Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 3, 574-584
2026
Publisher: Learning Gate
DOI: 10.55214/2576-8484.v10i3.12505

A proposed AI of the AI ecosystem for the orchestration, evaluation, and governance of autonomous AI systems

Ronald Brisebois[1], Apollinaire Nadembega[1], Alain Abran[2*]
[1]InMedia Technologies, Montréal, Canada; ronald.brisebois@inmediatech.co (R.B.). anadembega@gmail.com (A.N.).
[2]École de technologie supérieure, Université du Québec, Montréal, Canada; alain.abran@etsmtl.ca (A.A.).

Abstract: Coordinating autonomous and heterogeneous artificial intelligence (AI) systems is difficult because their capabilities, domain specializations, interfaces, and output structures vary widely and rapidly evolve. The purpose of this study is to address these issues through the design of an *AI of the AI* (AoA) ecosystem as a hierarchical system-of-systems framework that orchestrates and challenges multiple independent AI agents while preserving their operational autonomy. In this approach, the AoA integrates three interconnected layers: (i) a configuration and longitudinal performance-tracking layer that maintains operational parameters, version histories, and domain-specific performance profiles; (ii) a moderated collaboration and evaluation layer that enables indirect coordination through standardized response schemas, structured competition, and cross-agent benchmarking; and (iii) a context-dependent authority and data-source evaluation layer that weights AI outputs based on source credibility, reference quality, and domain relevance, supported by federated ontologies for semantic alignment and conflict reconciliation. In conclusion, this ecosystem also includes multi-stage alert mechanisms that detect drift, inconsistencies, conflicts, and emergent patterns, enabling continuous self-assessment, adaptive governance, and iterative ranking updates. In terms of practical implications, this AoA is designed to interoperate with existing information technology (IT) infrastructure and is additionally forward-compatible with emerging quantum computing platforms, providing a scalable foundation for orchestrating and governing heterogeneous AI ecosystems.

Keywords: *AI benchmarking, AI governance, AI orchestration, Heterogeneous AI integration, Multi-agent systems, System-of-systems architecture.*

1. Introduction

Multiple artificial intelligence (AI) applications exhibit heterogeneous capabilities, domain-specific expertise, and inherent limitations. AI broadly refers to computational systems capable of performing tasks that traditionally require human intelligence, including reasoning, decision-making, perception, and natural language understanding. Contemporary AI examples include large-language models such as ChatGPT, Gemini, and Claude, and Chinese AI models like DeepSeek, which demonstrate competitive reasoning and generative performance at comparatively lower computational costs.

This diversity raises a central question: How can one identify the most relevant and robust response to a problem by leveraging the complementary capabilities of autonomous and heterogeneous AI systems? Rather than constituting a barrier, this heterogeneity can be harnessed within a structured framework to enhance decision-making, problem-solving, and collective intelligence.

The proposed AI of the AI (AoA) ecosystem addresses this challenge by orchestrating multiple independent AI systems through moderated competition, indirect collaboration, and authority-weighted evaluation, while enforcing standardized output formats. The AoA combines principles from systems engineering, multi-agent systems, and ensemble learning and is adapted to current AI models accessible through standardized application programming interfaces (APIs). Importantly, the AoA is designed to

History: Received: 16 January 2026; Revised: 20 February 2026; Accepted: 25 February 2026; Published: 24 March 2026
* Correspondence: alain.abran@etsmtl.ca

interconnect seamlessly with existing computational infrastructures and is forward compatible with emerging quantum computing platforms, enabling high-performance computation, massive parallelism, and complex problem-solving capabilities.

AoA is positioned at the intersection of several research traditions:

- System of Systems (SoS) paradigm: The AoA extends classical SoS principles by integrating autonomous and heterogeneous AI applications to generate emergent outcomes while preserving operational independence. This structure allows collective intelligence to emerge without centralized control, thereby providing a foundation for hybrid classical-quantum computations.
- Multi-Agent Systems (MAS): Unlike traditional MAS approaches, AoAs do not require co-designed agents or direct inter-agent communication. Collaboration emerges indirectly via standardized outputs, moderated competition, domain-specific ranking, and cross-evaluation mechanisms, enabling scalable interactions among diverse AI models, including potential quantum-enhanced agents.
- Ensemble learning: Conceptually similar to ensemble methods, AoA combines multiple models to improve overall performance. Unlike classical ensembles, it operates in a heterogeneous black-box context without centralized control over model internals or training, supporting independently evolving AI systems and quantum-accelerated evaluations.

