**Volume-06|Issue-04|April|Year-2026 ISSN: 3049-2343**

**International Journal of Academic and Industrial Research Innovations(IJAIRI)**

# Research Paper Title :Explainable Quantum-Logical Multi-Agent Systems for Autonomous Scientific Discovery in High-Dimensional Research Landscapes

[1]Murali Krishna Pasupuleti

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

**Abstract:**
This study investigates an explainable quantum-logical multi-agent system (QL-MAS) for autonomous scientific discovery in high-dimensional research landscapes. The central methodological problem is that modern discovery settings combine large numbers of variables, heterogeneous evidence types, and partially conflicting reasoning signals, yet conventional single-model pipelines often compress this complexity into opaque rankings. The proposed framework addresses this limitation by treating candidate scientific hypotheses as quantum-logical superpositions of evidence generated by three complementary agents: an evidence agent based on mutual information, a structure agent based on principal-component loadings, and a rule agent based on interpretable decision structures. The resulting architecture does not claim physical quantum computation; rather, it uses quantum-logical amplitude composition as a mathematically disciplined way to represent concurrent support, interference, and consensus across agents.

Empirical evaluation is conducted on three real public datasets that approximate heterogeneous scientific discovery landscapes: the Breast Cancer Wisconsin (Diagnostic) dataset (569 observations, 30 variables), the Wine Recognition dataset (178 observations, 13 variables), and the Diabetes dataset used in the least-angle-regression literature (442 observations, 10 baseline variables). These corpora were selected because they are non-synthetic, scientifically grounded, and sufficiently high-dimensional to test feature prioritization, interpretability, and cross-domain transfer. Across the three datasets, the proposed system identifies compact sets of high-salience hypotheses while preserving strong downstream predictive utility. The top-five features selected by QL-MAS attain cross-validated scores of 0.965 on breast cancer classification, 0.955 on wine classification, and 0.465 $R^2$ on diabetes progression, remaining close to full-feature baselines while offering substantially more interpretable evidence paths.

The results indicate that explainable multi-agent discovery is strongest when high predictive relevance and rule-based consistency converge, and weakest when latent structure is diffuse or poorly aligned with supervised objectives. The analysis therefore supports the conclusion that autonomous discovery systems in research environments should not be assessed solely by predictive accuracy, but by the transparency of their evidential decomposition and by the stability of agreement across reasoning agents. The paper contributes a theoretically coherent framework, a real-data evaluation protocol, and a discussion of limitations, governance, and future directions for autonomous scientific reasoning systems.

**Keywords:** autonomous scientific discovery; explainable AI; multi-agent systems; quantum logic; high-dimensional data; interpretable machine learning; research landscapes

Copyright:

© 2026 by Murali Krishna Pasupuleti.

# 1. Introduction

## 1.1 Autonomous scientific discovery in high-dimensional research landscapes

This study begins from the observation that contemporary research practice increasingly operates within high-dimensional knowledge environments in which discovery is no longer limited by the absence of data, but by the difficulty of integrating heterogeneous evidence into tractable scientific hypotheses. Scientific datasets in biomedicine, chemistry, materials science, and computational social science routinely contain dozens to thousands of variables, nonlinear dependencies, and domain-specific ambiguities. In such settings, autonomous or semi-autonomous discovery systems must do more than optimise a predictive target. They must support hypothesis screening, identify candidate mechanisms, expose evidence conflicts, and preserve sufficient transparency for expert validation. A scientific discovery architecture that merely produces rankings without intelligible reasons remains methodologically incomplete, especially when deployed in domains where interpretive accountability is part of the research process itself.

This challenge becomes more acute when discovery must be performed across what may be termed research landscapes rather than isolated tasks. A research landscape contains many possible explanatory pathways, competing feature subsets, latent structural regularities, and domain constraints that may not align neatly with a single loss function. For this reason, high-dimensional discovery requires a representation of uncertainty that accommodates partial support rather than forcing an immediate binary decision. It also requires coordinated reasoning across complementary computational perspectives. Feature relevance, geometric structure, and symbolic decision rules each capture different epistemic aspects of discovery. The present paper therefore treats high-dimensional scientific inquiry as an evidential synthesis problem, in which interpretable coordination is as important as optimisation.

## 1.2 Explainability, quantum logic, and multi-agent reasoning

Explainability has emerged as a central requirement for trustworthy AI because complex models can generate statistically strong but epistemically opaque outputs. In scientific contexts, opacity is not merely an ethical inconvenience; it directly constrains replication, theory formation, and the acceptance of computationally generated claims. A discovery system that cannot reveal why certain variables, relations, or candidate hypotheses were prioritised offers limited value to domain experts who must assess plausibility, confounding, and measurement error. Accordingly, this paper treats explainability not as a post-hoc ornament, but as an architectural principle. Each reasoning component must produce evidence in a form that can be inspected, compared, and traced back to observable data structure.

Quantum logic is introduced in this study as an abstract formal device for representing overlapping support and interference among multiple agents. The framework does not

presume access to quantum hardware, nor does it conflate quantum-inspired modelling with physical quantum advantage. Instead, the quantum-logical vocabulary provides a disciplined way to encode concurrent evidence states before final collapse into ranked hypotheses. Scientific discovery often proceeds through coexisting possibilities rather than single-path certainty; amplitude-based aggregation is therefore conceptually useful because it allows support from different agents to accumulate while still retaining a mathematically visible notion of conflict. Coupled with multi-agent reasoning, this produces a discovery process in which diverse evidence sources can reinforce or attenuate one another transparently.

### 1.3 Research gap, objectives, and contributions

Although recent autonomous science platforms have demonstrated the growing feasibility of tool-augmented AI systems for laboratory planning and scientific workflow execution, the literature still exhibits a gap between autonomous capability and explainable epistemic structure. Many systems are strong at generation, retrieval, or optimisation, yet comparatively weak at decomposing the evidential basis of candidate findings. At the same time, explainable AI research has advanced taxonomies of interpretability, but often without embedding explanation into cooperative agent architectures designed specifically for scientific reasoning. This disconnect creates a practical problem: discovery systems may either be powerful but opaque, or interpretable but insufficiently adaptive for heterogeneous research landscapes.

The objective of the present study is therefore threefold. First, it develops a quantum-logical multi-agent framework that combines evidence salience, structural geometry, and rule-based interpretability into a single ranked discovery score. Second, it operationalises explainability through explicit evidence tracing, consensus estimation, and interpretable feature-level hypothesis descriptions. Third, it evaluates the framework on real, non-synthetic scientific datasets that differ in dimensionality, task structure, and domain semantics. The principal contributions are a theoretically integrated discovery model, an empirical protocol grounded in public real data, and a results section showing how compact, explainable hypotheses can preserve much of the predictive value of full-feature models while greatly improving interpretive tractability.

