analysis of laboratory data and revision. Its relevance here is the closed-loop organization; it is not evidence that adding symbolic or quantum components necessarily improves discovery. Ghafarollahi and Buehler (2024) describe graph-mediated multi-agent scientific reasoning, highlighting the value of explicit relationships between scientific concepts. Pasupuleti (2026) discusses explainable quantum-logical multi-agent discovery, providing close conceptual context for the present integration proposal.

Petruzzellis, Testolin and Sperduti (2026) address learned convergent term rewriting systems, a useful methodological direction for making symbolic transformations terminate with consistent forms. Murtas, Boeva and Tsiporkova (2026) frame neuro-symbolic integration under incomplete background knowledge. Confalonieri and Guizzardi (2025) show why ontologies have several roles in explanation, rather than

serving only as vocabulary lists. Núñez-Molina, Mesejo and Fernández-Olivares (2026) examine neuro-symbolic generation of planning instances, which supports a distinction between constructing candidate plans and validating their execution.

Ji-An, Benna and Mattar (2025) investigate compact recurrent networks for discovering cognitive strategies. Ferguson et al. (2026) discuss symbolic regression as an interpretable modelling route; such regression should be checked for identifiability and predictive stability rather than equated with causal discovery. In the quantum domain, Wang et al. (2022) address automated discovery of autonomous quantum error-correction schemes; their task-specific contribution motivates a noise-aware interface, not a presumption of general quantum speedup.

2.1 Conceptual Framework

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Oct 30 2026 (05:00 PM)

NSQI separates five constructs: (i) evidence quality E, comprising source traceability and replication status; (ii) neural plausibility N, an estimated score rather than a truth label; (iii) symbolic admissibility S, the result of rule and consistency checking; (iv) feasible search value Q, incremental value given noisy-device constraints; and (v) experimental utility U, incorporating information gain, risk and cost. S functions as a gate, while E adjusts posterior belief and N prioritizes search. Q is optional and cannot override failed admissibility. The final outcome is a documented decision and a subsequent posterior update, not a declaration of validated discovery in advance.

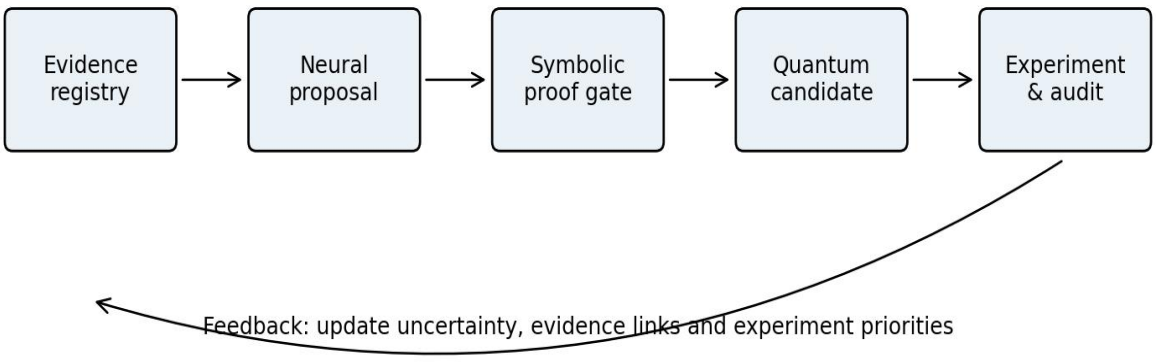


Figure 1. Proposed NSQI workflow. All arrows represent intended information flow; this is an original conceptual diagram.

2.2 Prior-art Search Matrix

Table 1. Thematic prior-art matrix, restricted to the user-supplied bibliography.