The operational framework relies on prompt standardization, authority-weighted evaluation, and comparative performance analysis, enabling systematic benchmarking, optimal output selection, and the progressive evolution of collective intelligence. By standardizing and cross-assessing domain-specific outputs, the AoA fosters emergent intelligence, where aggregated and validated responses surpass the capabilities of individual AI systems.

The remainder of this paper is organized as follows. Section 2 reviews related works on collaborative AI frameworks, moderated competition, and autonomous AI applications. Section 3 presents the AoA ecosystem, including its core concepts, definitions, and layers. Section 4 presents the operational mechanisms and principles of evaluation. Section 5 concludes the study and outlines future research directions.

2. Related Work

Recent research has explored various approaches to collaborative and autonomous AI systems. Several frameworks aim to combine multiple AI tools to deliver improved and reliable responses to users. However, these efforts are often fragmented, domain-specific, and insufficiently validated.

2.1. AI Collaborative Frameworks

This subsection focuses on collaborative AI frameworks.

Joshi [1] analyzed emerging AI agent frameworks designed to support autonomous decision-making and multi-agent collaboration across leading platforms, including LangGraph, CrewAI, OpenAI Swarm, AutoGen, and IBM Watsonx.ai. This study highlighted the ecosystem design, strengths, limitations, and domain-specific applications of these frameworks. Although it provides a valuable overview of the current landscape, it primarily catalogs existing frameworks without systematically comparing their performance, scalability, or interoperability. Furthermore, issues related to trust, transparency, and accountability have not yet been explored in depth.

Yim [2] proposed a framework to conceptualize AI literacy by analyzing existing theoretical models and introducing an inclusive, intelligence-based AI literacy framework for primary education. The framework advocates a shift from a traditional humanistic perspective to a post-humanistic approach that incorporates AI thinking and nonhuman agency. Although this contribution is conceptually innovative, the framework remains theoretical and has not yet been validated in real-world classroom environments. Additionally, although this article critiques constructionist approaches, it does not provide a detailed analysis of their pedagogical limitations or comparative empirical evidence.

Song et al. [3] proposed a framework to guide the design of inclusive AI learning environments based on three principles of Universal Design for Learning (UDL): engagement, representation, action, and expression. Although this framework addresses inclusivity in AI education, its effectiveness in various educational contexts remains unclear. Moreover, it assumes a level of teacher preparedness that may not align with real-world conditions, particularly in under-resourced settings. The lack of robust assessment mechanisms also hinders a clear evaluation of scalability and overall impact.

Geske et al. [4] introduced an AI-supported collaborative decision-making (CDM) framework for airline disruption management to enhance operational efficiency and responsiveness. The authors emphasized the need for data-driven CDM, AI-enabled management of complex systems, and performance-oriented disruption management. However, the proposed framework does not sufficiently address challenges related to human–AI interaction, including trust, transparency, and accountability. Additionally, the model remains largely theoretical with limited validation in live operational environments.

Holmström [5] proposed an AI readiness framework to assess organizational preparedness for integrating AI into broader digital transformation initiatives. The framework considers technological infrastructure, leadership commitment, data maturity, and workforce adaptability. Although it provides a structured planning tool, it remains largely theoretical and has yet to demonstrate its relevance or reliability across different sectors. Its design overlooks key contextual factors, such as regulatory constraints, cultural dynamics, and essential ethical or human-centered considerations, including algorithmic bias and workforce implications, resulting in a framework that risks being disconnected from practical reality.

Tang et al. [6] proposed a cloud–edge collaborative gaming framework that integrated AI-powered foveated rendering and super-resolution techniques to optimize visual performance and reduce bandwidth consumption. Although the framework addresses challenges related to latency, scalability, and energy efficiency, its evaluation is primarily limited to simulations or controlled experiments with minimal validation in large-scale real-world deployments.

Overall, existing collaborative AI frameworks provide valuable conceptual contributions but exhibit several recurring limitations.

- a real-world deployment evidence,
- geographic and cultural bias,
- insufficient consideration of ethical and human-centered dimensions,
- limited adaptability to diverse operational contexts.

These gaps highlight the need for integrative ecosystems that are capable of benchmarking, coordinating, and governing heterogeneous AI systems across domains.