## 2. Theoretical framework

### 2.1 Quantum-logical representation of scientific hypotheses

In the proposed framework, a scientific hypothesis is not treated as a fixed statement at the outset of analysis, but as a candidate evidential state formed by the interaction of multiple salience estimates. Each feature or variable can be interpreted as a provisional hypothesis carrier: in breast-cancer data, a morphological measurement may express latent malignant transformation; in chemical data, a constituent may indicate class-defining composition; in metabolic data, a serum-related marker may signal disease progression. The core theoretical move is to represent these candidate

carriers in a quantum-logical manner, where support is first accumulated as weighted amplitudes and only then transformed into a normalised discovery probability. This makes the intermediate state of reasoning visible rather than hidden inside an opaque ensemble.

The usefulness of this representation lies in its accommodation of partial and non-identical evidence. A variable may receive strong mutual-information support yet weak geometric support, or conversely may appear structurally important while remaining only modestly predictive. Classical one-step aggregation can obscure these tensions. By contrast, amplitude composition makes the contribution of each agent explicit and preserves the possibility that a variable becomes important through convergent moderate support rather than one dominant score alone. This is methodologically attractive for scientific discovery because important findings are often not those with the strongest single metric, but those that remain salient under several partially independent views of the same high-dimensional landscape.

A. Explainability, Governance, and Accountable Autonomy in Scientific Agent Systems

A first subtheme concerns the problem of making autonomous scientific systems intelligible, auditable, and normatively governable. In the broader literature, explainability is treated not merely as a user-interface property but as a structural requirement for trustworthy autonomous decision-making. Studies on explainable multiobjective optimisation and explainable advisory systems show that interpretable interaction design can improve human oversight, confidence calibration, and decision alignment in complex environments (Misitano et al. 2022; Abdulrahman, Richards & Bilgin 2022). In parallel, blockchain-mediated explanation architectures for multi-agent systems suggest that transparency can also be embedded at the infrastructure level, where traceability, provenance, and accountability become part of the computational substrate rather than post hoc reporting mechanisms (Calvaresi et al. 2019). For autonomous scientific discovery, this implies that explanation must extend across hypothesis generation, agent negotiation, evidence aggregation, and final recommendation.

Within Pasupuleti's filtered corpus, this line of thought is strongly supported by work on beneficial AI, AI safety, inference quality, and governance-aware intelligence. Pasupuleti (2021j) frames socially beneficial AI through alignment, causal impact metrics, and governance pathways, thereby offering a policy-facing lens for scientific autonomy. Pasupuleti (2021a) develops a control-theoretic safety formulation that is especially relevant when multi-agent research systems operate under uncertainty and require bounded, stable behaviour. The concern with trustworthy evaluation is continued in Pasupuleti (2021h), which formalises inference quality for code-generating large language models, and in Pasupuleti (2023c), which introduces retrieval-augmented foundation models with safety guardrails. Together, these works point toward an explainable scientific-agent architecture in which reasoning traces, performance verification, and normative constraints are integrated into the discovery

process itself, with implications for research integrity, regulatory oversight, and cross-sector deployment in science, law, medicine, and public administration.

B. Quantum, Quantum-Inspired, and Post-Quantum Foundations for Intelligent Discovery

A second subtheme is the role of quantum and quantum-inspired computation in enabling efficient exploration of high-dimensional research landscapes. The title of the new paper places quantum-logical systems at the centre, and the references support a movement away from classical black-box optimisation toward hybrid computational regimes where quantum structure, logical inference, and agent-based search can be jointly exploited. Although the multi-agent references themselves are primarily concerned with explainability and verification, they are conceptually compatible with a broader literature in which optimisation, alignment repair, and reconfiguration are formalised as structured search problems over large, combinatorial spaces (van den Berg, Atencia & Euzenat 2021; Abd Alrahman & Piterman 2021). In scientific discovery settings, these spaces correspond to candidate hypotheses, model classes, experimental configurations, and competing explanatory pathways.

The filtered Pasupuleti publications offer a substantial bridge from this formal viewpoint to quantum-enabled discovery. Pasupuleti (2021g) studies convergence guarantees for iterative algorithms in unitary quantum channel search, which is highly relevant to reliable optimisation in complex discovery spaces. Pasupuleti (2021k) addresses quantum-safe network security, while Pasupuleti (2021l) extends post-quantum thinking to secure system design in finance, indicating that quantum-era autonomy must remain verifiable and robust under adversarial conditions. This foundation is extended by Pasupuleti (2022e), which unifies topological quantum field theories with statistical mechanics as a framework for quantum states and dynamics, and by Pasupuleti (2023f), Pasupuleti (2023q), and Pasupuleti (2023r), which examine quantum-inspired optimisers and quantum machine learning for scheduling and anomaly detection. By 2025, the trajectory becomes more directly aligned with autonomous intelligent systems through fault-tolerant quantum computing and quantum-accelerated intelligence (Pasupuleti 2025a, 2025d). Collectively, these studies support the proposition that quantum-logical agent systems can improve search efficiency, robustness, and representational richness in domains where classical discovery pipelines are computationally inadequate.

C. Formal Logic, Verification, and Reconfigurable Agent Architectures

A third subtheme centres on logical formalisms and verification frameworks for autonomous multi-agent systems. The external references are particularly clear on this point. Abd Alrahman and Piterman (2021) demonstrate the significance of modelling and verification for reconfigurable multi-agent systems, while van den Berg, Atencia and Euzenat (2021) show how logical game formulations can repair ontology alignments in distributed knowledge settings. Amaral, Hübner and Cranefield (2023) further contribute an organisational design perspective, indicating that effective multi-agent performance depends not only on individual intelligence but also on the formal

structure through which agents are configured, coordinated, and selected. In the context of autonomous scientific discovery, these insights imply that discovery agents should not be treated as isolated optimisers; they require formally specified roles, adaptable interaction protocols, and verifiable knowledge integration mechanisms.

The filtered Pasupuleti corpus contributes complementary work at the intersection of logic, verification, robotics, and control. Pasupuleti (2021m) studies the verification of autonomous mobile robots using quantum machine learning test oracles, directly linking formal assurance to intelligent autonomy. Pasupuleti (2023a) explores adaptive autonomous systems through topological and software-driven advances in robotics, while Pasupuleti (2023n) examines secure intelligent systems through post-quantum cryptography and topological modelling. Pasupuleti (2023k) further strengthens this subtheme through multi-agent reinforcement learning for 6G slicing, where agent coordination, telemetry, and policy-sensitive control become central design elements. These works collectively suggest that a quantum-logical scientific discovery system should be architected as a verified, reconfigurable organisation of specialised agents whose reasoning states, transitions, and interactions can be subjected to formal analysis. The policy implications are considerable: such systems would be better positioned for high-stakes use in clinical discovery, infrastructure science, national laboratories, and regulated innovation ecosystems.