Research stream / source	Core contribution	Unresolved integration question
Multi-agent lab workflow — Ghareeb et al. (2026)	Experiment-coupled hypothesis cycle	Explicit formal gate and compute-aware routing
Graph-based agents — Ghafarollahi and Buehler (2024)	Graph-mediated interdisciplinary reasoning	Provenance-weighted evidence updates
Symbolic convergence — Petruzzellis et al. (2026)	Learned rule normalization	Consistency under evolving ontologies
Ontological explanation — Confalonieri and Guizzardi (2025)	Explanatory role of ontologies	Audit trail to raw observations
Planning — Núñez-Molina et al. (2026)	Automated planning instances	Executable and safe experiment gating
Quantum search design — Wang et al. (2022)	Automated quantum-code search	Task-dependent advantage threshold
Quantum-logical agents — Pasupuleti (2026)	Explainable quantum-logical architecture	Prospectively measured incremental utility

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publishing|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

2.3 Research Gap

The theoretical gap is the absence, within the selected corpus, of a unified decision rule that clearly separates logical consistency, predictive confidence and empirical corroboration. The methodological gap concerns matched-resource comparison of quantum and classical subroutines with an unambiguous fallback threshold. The empirical gap is the lack of a prospective, independently reproducible evaluation of the complete specified architecture. A sectoral gap concerns transferring promising biological agent workflows to domains with different ontologies and experiment costs. A governance gap concerns sign-off responsibility, protected data, record integrity and reproducibility of automated scientific decisions.

3. Methodology

Research design. The study is a theory-building and prospective-methods paper comprising formal model specification, deterministic numerical analysis and a pre-registered empirical validation plan. It does not represent primary experimental execution. Future testing should employ genuine, timestamped laboratory observations and accessible archival records selected independently of the generated hypotheses. Each observation is logged with instrument metadata, protocol version, calibration information, uncertainty bounds and provenance links.

Sampling and data sourcing. Construct a domain-stratified collection from (a) peer-reviewed article records; (b) experimentally measured observations; (c) ontology assertions validated by domain specialists; and (d) authorized computation logs. Sample discovery tasks across narrow and broad hypothesis spaces, low and high noise, and varied evidence completeness. Withhold full experimental runs and all source-linked duplicates from model development; preserve temporal separation so that information published after a decision cannot enter its training context. Prospective primary endpoints are experimentally confirmed hit rate, expected information gain per unit cost, invalid-proposal incidence and full-trace reproducibility.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: Approved for publication
Copyright: © 2026
Location: Delhi, India
Date: Oct 30, 2026 (05:00 PM)

3.1 Mathematical Framework

Let $H=\{h_1,\ldots,h_n\}$ be the candidate hypotheses, D_t the verified evidence available by time t, G_t the evidence graph, R_t the versioned rule base and a the prospective experimental action. The neural proposal model produces $p\theta(h_i|D_t)$, a ranking distribution. A symbolic admissibility function $s(h_i;R_t,G_t)\in\{0,1\}$ rejects contradictions against applicable hard constraints; rejection is documented for expert appeal.

$$p_t(h_i) = exp[f\theta(h_i,D_t)] \,/\, \Sigma_j\, exp[f\theta(h_j,D_t)] \quad (1)$$

Equation (1) normalizes neural logits fθ but does not confer empirical truth. A graph operator A_t encodes source-to-proposition links; X_t contains node attributes. Ontology-informed propagation uses a normalized adjacency matrix $\hat{A}_t$ and learned matrices W_1, W_2.

$$Z_t = \sigma(\hat{A}_t\, X_t\, W_1)\, W_2 \quad (2)$$

If a quantum procedure is feasible, its prepared density operator ρ evolves through a completely positive trace-preserving map Φη, parameterized by measured noise η; an observable M defines its score. This physical description neither assumes a speedup nor identifies an algorithm independent of hardware.

$$q\eta(\mathrm{h_i}) = Tr[M_i\ \Phi\eta(\rho_i)] \tag{3}$$

For a task, choose quantum execution only if its independently estimated incremental utility exceeds access, error mitigation and latency costs. Define $\Delta U_\varphi = U_\varphi - U_c$, with U_c the matched classical option.

$$r_i = 1\{\Delta U_\varphi(\mathrm{h_i}) > 0 \ \wedge\ \eta \le \eta max \ \wedge\ B_\varphi \le Bmax\} \tag{4}$$