2.2. AI Frameworks of Moderated Competition

This subsection focuses on AI frameworks based on moderated competition, drawing on the principles of algorithmic accountability, organizational theory, and market design.

Alsheibani et al. [7] examined AI adoption among Australian firms using the Technology–Organization–Environment (TOE) and Diffusion of Innovation (DOI) frameworks. This study identifies technological readiness, organizational culture, and regulatory pressure as key adoption drivers. However, the analysis does not sufficiently address the limitations of these frameworks in capturing the fluid and cross-sectoral nature of AI technologies. Moreover, the geographic focus restricts the generalizability of our findings.

Al-Surmi et al. [8] investigated the integration of marketing strategy, IT strategy, and AI capabilities using a hybrid methodological approach that combined structural equation modeling and artificial neural networks. While this study demonstrates performance improvements in specific organizational contexts, it does not examine how AI-based decision-making evolves or adapts to dynamic market conditions, thus limiting insights into long-term sustainability.

Fousiani et al. [9] conducted a literature review linking organizational psychology and technology acceptance, highlighting how competitive climates can reduce AI acceptance owing to stress and job insecurity. The authors showed that leadership styles that emphasize responsibility can mitigate these effects. However, this study relies heavily on Western-centric leadership models and does not sufficiently consider structural organizational factors or cross-cultural variability.
Badghish and Soomro [10], building on the TOE framework, analyzed AI adoption among small and medium-sized enterprises (SMEs) and its impact on sustainable business performance. Although the findings identify key technological, organizational, and environmental drivers, the study remains context-specific and does not address industry-level or regional differences in AI adoption challenges.

2.3. Autonomous AI Applications

The rapid advancement of artificial intelligence (AI) technologies has enabled the emergence of autonomous systems capable of performing complex tasks with minimal human intervention. Autonomous AI applications are increasingly being deployed in domains such as transportation, robotics, healthcare, finance, and smart infrastructure.

Ma et al. [11] provided a comprehensive survey of AI applications in autonomous vehicle development, categorizing techniques into perception, decision making, and control. Although this study offers a clear taxonomy and highlights interdisciplinary challenges, it emphasizes the need for explainable AI, safety validation, and hybrid approaches.

Casetti et al. [12] examined AI- and machine-learning (ML)-based services for 6G-connected autonomous vehicles, emphasizing edge computing, federated learning, and adaptive control. Although the proposed frameworks exhibit potential, their practical value is severely limited owing to the absence of large-scale deployment data. Without evidence of how they perform in real-world conditions, they risk retaining largely theoretical constructs with limited usefulness for practitioners or decision makers.

Maddali [13] explored autonomous AI agents for real-time data transformation and extracted, transformed, and loaded (ETL) automation, highlighting reinforcement learning and self-adaptive algorithms. Although conceptually robust, the proposed ecosystem has not been extensively tested in large-scale or heterogeneous production environments.

Veitch and Andreas Alsos [14] reviewed human–AI interactions in autonomous ship systems, emphasizing trust, situational awareness, and human-centered design. However, this study relies primarily on simulations and conceptual models, with limited empirical evidence from operational systems.

Alanne and Sierla [15] reviewed machine-learning applications in smart buildings, focusing on energy management, predictive maintenance, and occupant behavior modeling. Despite its breadth, this review lacks comparative performance analyses across building types and gives limited attention to interoperability challenges and data-privacy concerns.

Overall, autonomous AI applications face persistent challenges, including limited contextual awareness, opacity of decision-making processes, and heavy dependence on large, high-quality datasets. These limitations raise concerns regarding accountability, trust, bias, and generalizability, underscoring the need for an ecosystem capable of coordinating, evaluating, and governing autonomous AI systems across domains.

3. Proposed AoA Ecosystem

3.1. Concepts and Definitions

This subsection presents a set of concepts and terms adopted for the proposed AoA ecosystem.

- AI of AI (AoA): An integrated evolutionary ecosystem composed of autonomous and heterogeneous AI systems. Each AI is queried using tailored prompts and constrained to standardized response formats. The AoA operates as a system in which components preserve autonomy while contributing to a collective objective.