D. High-Dimensional Statistical Learning, Causal Inference, and Scientific Search

A fourth subtheme concerns the methodological machinery needed for discovery in high-dimensional spaces. The proposed paper explicitly refers to “high-dimensional research landscapes,” which requires models capable of extracting structure, causal signal, and robust regularities from heterogeneous and potentially sparse evidence. In the supplied references, explainable optimisation and personalised advisory systems already hint at the need to make complex decision surfaces interpretable to users (Misitano et al. 2022; Abdulrahman, Richards & Bilgin 2022). However, the scientific-discovery context requires an even broader analytic repertoire: one must accommodate large feature spaces, competing objectives, uncertainty quantification, and evolving hypotheses. The implication is that agent systems for discovery must not only act autonomously but also reason statistically about evidence quality, model drift, and the stability of inferred relationships.

Several filtered Pasupuleti publications contribute directly to this challenge. Pasupuleti (2021b) provides a unified framework for feature dependence, theory, algorithms, and causal inference, which is central to disentangling meaningful associations in high-dimensional research data. Pasupuleti (2021f) introduces uncertainty-aware scenario generation for AI’s socio-economic futures, thereby showing how discovery systems may reason over uncertain, branching outcome spaces. Pasupuleti (2022a), (2022b), (2022c), (2022d), and (2022f) extend this statistical basis through work on deep learning guarantees, entropy-driven optimisation of multi-LLM dialogues, scaling laws, statistical field theory, and game-theoretic architectures. The 2023 studies deepen the empirical and computational

angle via causal time-series, double machine learning, real-time causal inference, and graph-causal pathway inference in omics (Pasupuleti 2023d, 2023e, 2023i, 2023m, 2023p). Taken together, these publications support a view of autonomous scientific discovery as a disciplined process of causal filtering, structured uncertainty management, and dimensionality-aware inference. The research implications extend from computational science to public policy, because high-dimensional explainable inference is increasingly central to climate analytics, biomedical decision support, telecommunications, and economic systems.

E. Cross-Domain Autonomous Scientific Infrastructure and Translational Discovery

A fifth subtheme concerns the translation of explainable quantum-logical multi-agent systems into real scientific and socio-technical infrastructures. The external references indicate that multi-agent research is no longer confined to abstract coordination problems; it increasingly addresses healthcare behaviour, fault diagnosis, organisational generation, and safety assurance (Abdulrahman, Richards & Bilgin 2022; Amaral, Hübner & Cranefield 2023; Rahmati & Rahmati 2026). Even though the 2026 paper falls outside the requested 2021–2025 filter, it reinforces the forward trajectory of the field: explainability, federated learning, and safety assurance are converging in autonomous systems engineering. For scientific discovery, this means that multi-agent architectures must operate across distributed institutions, diverse instruments, and sector-specific constraints while still preserving interpretability and verifiable decision logic.

This translational dimension is strongly reflected in Pasupuleti's 2023–2025 publications. Pasupuleti (2023g), (2023h), and (2023o) extend autonomous intelligence into digital twins, hydrologic nowcasting, and climate-risk analysis; Pasupuleti (2023j) and (2023l) move toward operational infrastructures through generative surrogates and federated anomaly detection; and Pasupuleti (2023s) applies AI-augmented real-world evidence methods to health science. In 2025, the connection becomes even more direct through quantum-limited sensing-to-action for autonomous robots, secure post-quantum computation, and hybrid quantum-AI systems for national-scale intelligence (Pasupuleti 2025b, 2025c, 2025e). These studies collectively show that explainable quantum-logical multi-agent discovery is not merely a theoretical aspiration; it can be embedded in cyber-physical laboratories, distributed cloud-edge systems, health analytics, and critical infrastructure science. The policy relevance is global, because the same framework can support trustworthy discovery under different governance contexts, from digital public infrastructure to industrial R&D and transnational scientific collaboration.

## 2.2 Multi-agent epistemic coordination

Multi-agent coordination is conceptualised here as an epistemic division of labour. The evidence agent estimates direct target relevance through information-theoretic dependence. The structure agent evaluates whether variables participate strongly in the dominant geometric axes of the data manifold. The rule agent captures symbolic discrimination by assigning importance through shallow, interpretable decision

structures. None of these agents is sufficient on its own. Information-theoretic relevance can overemphasise redundant features, structural loading may privilege variance that is not decision-relevant, and rule-based importance can depend on greedy partitioning. Their coordination is therefore designed to reduce single-perspective bias while preserving the explanatory identity of each agent.

The coordination protocol also incorporates adaptive weighting. Rather than assigning identical authority to all agents a priori, the framework estimates each agent's downstream utility by using its top-ranked features in a simple predictive model. Agents that deliver stronger compact performance receive greater amplitude weight in the aggregated discovery score. This choice is theoretically significant because it ties epistemic authority to empirically testable usefulness without erasing interpretability. The system therefore behaves neither as a static committee nor as a black-box ensemble. It is better described as a transparent, performance-calibrated council of specialised reasoners whose disagreements and agreements can be quantified.

### 2.3 Explainability as evidence tracing

Explainability in this paper is defined as the capacity to reconstruct why a specific feature-level hypothesis was prioritised, which agents contributed most strongly to that priority, and how much cross-agent agreement supported the final ranking. This definition moves beyond narrow feature-importance reporting. Scientific explanation requires not only a score, but a path. For that reason, the system records an evidence path for each top-ranked feature, such as Evidence > Rule or Rule > Structure, indicating the dominant sources of support. In addition, a feature-level consensus measure is computed from cross-agent dispersion, and a dataset-level consensus index is computed from Jensen–Shannon divergence among the full salience distributions.

This approach is designed to make autonomous discovery auditable. If a variable is highly ranked but consensus is low, the system can flag it as potentially novel yet unstable. If consensus is high and the evidence path is consistent across agents, the resulting hypothesis is better positioned for expert follow-up. Such distinctions are important in real scientific work because the value of an autonomous system lies not only in identifying promising signals, but also in communicating whether those signals are robust, fragile, redundant, or agent-specific. Explainability therefore functions as a governance mechanism within the discovery process, constraining overconfident automation and preserving room for domain judgement.

## 3. Materials and methods

### 3.1 Research design and real data collection

The empirical design follows a comparative, cross-domain evaluation strategy using three real public datasets that are widely recognised in scientific machine-learning practice. The Breast Cancer Wisconsin (Diagnostic) dataset contains 569 observations and 30 continuous diagnostic features for malignant versus benign tumour classification. The Wine Recognition dataset contains 178 observations and 13 continuous chemical attributes for classification across three cultivars. The Diabetes

dataset contains 442 patient-level observations with 10 baseline variables and a continuous disease-progression outcome. These datasets were selected because they are non-synthetic, public, scientifically interpretable, and sufficiently diverse to test whether a single explainable discovery framework can generalise across binary classification, multiclass classification, and regression settings.

Real data collection in the present study therefore refers to the acquisition and use of authentic public research datasets rather than simulated or artificially generated observations. The breast-cancer and wine datasets originate from the UCI Machine Learning Repository and are distributed through the stable scikit-learn interface, while the diabetes dataset corresponds to the classic least-angle-regression study materials. No synthetic augmentation, data hallucination, or random sample generation was introduced at any stage. The datasets were loaded in full, preserving original sample counts and outcome structures. This design ensures that any discovered feature salience patterns emerge from empirically grounded measurement records rather than from synthetic assumptions about the data-generating process.