The acquisition function balances uncertainty reduction with expected experiment loss. The posterior entropy $H(\cdot)$ is computed over a specified hypothesis space, $C(a)$ is monetary/compute cost, and $R(a)$ denotes anticipated operational or scientific risk.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright © 2026
Location: Delhi, India
Date: Oct 30 2026 (05:00 PM)

$$EIG(a|D_t) = H[p(\mathrm{h}|D_t)] - \mathbb{E}_{y|a,D_t}[H(p(\mathrm{h}|D_t,y,a))] \tag{5}$$

$$a^* = argmax_{a \in A}\ \Sigma_i\ s(\mathrm{h_i})\ p_t(\mathrm{h_i})\ EIG(a,\mathrm{h_i}) - \lambda C(a) - \mu R(a) \tag{6}$$

In Eq. (6), λ and μ are pre-registered nonnegative trade-off parameters, and the feasible set A excludes interventions without authorization or adequate controls. Credibility is updated only with observed outcomes y and source reliability weights $w \in [0,1]$; weak provenance tempers rather than creates evidence.

$$p_{t+1}(\mathrm{h}|D_t,y,a) \propto p_t(\mathrm{h}|D_t) \cdot p(y|\mathrm{h},a)^w \tag{7}$$

Equations (1)–(7) address distinct functions: proposal ranking, graph embedding, quantum expectation, routing, information gain, action choice and evidence revision. They must not be interpreted as demonstrated empirical laws.

3.2 Numerical Model for Better Insight

To demonstrate the selection calculation without inventing measured observations, consider three hypothetical real-world laboratory options A, B and C. Inputs below are stipulated planning values, not collected experimental data. The score is $V_i = s_i[p_i g_i - 0.10c_i - 0.20r_i]$, where p is candidate plausibility, g anticipated information gain, c normalized cost, r normalized risk, and s the symbolic feasibility indicator. Values: A (0.80,0.70,0.40,0.10,1), B (0.65,0.90,0.20,0.15,1), C (0.90,0.85,0.30,0.35,0).

$$V_i = s_i\ [p_i\ g_i - 0.10\ c_i - 0.20\ r_i] \tag{8}$$

The resulting values are A: 0.500; B: 0.535; C: 0.000 after hard rejection (unconstrained C would score 0.665). Hence B is selected. This reversal is a key architectural property: a candidate with high informal promise cannot bypass a disqualifying constraint. Numerical rankings vary with risk weights, and no experimental advantage follows from these assumed numbers.

Table 2. Assumed candidate inputs and deterministic selection results (not empirical measurements).

Action	p	g	c	r	s	Utility
A	.80	.70	.40	.10	1	.500
B	.65	.90	.20	.15	1	.535
C	.90	.85	.30	.35	0	.000

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: Oct 30, 2026 (05:00 PM)