- Moderated Competitive Evaluation: A process in which multiple AI systems independently address the same task. Their outputs, formatted according to predefined criteria, were compared to identify convergence, divergence, and relative performance within a specific domain.
- Emergence of the Optimal Response: The optimal response is derived through authority-weighted consensus building, multi-criteria scoring, and periodic ranking updates rather than simple majority voting.
- Domain-based Performance: Comparative evaluation generates performance profiles for each AI system, reflecting accuracy, consistency, and contextual relevance within specific domains.
- Learning and Evolution: Continuous feedback on individual and comparative performance challenges and enables incremental system learning and adaptive refinement over time.

3.2. Ecosystem Overview

The proposed AI of AI (AoA) ecosystem is organized into three interconnected layers designed to orchestrate, evaluate, and optimize multiple autonomous AI systems.

1. AI Configuration and Performance Tracking and Evolution Layer:
   To maintain the operational parameters, version histories, and performance information for all integrated AI agents. It monitors outputs, consistency, and domain-specific effectiveness, enabling adaptive reconfiguration, evolutionary learning, and predictive modeling. High-performing agents are dynamically prioritized, while underperforming agents are flagged for retraining or deprecation.
2. Moderated Collaboration and Evaluation Layer:
   Orchestrates indirect collaboration and regulates competition among autonomous AI systems. Cross-agent benchmarking, standardized output protocols, and ontology-aware evaluations ensure comparability, despite heterogeneous systems. Outputs are aggregated using authority-weighted scoring and conflict resolution methods, enabling emergent collaboration where competitive interactions drive collective improvement without direct knowledge sharing.
3. Context-Dependent Authority and Data Source Evaluation Layer:
   To evaluate the credibility and relevance of external information sources and AI-generated content. The outputs were weighted according to the author's credibility, reference quality, and domain relevance. Federated ontologies support semantic alignment, conflict reconciliation, and cross-domain mapping. Multistage alert mechanisms detect inconsistencies, anomalies, and emerging patterns by feeding a meta-intelligent governance layer into the adaptive oversight.

Overall, these layers provide a hierarchical, yet interdependent ecosystem that combines real-time performance tracking, moderated competition, and authority-based evaluation. AoA produces optimized, semantically coherent, and domain-adapted collective intelligence, supporting scalable and transparent coordination of autonomous AI ecosystems while remaining compatible with existing infrastructure and future quantum computing platforms.