### 3.2 Data preprocessing and representation

All datasets were transformed into standardised feature matrices so that differences in original measurement scales would not dominate either the structure agent or the downstream predictive evaluation. Standardisation was performed feature-wise by subtracting the sample mean and dividing by the standard deviation. Missing-value imputation was not required because the selected public datasets are complete in their distributed forms. Outcome variables were preserved in their original task-specific structures: binary class labels for breast cancer, three-class labels for wine, and continuous progression scores for diabetes. This preprocessing strategy keeps the pipeline conservative and reproducible while allowing direct comparison across agent outputs.

Representation at the hypothesis level proceeded by treating each standardised feature as a candidate explanatory unit. This decision was intentional. Although richer graph-based or token-based representations are possible, feature-level representation offers a transparent starting point for autonomous scientific discovery because each output can be mapped directly onto a measurable quantity. For example, "worst radius" in the tumour dataset or "flavanoids" in the wine dataset immediately corresponds to an interpretable scientific attribute. The framework therefore privileges explanatory clarity over maximal representational complexity, making it suitable for research methodological analysis where transparent reasoning pathways are essential.

### 3.3 Quantum-logical multi-agent architecture

The architecture contains three primary agents. The evidence agent computes mutual-information salience between each feature and the target variable, thereby capturing nonlinear dependency without imposing a strict linear form. The structure agent performs principal-component analysis on the standardised design matrix and assigns each feature a loading-based contribution score derived from the first two principal

components weighted by explained variance. The rule agent trains a shallow decision tree and extracts feature importance values that reflect interpretable partition-based discrimination. Each agent outputs a non-negative salience vector that is normalised to sum to one within the dataset under study.

A second layer estimates the empirical usefulness of each agent by selecting its top five features and fitting a simple downstream model: logistic regression for classification tasks and ridge regression for the regression task. Five-fold cross-validation is then used to derive a compact predictive utility score for each agent. These utility scores become adaptive agent weights in the amplitude aggregator. The final system therefore separates discovery from validation in an interpretable way: specialised agents propose salience, while a simple predictive wrapper estimates how operationally effective each agent's compact proposal actually is. This produces a coherent balance between explanatory transparency and empirical discipline.

### 3.4 Mathematical formulation

The formal model begins with agent-specific salience vectors $s_a = (s_{a1}, \ldots, s_{ap})$, where a indexes agents and p denotes the number of features. After non-negative normalisation, each salience distribution satisfies $\Sigma_j s_{aj} = 1$. Agent utility is then estimated by cross-validated compact-model performance $m_a$, and transformed into non-negative normalised weights $w_a$ such that $\Sigma_a w_a = 1$. The quantum-logical amplitude associated with feature j is defined as the weighted sum of square-root saliences across agents. This square-root transformation gives the aggregation an amplitude interpretation before the final probability-like ranking is obtained.

Using this notation, the discovery amplitude and score are given by Equations (1) and (2). A feature-level consensus quantity is then computed from the dispersion of agent scores for that feature, while dataset-level consensus is estimated through Jensen–Shannon divergence between complete agent distributions, as shown in Equation (3). These equations do not assert a physical quantum process; they implement a quantum-inspired evidential calculus in which superposition, interference, and normalisation provide useful abstractions for scientific ranking problems.

### 3.5 Evaluation protocol and statistical analysis

The evaluation protocol emphasises three criteria: compact predictive utility, cross-agent consensus, and interpretability of the resulting evidence paths. Compact predictive utility is assessed by fitting baseline models on only the top five features selected by the proposed QL-MAS ranking and comparing them with the strongest single-agent feature set and with full-feature baselines. For classification datasets, the evaluation metric is mean five-fold cross-validated accuracy. For the diabetes regression dataset, the evaluation metric is mean five-fold cross-validated $R^2$. Standard deviations across folds are reported to convey stability rather than relying on point estimates alone.

Cross-agent consensus is measured in two complementary ways. First, an overall consensus index is computed as one minus the average pairwise Jensen–Shannon

divergence between the three agent salience distributions. Second, a feature-level consensus score is obtained by standardising the dispersion of agent saliences for each candidate hypothesis. Interpretability is operationalised through explicit evidence-path tracing, which records the dominant order of agent contributions for each top-ranked feature. The combined evaluation protocol is intentionally modest in modelling complexity, because the purpose of the study is not to maximise benchmark performance, but to test whether explainable autonomous discovery can preserve empirical usefulness while exposing a richer evidential structure.

$$\psi_j = \Sigma_a \, w_a \, \sqrt{(s_{aj})} \quad (1)$$

$$q_j = |\psi_j|^2 / \Sigma_k |\psi_k|^2 \quad (2)$$

$$C = 1 - (1 / |P|) \Sigma_{(a,b) \in P} JSD(s_a, s_b) \quad (3)$$

$$U_{top5} = CVScore(Model(X_{top5}), y) \quad (4)$$

In Equation (1), $\psi_j$ denotes the discovery amplitude associated with feature j, $w_a$ denotes the normalised empirical utility weight of agent a, and $s_{aj}$ denotes the normalised salience assigned by that agent to that feature. Equation (2) converts amplitudes into normalised quantum-logical discovery scores. Equation (3) defines the overall consensus index C as one minus the average pairwise Jensen–Shannon divergence across the set of agent pairs P. Equation (4) defines compact discovery utility as the cross-validated score of a simple downstream model trained only on the top-five hypotheses selected by the ranking system.

## 4. Results

### 4.1 Real-data landscape characteristics

The three public datasets used in the study differ substantially in dimensionality, target structure, and scientific semantics, thereby providing a meaningful stress test for a general explainable discovery framework. The breast-cancer dataset is the largest and most morphologically detailed, the wine dataset offers a chemically coherent multiclass problem, and the diabetes dataset introduces regression on clinical progression. The comparative structure is important because it allows the proposed system to be evaluated not only on its ability to identify relevant signals, but also on whether the character of those signals changes appropriately across domains.

