Human-Computer Interaction & Emerging Technologies, Vol. 195, 2025, 132–141
https://doi.org/10.54941/ahfe1006231

AI Agents as Knowledge Navigators: A Conceptual Framework for Multi-Agent Systems in Scientific Knowledge Management

Paolo Gemelli, Laura Pagani, Mario Ivan Zignego, and Alessandro Bertirotti

Università Studi di Genova, Dipartimento Architettura e Design, Genova, Italy

ABSTRACT

The rapid and continuous expansion of scientific literature has led to an unprecedented increase in the volume of knowledge produced, significantly complicating its organization, retrieval, and effective utilization. Researchers face considerable challenges in managing this vast information landscape, particularly in terms of identifying relevant studies, maintaining contextual integrity, and integrating knowledge across multiple disciplines. Traditional database-driven search engines and static indexing methods often fall short in addressing these issues, as they lack the capacity to dynamically interpret scientific discourse, establish meaningful cross-disciplinary connections, and facilitate real-time knowledge synthesis. To overcome these limitations, this study proposes a novel approach to scientific knowledge management through a multi-agent artificial intelligence (AI) system. This system is designed to enhance interactive, context-aware, and dynamic information retrieval by leveraging a network of AI agents, each specialized in distinct scientific domains. These agents operate collaboratively, employing advanced adaptive learning mechanisms, context-sensitive reasoning, and cooperative problem-solving strategies to improve the organization and accessibility of scientific knowledge. The multi-agent framework integrates state-of-the-art Natural Language Processing (NLP) techniques, transformer-based architectures, and knowledge graph methodologies to provide a more nuanced understanding of scientific texts. By doing so, it enables automated cross-domain conceptual linking, validation of theoretical and experimental claims, and the seamless integration of newly acquired information into existing knowledge structures. Moreover, the system incorporates reinforcement learning mechanisms to continuously optimize its retrieval and synthesis processes based on user interactions and evolving research trends. Beyond its immediate applications in knowledge retrieval, the proposed system fosters a paradigm shift in how scientific research is conducted, promoting more effective interdisciplinary collaboration and accelerating the development of innovative ideas. By facilitating automated synthesis of vast scientific corpora, it enables researchers to explore novel hypotheses, detect previously unrecognized connections between fields, and refine theoretical models with enhanced precision. Additionally, the ability of AI agents to autonomously process complex information reduces cognitive load on researchers, allowing them to focus on higher-order analytical tasks and creative problem-solving. This study contributes to the field of scientific knowledge management by introducing a scalable and adaptive AI-driven framework capable of supporting research in high-information-density environments. By bridging the gaps between disparate scientific disciplines and facilitating intelligent, real-time knowledge synthesis, the proposed system has the potential to revolutionize the way researchers interact with and utilize scientific information, ultimately advancing the efficiency and impact of knowledge discovery in the modern research landscape.

Keywords: Multi-agent AI systems, Scientific knowledge management, Interdisciplinary collaboration

INTRODUCTION

The exponential expansion of scientific literature poses significant challenges in knowledge retrieval and organization. With an annual publication rate of approximately 2.5 million new research papers (Bettin et al., 2022) methods struggle to provide contextual understanding and dynamic interactions necessary for effective scientific discourse. This paper introduces a conceptual framework for a multi-agent AI system designed to enhance scientific knowledge organization and retrieval through collaboration, contextual reasoning, and adaptability.

CHALLENGES IN SCIENTIFIC KNOWLEDGE MANAGEMENT

Key obstacles in current knowledge management approaches include information overload, where researchers spend an estimated 12 hours per week searching for relevant literature (Eppler, 2004). Contextual limitations also pose significant challenges, as conventional search methods fail to capture nuanced relationships between topics. Furthermore, barriers to cross-disciplinary research remain prevalent, as existing systems lack mechanisms to integrate findings from multiple domains effectively. Addressing these challenges requires innovative AI-driven solutions capable of managing vast amounts of scientific information while preserving contextual integrity and fostering interdisciplinary collaboration.