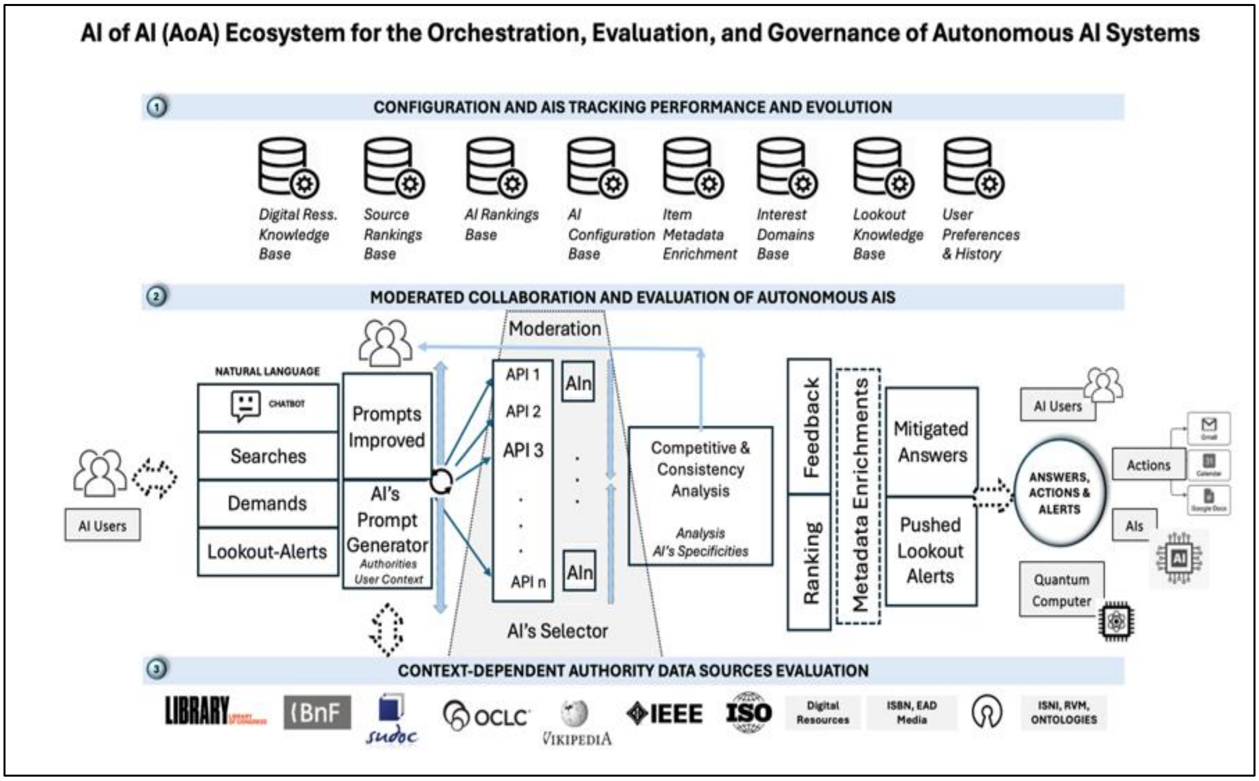

Figure 1.
Layered overview of the AI of AI (AoA) ecosystem.

The configuration, AIs tracking performance, and the evolution layer monitor the agent evolution. The Moderated Collaboration layer manages indirect cooperation and output aggregation. The Authority and Data Source Evaluation layer ensures semantic alignment, source credibility, and context-sensitive prioritization. These three layers are described in the following subsections.

3.3. Configuration and AI Performance Tracking and Evolution

The AoA ecosystem implements an integrative framework that is designed to coordinate multiple autonomous AI agents through moderated competition and indirect collaboration. Inspired by systems engineering, ecosystem**s** emphasize scalability, robustness, and adaptability across heterogeneous AI components.

The key components of this layer include – see Figure 1.

1. Digital Resource Knowledge Base: A centralized repository for structured knowledge, documents, and digital content that supports efficient semantic retrieval.
2. Source Rankings Base: An authority-weighted evaluation framework assessing author credibility, reference quality, and source reliability.
3. AI Rankings Base: A performance-based ranking system enabling moderated competition and comparative evaluation.
4. AI Configuration Base: A repository of operational parameters, interrogation strategies, and coordination protocols.
5. Item Metadata Enrichment: Automated semantic enrichment using cross-domain ontologies and context-aware annotations.
6. Interest Domains Base: Domain-specific knowledge profiles supporting personalized and context-adapted responses.

DOI: 10.55214/2576-8484.v10i3.12505

7. Lookout Knowledge Base: An alert and monitoring system capturing deviations, inconsistencies, conflicts, and emergent patterns.
8. User Preferences and History: Interaction history and preference modeling to tailor outputs and refine system behavior.

The Configuration and AI Performance Tracking layer can be formally represented as a tuple:

$$L_{config} = \langle A, D, S, R, O, M, P, \Lambda \rangle$$

where:

- A denotes the set of autonomous AI agents,
- D represents domain-specific knowledge profiles,
- S corresponds to authority-weighted source evaluations,
- R denotes dynamic AI performance rankings,
- O represents federated ontologies and semantic mappings,
- M captures enriched metadata and annotations,
- P denotes user preferences and interaction histories,
- Λ represents alerting, monitoring, and performance evolution mechanisms.

Semantic alignment mechanisms, including concept harmonization, ontology mapping, and conflict reconciliation, mitigate interpretative inconsistencies and support coherent knowledge exchange. Standardized response, annotation, and evaluation protocols ensure reproducibility, comparability, and consistency in cross-domain assessments. Additionally, meta-intelligent alert mechanisms facilitate continuous self-evaluation, adaptive benchmarking, and longitudinal performance monitoring, thereby strengthening epistemic governance, accountability, and the overall resilience of the federated knowledge ecosystem.

While the Configuration and Performance Tracking layer ensures semantic coherence, governance, and longitudinal evaluation, the Moderated Collaboration and Evaluation layer operationalizes these foundations by orchestrating competitive interaction, cross-assessment, and the aggregation of outputs produced by autonomous AI agents.

This layer maintains longitudinal state and governance metadata, enabling adaptive reconfiguration, performance evolution, and semantic coherence across heterogeneous AI systems.

3.4. Moderated Collaboration and Evaluation of Autonomous AIs

Three primary interaction types were supported, namely search, suggestion, and lookout alerts. Each AI is accessed through its native API or interface with prompts adapted to system-specific characteristics, while enforcing standardized output structures.