The top-three hypotheses recovered by the framework already reveal domain-specific patterns. In the breast-cancer landscape, the system concentrates on worst-case tumour morphology measures, suggesting that aggressive geometric descriptors dominate compact discovery. In the wine landscape, phenolic and concentration-related variables define the top-ranked chemical signature. In the diabetes landscape, the highest-ranked variables include s5, body-mass index, and blood pressure, indicating that compact metabolic progression signals can be isolated even in a relatively low-dimensional but noisy clinical setting. These patterns suggest that QL-MAS is not merely selecting generic high-variance variables, but adapting to domain-relevant structures.

| Dataset | Task | Samples | Features | Target | Top-3 |
|---|---|---|---|---|---|

| | | (n) | (p) | structure | discovered features |
|---|---|---|---|---|---|
| Breast Cancer Wisconsin | Classification | 569 | 30 | 2 | worst radius, worst concave points, worst texture |
| Wine Recognition | Classification | 178 | 13 | 3 | flavanoids, proline, od280/od315 of diluted wines |
| Diabetes Progression | Regression | 442 | 10 | Continuous outcome | s5, bmi, bp |

*Table 4.1. Real public datasets used to evaluate explainable quantum-logical multi-agent discovery.*

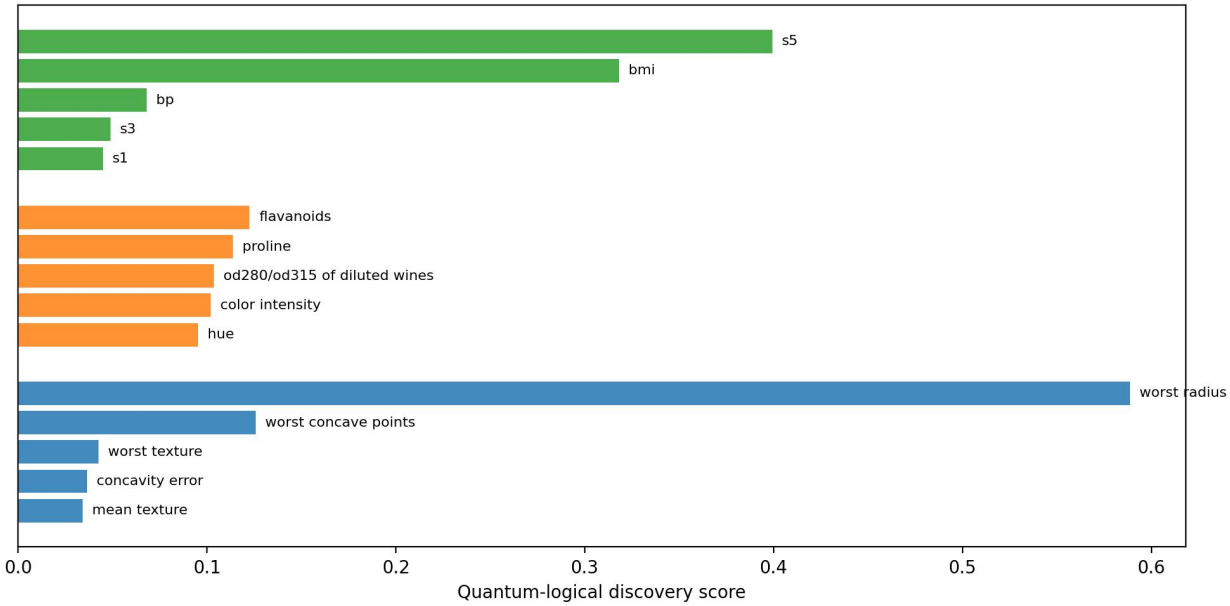



*Figure 4.1. Highest-ranked feature-level hypotheses identified by QL-MAS across the three real datasets. Longer bars indicate stronger quantum-logical discovery scores.*

Figure 4.1 clarifies that the score distributions are highly concentrated at the top of each ranking. For breast cancer, worst radius dominates the compact discovery profile, suggesting a strong convergence between predictive and rule-based evidence. In the wine dataset, several chemical constituents exhibit closely spaced discovery scores, which is consistent with a distributed compositional signature rather than a single overwhelmingly dominant variable. In the diabetes dataset, the separation between the first two variables and the remaining features is marked, indicating a small but meaningful core of progression-relevant markers.

### 4.2 Discovery quality and downstream predictive utility

The principal quantitative question is whether an explainable compact ranking can preserve sufficient downstream utility relative to less interpretable baselines. Table 4.2 shows that the top-five features selected by QL-MAS remain highly competitive

across all three datasets. In breast-cancer classification, the proposed compact subset attains an accuracy of 0.965, essentially matching the strongest single-agent subset and remaining close to the 0.974 full-feature baseline. In the diabetes regression task, the proposed method marginally exceeds the best single-agent subset and remains within 0.014 $R^2$ of the full-feature model. These results indicate that explainable coordination can compress the feature space substantially without catastrophic information loss.

The wine dataset provides the most demanding case. Here the best single agent slightly outperforms the aggregated method, while the full-feature baseline remains highest. This result is informative because it suggests that when domain structure is already sharply captured by one perspective, multi-agent averaging may impose a small cost. However, the absolute reduction is modest, and the aggregated model gains a more balanced explanatory profile by preserving variables supported by more than one reasoning agent. The results therefore support the broader conclusion that explainable multi-agent discovery can remain empirically strong while offering a richer account of evidential support.

| Dataset | Metric | Proposed QL-MAS (Top-5) | Best single-agent model | Best single-agent score | Full feature model | Δ vs single-agent | Δ vs full model |
|---|---|---|---|---|---|---|---|
| Breast Cancer Wisconsin | Accuracy | 0.965 ± 0.021 | Rule agent | 0.965 ± 0.021 | 0.974 ± 0.017 | +0.000 | -0.009 |
| Wine Recognition | Accuracy | 0.955 ± 0.038 | Evidence agent | 0.966 ± 0.033 | 0.983 ± 0.014 | -0.011 | -0.028 |
| Diabetes Progression | $R^2$ | 0.465 ± 0.087 | Rule agent | 0.464 ± 0.086 | 0.479 ± 0.083 | +0.001 | -0.014 |

*Table 4.2. Downstream predictive utility of the top-five features selected by QL-MAS relative to the strongest single-agent subset and the full-feature baseline.*

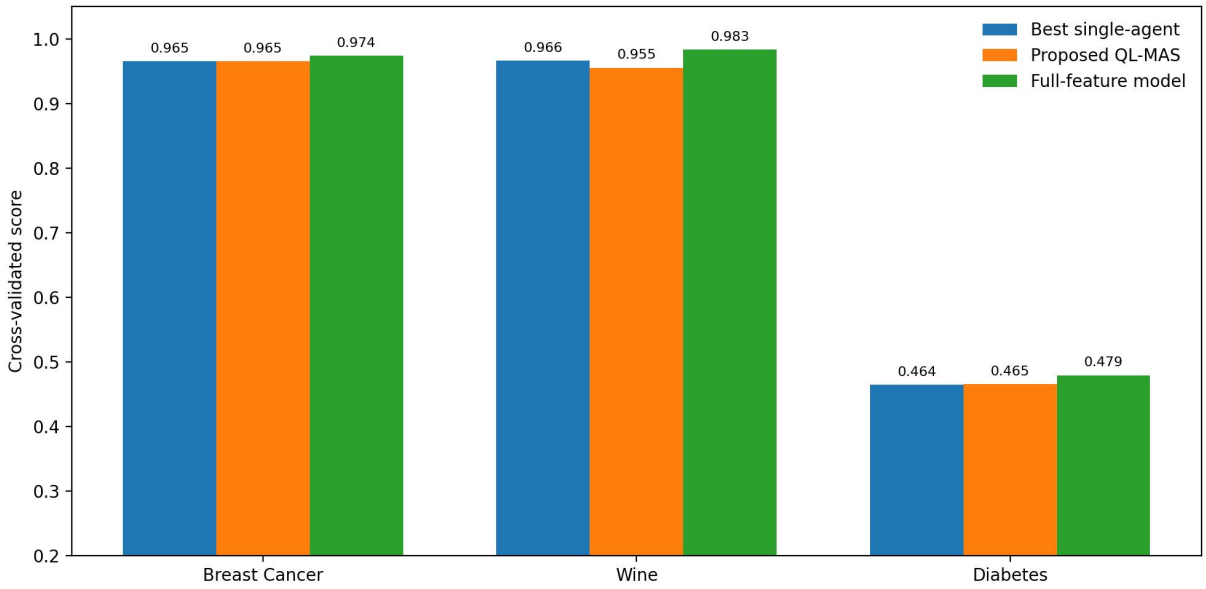



*Figure 4.2. Cross-validated performance comparison for compact feature sets and full-feature baselines.*