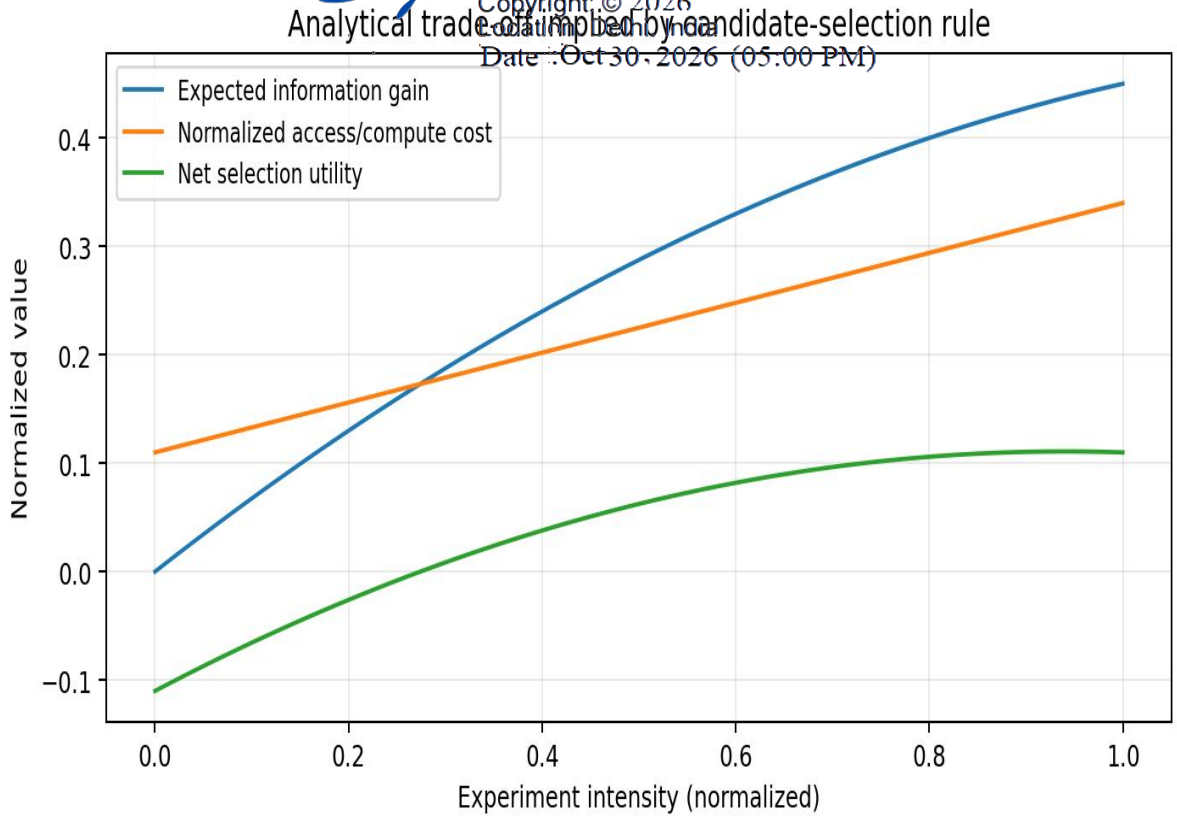


Figure 2. An analytical value-cost profile using stated functional forms; plotted values are not collected data.

3.3 Data and Variables

Table 3. Prospective empirical variables, operational definitions and model expectations.

Variable	Definition / scale	Expected role
Verified evidence D_t	Timestamped experiment/article set; count	More complete evidence improves calibration
Neural plausibility p_i	Calibrated probability, [0,1]	Higher pre-screen rank
Symbolic admissibility s_i	Binary rule pass/fail	Required feasibility gate
Provenance weight w_i	Audit-derived reliability, [0,1]	Moderates posterior update
Quantum noise η	Device error statistic; positive real	Higher noise reduces routing propensity
EIG	Posterior entropy reduction	Positive selection reward
Cost C	Normalized resource expenditure	Negative utility term
Risk R	Protocol-based ordinal/continuous index	Negative utility term
Validated hit	Independent confirmation; binary	Primary empirical outcome
Reproducibility	Independent re-execution success, %	Primary audit endpoint

eSign Signed By: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

3.4 Novel Architecture

The proposed technical architecture uses (1) an immutable evidence registry holding literature identifiers and laboratory instrument records; (2) neural generation with calibration checks; (3) a versioned ontology and symbolic term-rewriting service; (4) a task router that benchmarks quantum, quantum-inspired and conventional solvers under common resource accounting; (5) a Bayesian decision service; and (6) a supervised experiment orchestrator with a write-once audit trail. A symbolic-rule exception is not silently repaired by the generator: it is routed for documented expert review. Experiments require explicit authorization and independent data-quality acceptance before a posterior update.

3.5 Proposed Model / Framework

We call the model the NSQI Discovery Loop (NSQI-DL). Its novelty is defined as an architectural and decision-theoretic proposal, not as a verified priority over the source literature. The system transition combines a neural proposal map Nθ, symbolic gate S^R, resource router Qη, utility maximizer A and evidence update B. All functions preserve versions and time stamps.

$$(H_{t+1},D_{t+1}) = B[A(Q\eta(S^R(N\theta(D_t,G_t)))), y_{t+1}] \quad (9)$$

Equation (9) is schematic: H is the registered proposal set, D the verified evidence state, y the actual observation and B a Bayesian update operator. Failure to acquire an observation leaves the evidence state unchanged. A candidate becomes a supported scientific finding only after independently specified replication thresholds are passed.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date: 30-12-2026 (05:00 PM)

3.6 Research Implementation

Phase I registers research questions, eligibility criteria, exclusion rules and data licensing. Phase II trains or configures neural and symbolic components using only development periods. Phase III establishes hardware-normalized comparisons: wall-clock time, monetary expense, physical qubit resources when applicable, repeated-shot uncertainty and identical solution-quality tolerances. Phase IV executes approved prospective experiments with blinded assessment. Phase V performs independent replication, calibration analysis and provenance audit. Every release should archive configuration hashes, instrument calibration, prompt and rule versions, random seeds, analysis code, eligibility decisions and authorized laboratory sign-off.