For moderated collaboration and evaluation of outputs from autonomous AI, the AoA **ecosystem** allows

- Multiple AI systems independently respond to identical queries with mandatory elements such as confidence scores, references, and structured summaries. The outputs were stored in a shared repository to enable systematic cross-comparisons.
- Responses were evaluated based on their coherence, robustness, and contextual relevance. Each AI receives domain-specific scores that inform the dynamic rankings. Feedback loops refine prompt design and interrogation strategies, enabling continuous improvement.
- Validated outputs are incorporated into a global knowledge repository, supporting cumulative learning.

End users interact with the AoA ecosystem by issuing searches, submitting explicit requests, or activating lookout alerts. User inputs are processed by a Prompt Generation module that contextualizes intent according to the domain of interest and adapts queries to the specific constraints and capabilities of the selected AI systems. Standardized prompt structures are generated and dispatched to multiple autonomous AIs through their native interfaces.

A moderation and AI selection mechanism dynamically governs agent participation, evaluation criteria, and governance policies based on domain relevance. AI outputs were subjected to competitive and consistency analyses, enabling cross-model comparisons, interrelation detection, and domain-specific performance ranking. Feedback loops update AI rankings, interrogation strategies, and governance parameters.

In parallel, the response metadata was enriched using federated ontologies and context-aware annotations to support semantic alignment, traceability, and longitudinal performance monitoring. The validated outputs are aggregated and propagated through lookout alerts when anomalies, conflicts, or emergent patterns are detected.

AoA outputs may be consumed by human users, automated execution systems, other AI agents, or quantum computing platforms, thereby enabling scalable, adaptive, and future-compatible collective intelligence.

3.5. Context-Dependent Authority Data Source Evaluation

Authority data sources encompass a diverse range of institutional, community-driven, and machine-readable knowledge repositories, including national and international institutions such as the Library of Congress, Bibliothèque Nationale de France (BnF), Sudoc, and the Online Computer Library Center (OCLC), as well as collaboratively curated resources such as Wikipedia (see Figure 1). Additional authoritative inputs include controlled vocabularies and identifier systems (e.g., Répertoire de vedettes-matière (RVM) and International Standard Name Identifier (ISNI)), domain-specific digital resources, and formal ontologies.

Within the proposed AoA ecosystem, these authority sources were not treated as uniformly reliable or universally applicable. Instead, each source was dynamically weighted according to multiple context-dependent criteria, including domain relevance, epistemic reliability, update frequency, provenance transparency, and task-specific applicability. The weighting process enables the system to privilege institutional or highly curated sources in contexts requiring bibliographic precision or formal validation while allowing broader or community-maintained sources to contribute to exploratory, semantic enrichment, or cross-domain inference tasks.

By integrating authority-weighted evaluation into the decision and ranking mechanisms, the AoA framework supports the nuanced cross-validation of information, reduces the risk of epistemic bias, and enhances the robustness and explainability of outputs. This context-sensitive authority assessment mechanism plays a central role in ensuring that the most appropriate and credible knowledge sources are mobilized for each problem domain and use case.

4. Discussion

4.1. Challenges

Despite the potential of AI in AI (AoA) ecosystems, several structural and operational challenges have emerged.

First, heterogeneity across AI interfaces, APIs, and internal reasoning paradigms is crucial. Contemporary AI systems vary significantly in prompt schemas, output formats, epistemic assumptions, degrees of explainability, complicated orchestration, and fair comparisons [10, 16, 17].

Second, achieving output standardization across independent and evolving systems is nontrivial. Even when tasks are semantically equivalent, differences in reasoning depth, abstraction level, and linguistic framing introduce variability that must be normalized through ontology-aware response protocols and postprocessing layers [6, 7, 11].

Third, the definition of the evaluation, aggregation, and ranking criteria by the domain of interest requires continuous calibration. Domain relevance, authority weighting, and performance evaluation are inherently contextual and must adapt to shifting knowledge landscapes and rapid AI version evolution [4, 12].

Finally, epistemic drift and source inconsistency pose challenges to long-term governance. As the model updates and external sources change, maintaining ranking stability and longitudinal comparability demands active monitoring and alert mechanisms [5, 18].