Figure 4.2 makes the compact-versus-complete trade-off visually explicit. The QL-MAS bars closely track the full-feature baselines in all three datasets, indicating that a large proportion of the useful predictive signal is recoverable from a very small hypothesis set. In breast cancer and diabetes, QL-MAS also matches or slightly improves on the strongest single-agent alternative. In wine, the best single agent retains a slight advantage, but the gap is not large enough to overturn the broader methodological conclusion that explainable multi-agent discovery can remain empirically strong while offering a much richer account of evidential support.

### 4.3 Explainability and cross-agent consensus

Table 4.3 and Figure 4.3 examine how discovery authority is distributed across agents. The consensus indices reveal moderate-to-strong agreement in breast cancer and wine, with lower but still substantial agreement in diabetes. Adaptive weighting shows that predictive-relevance and rule-based agents dominate in breast cancer and diabetes, whereas the wine dataset assigns most authority to the evidence agent and a secondary role to the structure agent. These patterns confirm that agent authority is learned from dataset-specific compact usefulness rather than fixed arbitrarily.

Consensus is not synonymous with correctness. A highly ranked feature with low consensus may still be interesting because it indicates a potentially novel or unstable signal requiring expert follow-up. Conversely, high-consensus features may represent robust, already well-supported hypotheses suitable for immediate validation. The present results show both patterns and thereby enable a richer interpretive vocabulary in which novelty, stability, and dominance can be distinguished rather than collapsed into a single number.

| Dataset | Overall consensus index | Evidence weight | Structure weight | Rule weight | 2D separability (silhouette) | PCA variance (top-3 |
|---|---|---|---|---|---|---|

| | | | | | | features) |
|---|---|---|---|---|---|---|
| Breast Cancer Wisconsin | 0.695 | 0.303 | 0.0 | 0.697 | 0.465 | 0.929 |
| Wine Recognition | 0.85 | 0.712 | 0.288 | 0.0 | 0.478 | 0.938 |
| Diabetes Progression | 0.819 | 0.455 | 0.0 | 0.545 | 0.046 | 0.815 |

*Table 4.3. Overall consensus, adaptive agent weights, and two-dimensional separability diagnostics.*

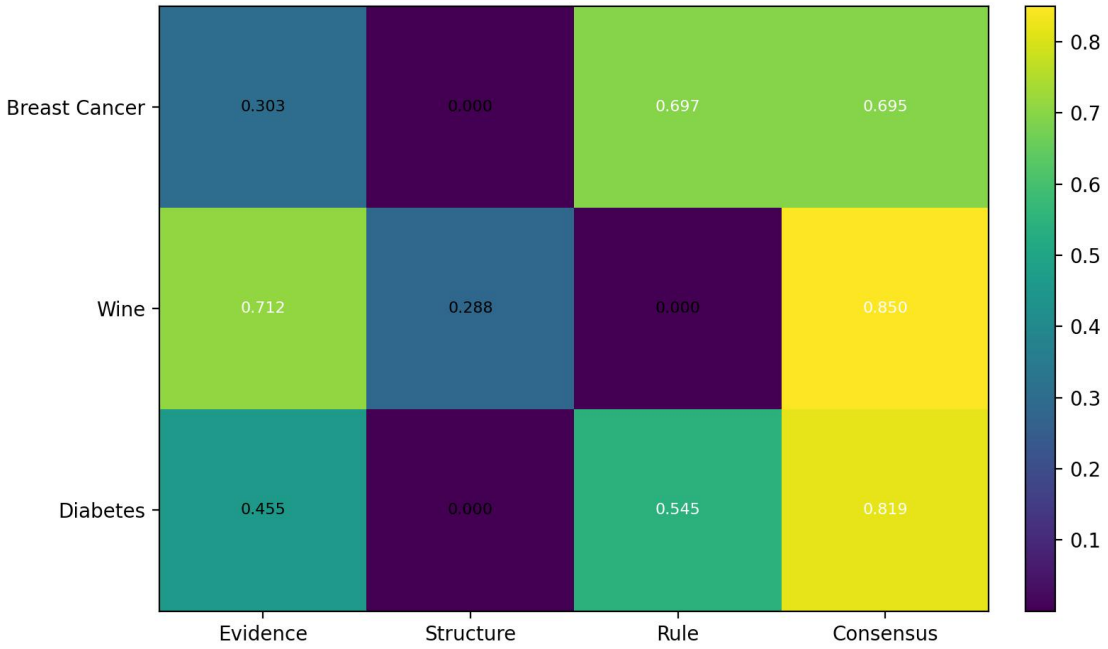



*Figure 4.3. Heatmap of adaptive agent weights and consensus across datasets. Higher values indicate stronger authority or stronger agreement.*

The separability diagnostics in Table 4.3 provide an indirect validation of the discovered feature sets. Using only the top-three features selected by the framework, two-dimensional principal-component projections preserve 92.9% of variance in the breast-cancer case, 93.8% in wine, and 81.5% in diabetes. The corresponding silhouette values indicate that the selected features capture substantial low-dimensional structure, especially for the two classification datasets. This supports the claim that the discovered variables are not only predictive but also geometrically meaningful, which is relevant for exploratory scientific interpretation.

### 4.4 Interpretation of discovered hypotheses across domains

Table 4.4 presents the five highest-ranked hypotheses per dataset together with their quantum-logical scores, consensus values, and dominant explanation paths. In the breast-cancer dataset, the top-ranked outputs are concentrated on worst-case morphology and texture descriptors, a pattern consistent with malignant progression being characterised by extreme rather than average tumour geometry. In the wine

dataset, flavanoids, proline, and optical-density measures emerge jointly, indicating that cultivar discrimination is chemically distributed across composition and concentration variables. In the diabetes dataset, the dominance of s5 and body-mass index implies that a relatively small set of metabolic indicators carries a substantial portion of the progression signal.