MULTI-AGENT SYSTEM ARCHITECTURE

The proposed system employs a multi-agent architecture with specialized AI agents designed for distinct scientific domains while maintaining collaborative capabilities. These agents operate on a modified transformer architecture with domain-specific adaptations and training protocols (Table 1).

Each agent plays a crucial role in ensuring the system's efficiency, providing domain-specific expertise while collaborating to enhance interdisciplinary knowledge synthesis. The proposed system employs a multi-agent architecture with specialized AI agents designed for distinct scientific domains while maintaining collaborative capabilities. These agents operate on a modified transformer architecture with domain-specific adaptations and training protocols.

The Theoretical Sciences Agent (TSA) focuses on fields such as physics, mathematics, and computational modelling, integrating LaTeX interpretation, equation parsing, and theorem verification. It ensures precision in mathematical notation processing and logical consistency in theoretical proofs. The Experimental Sciences Agent (ESA) specializes in validating experimental designs, comparing methodologies, and optimizing research protocols using deep learning models trained on laboratory datasets. It also employs causal inference models to detect bias and assess reproducibility.

Table 1: Performance metrics of different agents.

Agent	Focus	Performance Metrics
Theoretical Sciences Agent (TSA)	Physics, mathematics, and computational modelling. Integrates LaTeX interpretation, equation parsing, and theorem verification.	Mathematical Expression Parsing Accuracy $\geq 95\%$, Logical Consistency in Proof Verification $\geq 92\%$, Cross-Theorem Relationship Identification Accuracy $\geq 90\%$
Experimental Sciences Agent (ESA)	Validating experimental designs, comparing methodologies, and optimizing research protocols using deep learning models. Detects bias and assesses reproducibility.	Experimental Protocol Consistency Score $\geq 90\%$, Bias Detection Accuracy $\geq 88\%$, Reproducibility Assessment Score $\geq 90\%$
Engineering Sciences Agent (ENA)	Mechanical, aerospace, and robotics engineering. Uses computational models and machine learning techniques to interpret CAD schematics, predict structural integrity, and ensure system integration compliance.	CAD Interpretation Accuracy $\geq 93\%$, Structural Integrity Prediction Score $\geq 92\%$, System Integration Validation Score $\geq 90\%$
Natural Sciences Agent (NSA)	Earth sciences, environmental studies, and space exploration. Integrates geospatial data analysis, climate modelling, and ecological system assessments.	Geospatial Pattern Recognition Accuracy $\geq 90\%$, Climate Model Interpretation Precision $\geq 91\%$, Environmental Impact Forecasting Reliability $\geq 89\%$
Methodology Agent (MTA)	Enhances research rigor by evaluating experimental frameworks for logical consistency and reproducibility. Employs Bayesian optimization models for research structure refinement.	Methodological Consistency Evaluation Score $\geq 93\%$, Logical Fallacy Detection Accuracy $\geq 92\%$, Experimental Design Optimization Score $\geq 91\%$
Decision Sciences Agent (DSA)	Supports research strategy and interdisciplinary integration by optimizing resource allocation and risk assessment through decision-tree processing models.	Multi-Criteria Decision Analysis Score $\geq 89\%$, Strategic Resource Allocation Accuracy $\geq 88\%$, Risk Forecasting Precision $\geq 90\%$

The Engineering Sciences Agent (ENA) addresses challenges in mechanical, aerospace, and robotics engineering by leveraging computational models and machine learning techniques to interpret CAD schematics, predict structural integrity, and ensure system integration compliance.

The Natural Sciences Agent (NSA) focuses on earth sciences, environmental studies, and space exploration, integrating geospatial data analysis, climate modelling, and ecological system assessments. The Methodology Agent (MTA) enhances research rigor by evaluating experimental frameworks for logical consistency and reproducibility, employing Bayesian optimization models for research structure refinement. The Decision Sciences Agent (DSA) supports research strategy and interdisciplinary integration by optimizing resource allocation and risk assessment through decision-tree processing models.