DOI: 10.55214/2576-8484.v10i3.12505

4.2. Contributions

This study makes several theoretical and architectural contributions to the field of multi-agent and collaborative AI systems.

First, it introduces a methodological framework for benchmarking heterogeneous AI systems, extending beyond static performance evaluation to include authority-weighted consistency, domain sensitivity, and longitudinal stability [10, 18].

Second, the AoA ecosystem demonstrates emergent collaboration driven by regulated competition, in which indirect interaction among autonomous AIs leads to collective improvement without requiring explicit coordination or shared internal states [6, 9, 11].

Third, the ecosystem establishes an evolutionary intelligence model grounded in iterative learning, cross-evaluation, and adaptive ranking, enabling continuous optimization across domains of interest [12, 17].

Finally, the integration of federated ontologies and multi-stage alert mechanisms contributes a practical approach to semantic interoperability, conflict reconciliation, and epistemic governance in distributed AI ecosystems [5, 10, 18].

4.3. Future Directions

Several research and development directions emerge from this work.

A primary avenue is the automation of adaptive and optimized prompt generation, allowing the ecosystem to dynamically tailor interrogation strategies based on AI performance, domain sensitivity, and historical outcomes [16, 19].

Another avenue involves scalable integration of new AI agents while preserving coherence, ranking stability, and epistemic alignment, particularly as agentic and autonomous AI systems proliferate [17, 20].

Future research should also explore hybrid strategies that combine competition and collaboration, where selective cooperation is introduced under controlled conditions to enhance reasoning depth, reduce redundancy, and improve convergence in complex problem spaces [9, 11].

Finally, extending the Lookout Knowledge Base into a predictive meta-intelligence layer capable of forecasting performance degradation, epistemic divergence, or governance risks represents a critical step toward sustainable collective machine intelligence [5, 18].

5. Conclusion

The AI of the AI (AoA) ecosystem, explicitly inspired by the system-of-systems (SoS) paradigm, provides a scalable and modular framework for orchestrating heterogeneous AI systems. By combining moderated competition, standardized outputs, cross-evaluation, and authority-weighted ranking, the AoA enables the emergence of distributed collective intelligence capable of outperforming individual AI systems in terms of precision, relevance, and robustness.

Key Contributions:

- Formalization of a hierarchical AoA framework to compare, compete, validate, and continuously improve heterogeneous AI systems.
- Introduction of indirect collaboration mechanisms, where structured competition drives collective advancement.
- Establishment of an evolving domain-aware knowledge base maintained through iterative interrogation, evaluation, and cross-learning among multiple AI agents.
- Recognition of AI evaluation as a periodic and adaptive process accommodates the rapid evolution of AI versions and domain-specific ranking updates.
- Seamless integration with existing computational infrastructures to ensure interoperability.
- Forward-compatible design for emerging quantum computing platforms enables accelerated computation, massive parallelism, and enhanced problem-solving for complex tasks.

Applications:

- Short-term: Dynamic AI benchmarking, consolidated multi-AI response generation, resilient and authority-aware knowledge repositories, and integration with current IT infrastructure.
- Long-term: Distributed intelligent ecosystems capable of continuous adaptation, self-assessment, and epistemic governance across heterogeneous AI landscapes with the potential to leverage quantum-enhanced computation for advanced analytics and collective intelligence.

Overall, the AoA represents a flexible, scalable, and future-proof ecosystem that bridges the current computational infrastructure and emerging quantum technologies, thereby maximizing collective intelligence and problem-solving performance in heterogeneous AI ecosystems.

Transparency:

The authors confirm that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.

Copyright:

References

[1] S. Joshi, "Review of autonomous systems and collaborative AI agent frameworks," *International Journal of Science and Research Archive*, vol. 14, no. 2, pp. 961–972, 2025. https://doi.org/10.30574/ijsra.2025.14.2.0439

[2] I. H. Y. Yim, "A critical review of teaching and learning artificial intelligence (AI) literacy: Developing an intelligence-based AI literacy framework for primary school education," *Computers and Education: Artificial Intelligence*, vol. 7, p. 100319, 2024. https://doi.org/10.1016/j.caeai.2024.100319

[3] Y. Song, L. R. Weisberg, S. Zhang, X. Tian, K. E. Boyer, and M. Israel, "A framework for inclusive AI learning design for diverse learners," *Computers and Education: Artificial Intelligence*, vol. 6, p. 100212, 2024. https://doi.org/10.1016/j.caeai.2024.100212