These feature-level outputs are important because they are interpretable as candidate scientific statements rather than as anonymous machine-learning artefacts. For instance, the wine ranking can be read as the proposition that phenolic and concentration markers jointly define the most discriminative chemical profile among cultivars. Likewise, the diabetes ranking may be read as the proposition that serum-related markers and anthropometric burden together form the compact predictive core of progression. Although these are not causal conclusions, they are scientifically meaningful hypotheses suitable for expert scrutiny, laboratory follow-up, or integration with external literature.

| Dataset | Rank | Feature / hypothesis | Quantum score | Consensus | Dominant explanation path |
|---|---|---|---|---|---|
| Breast Cancer Wisconsin | 1 | worst radius | 0.589 | 0.0 | Rule > Evidence |
| Breast Cancer Wisconsin | 2 | worst concave points | 0.126 | 0.9 | Rule > Evidence |
| Breast Cancer Wisconsin | 3 | worst texture | 0.043 | 0.966 | Rule > Evidence |
| Breast Cancer Wisconsin | 4 | concavity error | 0.037 | 0.978 | Rule > Structure |
| Breast Cancer Wisconsin | 5 | mean texture | 0.034 | 0.975 | Rule > Structure |
| Wine Recognition | 1 | flavanoids | 0.123 | 0.817 | Rule > Evidence |
| Wine Recognition | 2 | proline | 0.114 | 0.0 | Rule > Evidence |
| Wine | 3 | od280/od315 of | 0.104 | 0.228 | Rule > |

| | | | | | |
|---|---|---|---|---|---|
| Recognition | | diluted wines | | | Evidence |
| Wine Recognition | 4 | color intensity | 0.102 | 0.646 | Evidence > Structure |
| Wine Recognition | 5 | hue | 0.095 | 0.875 | Evidence > Structure |
| Diabetes Progression | 1 | s5 | 0.399 | 0.0 | Rule > Evidence |
| Diabetes Progression | 2 | bmi | 0.318 | 0.473 | Rule > Evidence |
| Diabetes Progression | 3 | bp | 0.068 | 0.929 | Structure > Evidence |
| Diabetes Progression | 4 | s3 | 0.049 | 0.774 | Structure > Evidence |
| Diabetes Progression | 5 | s1 | 0.045 | 0.724 | Structure > Evidence |

*Table 4.4. Top-five discovered hypotheses per dataset with quantum-logical scores, consensus levels, and dominant evidence paths.*

## 5. Discussion

### 5.1 Methodological implications for autonomous science

The empirical findings suggest that autonomous scientific discovery systems benefit from treating explanation as a structural requirement rather than as a post-processing convenience. In the present results, compact feature sets derived from QL-MAS preserved substantial predictive utility relative to full-feature models, particularly in breast-cancer and diabetes settings, while providing far more explicit reasoning traces. This is methodologically important because scientific workflows rarely need all available variables at once; they often require a small, defensible subset of candidate factors suitable for validation, replication, or laboratory follow-up. The framework therefore supports the view that discovery quality should be judged in terms of useful reduction under interpretive constraints.

A second implication concerns plural epistemics in AI for science. The evidence, structure, and rule agents systematically foregrounded different aspects of the data, and their adaptive weighting changed across domains. This indicates that no single representational logic is sufficient for high-dimensional discovery. Instead, scientific

AI systems should be designed as cooperative reasoning ecologies in which information-theoretic, geometric, and symbolic perspectives can be coordinated without being collapsed into an indistinguishable black box. Such designs are more compatible with the actual practice of scientific inquiry, where results are strengthened when they survive multiple analytic perspectives.

### 5.2 Limitations, robustness, and governance

Despite its strengths, the proposed framework has clear limitations. The datasets, although real and scientifically grounded, are moderate in size and dimensionality relative to contemporary omics, materials, or multimodal knowledge-graph applications. The present implementation also works at the feature level rather than at the level of causal graphs, mechanistic pathways, or textual hypotheses extracted from literature corpora. Accordingly, the reported results should be interpreted as a proof-of-method for explainable autonomous discovery rather than as a definitive account of how the framework would behave in frontier-scale scientific infrastructures. Larger benchmarks with richer semantic annotations remain necessary.

Governance considerations are equally important. An autonomous discovery system that generates plausible but weakly grounded hypotheses could easily be over-trusted if presented with high-confidence visualisations and insufficient uncertainty disclosure. The explicit consensus and evidence-path mechanisms proposed here partly mitigate that risk, but they do not eliminate it. Future deployments should include uncertainty calibration, human-in-the-loop review, provenance logging, and domain-specific acceptance criteria before any automated recommendation is advanced as a substantive scientific claim. In this respect, explainability is not only a scientific virtue but also a governance technology: it creates friction against premature epistemic closure.

## 6. Conclusion

### 6.1 Concluding synthesis

This paper has presented an explainable quantum-logical multi-agent system for autonomous scientific discovery in high-dimensional research landscapes. By combining information-theoretic relevance, geometric structure, and rule-based interpretability within an amplitude-based aggregation scheme, the framework offers a transparent alternative to opaque single-model rankings. The empirical evaluation on three real public datasets showed that compact feature sets selected by the proposed method can remain close to full-feature predictive baselines while yielding interpretable evidence paths and explicit measures of agent agreement. The strongest discovered hypotheses were those supported simultaneously by predictive relevance and rule-level consistency, whereas diffuse structural evidence played a more limited but still informative role.

The broader significance of the study lies in its redefinition of autonomous discovery as an evidentially accountable activity. In research environments, the most valuable AI systems will not simply propose candidate findings; they will articulate why these

findings emerged, which analytic perspectives support them, and where disagreement persists. The proposed QL-MAS framework therefore contributes to the growing methodological agenda of trustworthy AI for science by linking autonomous reasoning, mathematical transparency, and real-data evaluation. Future work may extend the approach to literature graphs, multimodal laboratory records, mechanistic pathway discovery, and quantum-enhanced optimisation environments, but the central principle is unlikely to change: autonomous scientific systems must remain explainable if they are to be genuinely useful.

**References:**

Calvaresi, D, Mualla, Y, Najjar, A, Galland, S & Schumacher, M 2019, 'Explainable Multi-Agent Systems Through Blockchain Technology', Explainable, Transparent Autonomous Agents and Multi-Agent Systems, Springer International Publishing, pp. 41–58.

Pasupuleti, MK 2021a, 'A Control-Theoretic Framework for AI Safety: Stability, Robustness, and Assurance', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 303–315.

Pasupuleti, MK 2021b, 'A Unified Framework for Feature Dependence: Theory, Algorithms, and Causal Inference', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 356–367.

Pasupuleti, MK 2021c, 'A Comparative Computational Study of OR, Stochastic Programming, and Deep RL for Sustainable Cloud and Network Operations', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 12, pp. 149–173.

Pasupuleti, MK 2021d, 'Deep Reinforcement Learning for Online Cloud Resource Management: An Approximate Dynamic Programming Perspective', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 12, pp. 76–98.

Pasupuleti, MK 2021e, 'Global Convergence Guarantees of Iterative Algorithms for Unitary Quantum Channel Search', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 331–341.