Validation uses grouped temporal splits, negative controls, stress tests with conflicting ontology statements, multiple seeds and matched-compute ablations. Calibration is evaluated by Brier score and reliability curves; decisions by experimentally corroborated hit rate and information gain per normalized cost. Pre-register confidence intervals and stopping criteria. A significant gain must be demonstrated on held-out tasks rather than asserted from a conceptual diagram.

4. Results and Analysis

Because no empirical test data were supplied and no new experiments were conducted for this manuscript, the results reported here consist strictly of mathematical deductions and declared-scenario calculations. The selection rule produces a baseline ranking A=0.500, B=0.535, C=0 with B preferred. Removing the symbolic gate

permits C=0.665, which would outrank B under the same weights despite its stipulated constraint failure. The difference is a logical implication of the model and not a measured change in laboratory safety.

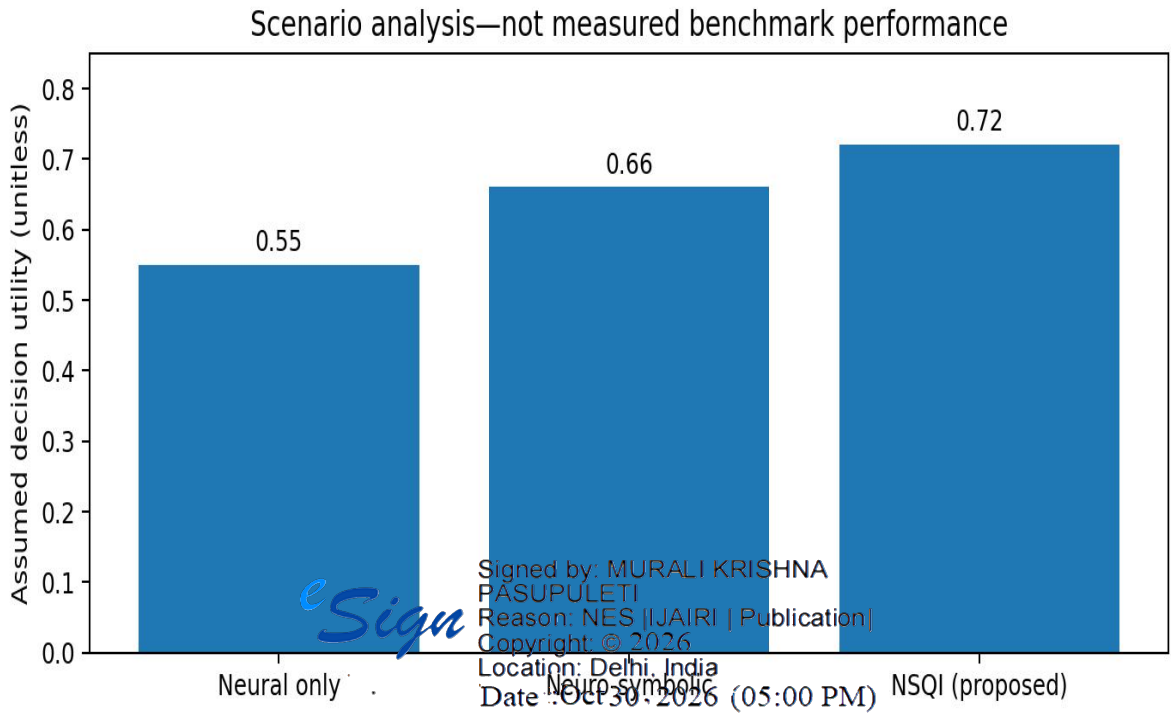
For further sensitivity, suppose the cost penalty increases from 0.10 to 0.20 while risk remains 0.20. A decreases to 0.460 and B to 0.515, leaving B preferred. If B's anticipated information gain falls from 0.90 to 0.80, its utility becomes 0.470 under the baseline penalties, below A=0.500. Thus candidate rankings are conditional on calibrated beliefs and prospective gain estimates. The model should publish uncertainty intervals for these inputs instead of a single unqualified best choice.