[4] A. M. Geske, D. M. Herold, and S. Kummer, "Integrating AI support into a framework for collaborative decision-making (CDM) for airline disruption management," *Journal of the Air Transport Research Society*, vol. 3, p. 100026, 2024. https://doi.org/10.1016/j.jatrs.2024.100026

[5] J. Holmström, "From AI to digital transformation: The AI readiness framework," *Business Horizons*, vol. 65, no. 3, pp. 329-339, 2022. https://doi.org/10.1016/j.bushor.2021.03.006

[6] X. Tang, Y. Xu, F. Ouyang, L. Zhu, and B. Peng, "A cloud-edge collaborative gaming framework using AI-powered foveated rendering and super resolution," *International Journal on Semantic Web and Information Systems*, vol. 19, no. 1, pp. 1-19, 2023. https://doi.org/10.4018/IJSWIS.321751

[7] S. Alsheibani, C. Messom, and Y. Cheung, "Re-thinking the competitive landscape of artificial intelligence," in *Proc. 53rd Hawaii Int. Conf. on System Sciences*, 2020.

[8] A. Al-Surmi, M. Bashiri, and I. Koliousis, "AI based decision making: Combining strategies to improve operational performance," *International Journal of Production Research*, vol. 60, no. 14, pp. 4464-4486, 2022. https://doi.org/10.1080/00207543.2021.1966540

[9] K. Fousiani, G. Michelakis, P. A. Minnigh, and K. M. De Jonge, "Competitive organizational climate and artificial intelligence (AI) acceptance: the moderating role of leaders' power construal," *Frontiers in Psychology*, vol. 15, p. 1359164, 2024. https://doi.org/10.3389/fpsyg.2024.1359164

[10] S. Badghish and Y. A. Soomro, "Artificial intelligence adoption by SMEs to achieve sustainable business performance: application of technology–organization–environment framework," *Sustainability*, vol. 16, no. 5, p. 1864, 2024. https://doi.org/10.3390/su16051864

[11] Y. Ma, Z. Wang, H. Yang, and L. Yang, "Artificial intelligence applications in the development of autonomous vehicles: A survey," *IEEE/CAA Journal of Automatica Sinica*, vol. 7, no. 2, pp. 315-329, 2020. https://doi.org/10.1109/JAS.2020.1003021

[12] C. Casetti, C. F. Chiasserini, F. Dressler, A. Memedi, D. Gasco, and E. M. Schiller, "AI/ML-based services and applications for 6G-connected and autonomous vehicles," *Computer Networks*, vol. 255, p. 110854, 2024. https://doi.org/10.1016/j.comnet.2024.110854

[13] R. Maddali, "Autonomous ai agents for real-time data transformation and etl automation," *International Journal of Research and Analytical Reviews*, vol. 10, no. 3, p. 12, 2023.

[14] E. Veitch and O. Andreas Alsos, "A systematic review of human-AI interaction in autonomous ship systems," *Safety Science*, vol. 152, p. 105778, 2022. https://doi.org/10.1016/j.ssci.2022.105778

[15] K. Alanne and S. Sierla, "An overview of machine learning applications for smart buildings," *Sustainable Cities and Society*, vol. 76, p. 103445, 2022. https://doi.org/10.1016/j.scs.2021.103445

[16] J. Li, Q. Zhang, and X. Wang, "Benchmarking heterogeneous AI interfaces: Toward unified evaluation frameworks," *Machine Learning Systems Journal*, vol. 6, no. 7, pp. 314–332, 2020.

[17] A. Gupta and P. Ramesh, "Collaborative multi-agent learning and evaluation in autonomous systems," in *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 2019.

[18] R. Harris and K. Singh, "AI governance and longitudinal comparability in multi-model ecosystems," *Journal of Machine Intelligence Governance*, vol. 2, no. 1, pp. 55–71, 2023.

[19] S. Lee and Y. Kim, "Adaptive prompt generation strategies for AI-on-AI interrogation," *AI Optimization and Techniques*, vol. 18, no. 2, pp. 99–117, 2024.

[20] Y. Zhang and H. Chen, "Scalable integration of autonomous agents with coherence preservation," *Journal of Distributed Intelligence Systems*, vol. 8, no. 4, pp. 87–105, 2023.