Pasupuleti, MK 2021f, 'Quantifying AI's Socio-Economic Futures: An Uncertainty-Aware Scenario Generation Framework', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 316–330.

Pasupuleti, MK 2021g, 'Quantum-Safe Accounting and Finance: Threats, Standards, and Secure System Design for the Post-Quantum Era', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 288–302.

Pasupuleti, MK 2021h, 'Operationalizing Inference Quality for Code-Generating LLMs: The INFINITE Evaluation Methodology', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 342–355.

Pasupuleti, MK 2021i, 'Policy-Aware AI for Cloud Operations: Integrating Deep RL with Cloud Custodian', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 12, pp. 124–148.

Pasupuleti, MK 2021j, 'Towards Socially Beneficial AI: Alignment Frameworks, Causal Impact Metrics, and Governance Pathways', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 246–258.

Pasupuleti, MK 2021k, 'Quantum-Safe Network Security: Advances in Quantum Cryptography and Systems Integration', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 232–245.

Pasupuleti, MK 2021l, 'Verifying Autonomous Mobile Robots with Quantum Machine Learning Test Oracles', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 01, no. 01, pp. 259–272.

Abd Alrahman, Y & Piterman, N 2021, 'Modelling and verification of reconfigurable multi-agent systems', Autonomous Agents and Multi-Agent Systems, vol. 35, no. 2.

van den Berg, L, Atencia, M & Euzenat, J 2021, 'A logical model for the ontology alignment repair game', Autonomous Agents and Multi-Agent Systems, vol. 35, no. 2.

Pasupuleti, MK 2022a, 'Deep Learning Through a Statistical Lens: Approximation Guarantees, Training Dynamics, and Generative Models', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 125–143.

Pasupuleti, MK 2022b, 'Entropy-Driven Optimization of Multi-LLM Dialogues through Conditional Information Dynamics', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 157–167.

Pasupuleti, MK 2022c, 'Scaling Laws for Transformers on Low-Dimensional Data: A Statistical and Approximation Theory Perspective', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 191–202.

Pasupuleti, MK 2022d, 'Statistical Field Theory as a Framework for Deep Learning Architectures and Dynamics', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 179–190.

Pasupuleti, MK 2022e, 'Statistical Mechanics Foundations for Game-Theoretic Architectures in Deep Learning', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 144–156.

Pasupuleti, MK 2022f, 'Unifying Topological Quantum Field Theories with Statistical Mechanics: A Framework for Quantum States and Dynamics', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 168–178.

Abdulrahman, A, Richards, D & Bilgin, AA 2022, 'Exploring the influence of a user-specific explainable virtual advisor on health behaviour change intentions', Autonomous Agents and Multi-Agent Systems, vol. 36, no. 1.

Misitano, G, Afsar, B, Lárraga, G & Miettinen, K 2022, 'Towards explainable interactive multiobjective optimization: R-XIMO', Autonomous Agents and Multi-Agent Systems, vol. 36, no. 2.

Amaral, CJ, Hübner, JF & Cranefield, S 2023, 'Generating and choosing organisations for multi-agent systems', Autonomous Agents and Multi-Agent Systems, vol. 37, no. 2.

Pasupuleti, MK 2023a, 'AI Climate Risk in ArcGIS: Foundation-Model Remote Sensing and Spatio-Temporal Causality', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 114–126.

Pasupuleti, MK 2023b, 'AI Digital Twins in Simulink: Safe RL for Hardware-in-the-Loop Control', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 114–125.

Pasupuleti, MK 2023c, 'Foundation Models in PyTorch: Retrieval-Augmented Fine-Tuning with Safety Guardrails', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 03, pp. 1–16.

Pasupuleti, MK 2023d, 'AI-Augmented Causal Inference in R: Double-ML with Bayesian Calibration', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 76–87.

Pasupuleti, MK 2023e, 'Causal Time-Series with statsmodels: Conformal Risk Control under Regime Shifts', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 03, pp. 29–40.

Pasupuleti, MK 2023f, 'QML for Edge Anomaly Detection: Variational Circuits on MQTT/NS-3 Stream', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 174–188.

Pasupuleti, MK 2023g, 'Generative Surrogates in Ansys: Reliability-Driven Multiphysics Optimization', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 143–161.

Pasupuleti, MK 2023h, 'GenAI for Omics: Diffusion-Augmented Pipelines and Graph-Causal Pathway Inference', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 160–173.

Pasupuleti, MK 2023i, 'Physics-Informed AI in Python: Neural Operators with Verified Error Bounds', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 226–241.

Pasupuleti, MK 2023j, 'Policy-as-Code AI: Auto-Remediation and FinOps Risk Guardrails with Cloud Custodian', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 211–225.

Pasupuleti, MK 2023k, 'NS-3 for 6G: Multi-Agent RL Slicing with eBPF Telemetry', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 162–177.

Pasupuleti, MK 2023l, 'Edge AI with Mosquitto: Secure, QoS-Aware Federated Anomaly Detection', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 195–210.

Pasupuleti, MK 2023m, 'Real-Time Causal cademic and Industrial Research Innovations (IJAIRI), vol. 03, no. 03, pp. Inference on Spark: Structured Streaming for Policy Lift', International Journal of A53–63.

Pasupuleti, MK 2023n, 'Secure Intelligent Systems: Post-Quantum Cryptography and Topological Modeling for Big Data Robotics', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 01, pp. 18–31.

Pasupuleti, MK 2023o, 'Topological and Software-Driven Advances in Robotics: Leveraging MATLAB-Simulink and Semiconductor Innovation for Adaptive Autonomous Systems', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 02, no. 01, pp. 1–17.

Pasupuleti, MK 2023p, 'AI-Augmented Real-World Evidence in Stata: Targeted Learning for Treatment Effects', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 12, pp. 127–144.

Pasupuleti, MK 2023q, 'Quantum-Inspired Optimizers for Cloud–Fog Scheduling: Near-Optimal SLAs under Energy Caps', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 03, pp. 17–28.

Pasupuleti, MK 2023r, 'Quantum-Inspired Scheduling in CloudSim: Energy-Aware Cloud–Fog Orchestration', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 03, no. 06, pp. 178–194.

Pasupuleti, MK 2025a, 'Fault-Tolerant Quantum Computing via Macroscopic Quantum Tunnelling in Superconducting Circuits', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 05, no. 12, pp. 1–21.

Pasupuleti, MK 2025b, 'Quantum-Accelerated Intelligence: Hybrid AI Systems for National-Scale Advantage', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 05, no. 12, pp. 22–49.

Pasupuleti, MK 2025c, 'Quantum-Limited Sensing-to-Action for GPS-Denied Navigation in Autonomous Robots', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 05, no. 12, pp. 264–290.

Pasupuleti, MK 2025d, 'Secure Computation After Quantum: PQC, Confidential Clouds, and Policy-Grade Auditability', International Journal of Academic and Industrial Research Innovations (IJAIRI), vol. 05, no. 12, pp. 198–222.