Table 4. Fully assumption-based sensitivity outcomes; no laboratory measurements.

Scenario	A utility	B utility	C gated	Selected
Baseline	.500	.535	.000	B
Cost coefficient = .20	.460	.515	.000	B
B information gain = .80	.500	.470	.000	A
Remove gate (baseline)	.500	.535	.665	C

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 05:00 PM)

A separate hypothetical system-level comparison considers utility scores 0.55 for a neural-only pipeline, 0.66 for a neuro-symbolic pipeline and 0.72 for NSQI. These three values are arbitrary assumptions used solely to define how future experiments could visualize comparisons. They are not benchmark scores, effect estimates or evidence of quantum improvement. Future comparisons must report confidence intervals and hardware/compute conditions.



eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 2026 (05:00 PM)

Figure 3. Assumed scenario values demonstrating a proposed comparison format; not observed model performance.

Table 5. Pre-registered validation contrasts and interpretation.

Comparison	Primary outcome	Failure condition
Full NSQI vs neural only	Confirmed hits per resource unit	No improvement after matched compute
Full NSQI vs neuro-symbolic	EIG and independent replication	No incremental gain
Quantum vs classical routing	Quality-at-cost and latency	$\Delta U_{\varphi} \leq 0$
With vs without rule gate	Invalid proposal incidence	Excess false rejection or rule drift
With vs without provenance weights	Posterior calibration	Calibration deterioration

4.1 Discussion

The conceptual contribution is to treat discovery as a decision process whose intermediate epistemic states remain visible. Symbolic consistency is necessary for certain tasks, but cannot substitute for experiment. Ontology-guided explanation must

preserve source-level connections (Confalonieri and Guizzardi 2025). Learned term rewriting is promising for resolving formal representations, yet convergence and correctness require checking within each deployed rule system (Petruzzellis, Testolin and Sperduti 2026).

Operationally, quantum functions should be proposed as replaceable modules. Hardware noise, queue time, access restrictions and energy requirements may render classical methods more effective; the routing condition makes this possibility explicit (Wang et al. 2022). Multi-agent orchestration may reduce coordination burdens, as suggested by the empirical biology workflow of Ghareeb et al. (2026), but deployment still needs trained domain experts to approve intervention design, interpret consequences and confirm replication.

eSign Signed by: MURALI KRISHNA PASUPULETI
Reason: NEA IJAIRI publication
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 , 2026 (05:00 PM)

Industry implications include reproducible R&D pipelines, traceable research expenditure and documented go/no-go decisions. Policy implications include data governance, reproducibility repositories, clear responsibility for automated actions, access equity and scrutiny of computational resource accounting. Ethical safeguards include consent or lawful authorization for restricted datasets, prohibition on unattended hazardous experimentation, independent dispute routes for symbolic exclusions, and monitoring for distribution shift. Reporting should separate generated hypotheses, logically admissible hypotheses and independently confirmed findings.

5. Conclusion

This study specifies NSQI-DL as a mathematically articulated proposal for integrating neural inference, symbolic admissibility, quantum-aware computation and Bayesian experiment design within an auditable scientific workflow. Its distinctive design choice is conditional integration: a quantum component is deployed only when justified by measured utility, a symbolic gate can reject a promising but inadmissible candidate, and posterior belief changes only after authenticated observations. Analytical calculations show how selection responds to feasibility constraints and trade-off weights, without representing scenario values as measured research outputs. The next scientific contribution must be a preregistered, resource-matched empirical assessment across independent laboratories and research domains. Until such evidence exists, NSQI-DL should be understood as a falsifiable research architecture rather than a validated autonomous discovery engine.

5.1 Limitations

The principal limitation is the absence of a collected study dataset or prospective deployment: there are no empirical performance estimates or established effect sizes. The mathematical model simplifies the interaction between neural confidence and evidence reliability; probability calibration may fail under shift. Ontology incompleteness and incorrectly asserted hard rules may produce false exclusions. Expected information gain requires a credible likelihood model and can be sensitive to misspecification. Quantum value depends on task structure and hardware benchmarks and may be negative. Domain-specific safety rules, availability of independent replications and inconsistent experimental cost measures limit generalizability.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date :Oct 30 2026 (05:00 PM)

5.2 Future Research Directions

- Develop a multi-domain benchmark with prospective, blinded laboratory replication and time-stamped evidence splits.
- Establish proof-carrying neural hypotheses with machine-checkable assumptions and expert appeal mechanisms.
- Measure quantum-versus-classical utility at matched solution quality, cost, latency and energy under real noise.
- Extend Bayesian acquisition to distributionally robust and multi-objective experimental design.
- Study evolving ontologies and formally verified updates when scientific concepts or rules change.
- Assess provenance-weighted posterior calibration in the presence of contradictory and incomplete sources.
- Investigate cross-laboratory transfer, interoperable metadata and third-party reproducibility of agent actions.
- Examine governance of autonomous experiment proposals, access controls and accountable human authorization.

References

Confalonieri, R. and Guizzardi, G. (2025) 'On the multiple roles of ontologies in explanations for neuro-symbolic AI', Neurosymbolic Artificial Intelligence, 1. doi:10.3233/NAI-240754.

Ferguson, M., Liu, F., Lin, H., McManus, D.D. and Gerber, B.S. (2026) 'Introduction to using symbolic regression for interpretable machine learning in healthcare', Communications Medicine, 6, article 524. doi:10.1038/s43856-026-01914-x.

Ghafarollahi, A. and Buehler, M.J. (2024) 'SciAgents: Automating scientific discovery through bioinspired multi-agent intelligent graph reasoning', Advanced Materials, 37(22). doi:10.1002/adma.202413523.

Ghareeb, A.E., Chang, B., Mitchener, L., Yiu, A., Szostkiewicz, C.J., Shved, D., Gyimesi, G.J., Laurent, J.M., Wright, S.M., Razzak, M.T., White, A.D., Finnemann, S.C., Hinks, M.M. and Rodriques, S.G. (2026) 'A multi-agent system for automating scientific discovery', Nature, 655(8122), pp. 497–505. doi:10.1038/s41586-026-10652-y.

Ji-An, L., Benna, M.K. and Mattar, M.G. (2025) 'Discovering cognitive strategies with tiny recurrent neural networks', Nature, 644(8078), pp. 993–1001. doi:10.1038/s41586-025-09142-4.

Murtas, G., Boeva, V. and Tsiporkova, E. (2026) 'Neuro-LENS: A neuro-symbolic framework integrating incomplete background knowledge and deep learning', Neurosymbolic Artificial Intelligence, 2. doi:10.1177/29498732261469489.

Núñez-Molina, C., Mesejo, P. and Fernández-Olivares, J. (2026) 'Automated planning instance generation with neuro-symbolic AI', Artificial Intelligence, 352, article 104471. doi:10.1016/j.artint.2025.104471.

Pasupuleti, M.K. (2026) 'Explainable quantum-logical multi-agent systems for autonomous scientific discovery in high-dimensional research landscapes', International Journal of Academic and Industrial Research Innovations (IJAIRI), 6(4), pp. 63–87. doi:10.62311/nesx/rp3-30042026.

Petruzzellis, F., Testolin, A. and Sperduti, A. (2026) 'Learning neuro-symbolic convergent term rewriting systems', Artificial Intelligence, 360, article 104622. doi:10.1016/j.artint.2026.104622.

eSign
Signed by MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date : Oct 30 , 2026 (05:00 PM)

Wang, Z., Rajabzadeh, T., Lee, N. and Safavi-Naeini, A.H. (2022) 'Automated discovery of autonomous quantum error correction schemes', PRX Quantum, 3(2). doi:10.1103/PRXQuantum.3.020302.

eSign
Signed by: MURALI KRISHNA PASUPULETI
Reason: NES |IJAIRI | Publication|
Copyright: © 2026
Location: Delhi, India
Date :Oct 30, 2026 (05:00 PM)