INTER-AGENT COLLABORATION AND VALIDATION

A key aspect of this project is its intersection with neuroscience, particularly in understanding how human cognition processes and integrates vast amounts of information. By modelling the AI agents based on principles derived from neural networks and cognitive neuroscience, the system aims to mimic the efficiency of human knowledge navigation. Insights from neuroscience, including memory consolidation, pattern recognition, and associative learning, inform the development of algorithms that enable the agents to synthesize complex scientific knowledge dynamically. Additionally, neurophysiological studies contribute to refining adaptive learning mechanisms within the AI framework, allowing it to evolve and optimize information retrieval over time.

The system integrates real-time information sharing, hierarchical decision-making, and simulated peer-review processes to ensure accuracy, adaptability, and knowledge synthesis across disciplines. Probabilistic reasoning models assess conflicting information reliability, while graph-based knowledge representation facilitates cross-disciplinary integration. Additionally, ontology-based frameworks contextualize knowledge across domains, enhancing the overall efficiency of the AI system. The system integrates real-time information sharing, hierarchical decision-making, and simulated peer-review processes to ensure accuracy, adaptability, and knowledge synthesis across disciplines. Probabilistic reasoning models assess conflicting information reliability, while graph-based knowledge representation facilitates cross-disciplinary integration. Additionally, ontology-based frameworks contextualize knowledge across domains, enhancing the overall efficiency of the AI system.

One of the key aspects of the proposed framework is the ability of AI agents to collaborate to enhance scientific knowledge management in an interdisciplinary and context-aware manner. Each agent specializes in a specific domain, such as theoretical, experimental, engineering, or natural sciences, but the true innovation lies in their coordinated interaction. Collaboration among agents occurs through hierarchical information exchange mechanisms, resolution of conflicting data, and cross-validation of results. For example, if the Experimental Sciences Agent (ESA) detects discrepancies in experimental data compared to established theories, it can consult the Theoretical Sciences Agent (TSA) to verify theoretical consistency, while the Methodology Agent (MTA) can assess the reliability of the experimental methodology used. This synergy is facilitated by a semantic graph-based knowledge representation, enabling agents to identify conceptual relationships across different disciplines. Additionally, the integration of probabilistic reasoning models helps quantify the level of uncertainty in generated inferences, improving the system's robustness.

Finally, the use of cooperative reinforcement learning strategies allows agents to dynamically adjust their behaviour based on user feedback and evolving scientific literature. In this way, collaboration among agents not only improves the quality of information retrieval and synthesis but also fosters the discovery of innovative interdisciplinary connections that might otherwise go unnoticed.

SCALABILITY AND COMPUTATIONAL COMPLEXITY IN THE MULTI-AGENT SYSTEM

The implementation of a multi-agent system based on machine learning, probabilistic reasoning, and large-scale knowledge representation presents significant challenges in terms of scalability and computational complexity.

Managing millions of scientific documents, each containing complex information structures, requires an optimized architecture both in terms of data processing and energy efficiency. To ensure adequate performance, the system can adopt a **distributed architecture** based on cloud computing and edge computing, leveraging parallelization and distributed computing to process large volumes of data in real-time. Additionally, the use of **optimized transformer models**, such as lightweight versions of BERT or GPT specialized for the scientific domain, can reduce computational consumption without compromising the quality of natural language processing. Another strategy to improve scalability is the implementation of a **hierarchical approach to information management**, where agents progressively filter and classify data, reducing the overall computational load. Furthermore, the integration of **knowledge compression techniques**, such as pruning semantic graphs and optimizing probabilistic models, can enhance system efficiency without sacrificing inference accuracy. However, to maintain the framework's effectiveness over time, continuous performance monitoring and dynamic resource adaptation will be necessary, for example, through **automated scaling mechanisms based on machine learning**, which adjust resource allocation based on demand and query complexity. Ultimately, the system's success will depend on its ability to balance computational power and analytical accuracy, ensuring efficient and sustainable scientific knowledge management.

COMPARATIVE ANALYSIS WITH EXISTING AI SYSTEMS

The proposed multi-agent AI system for scientific knowledge management stands out due to its modularity, adaptability, and interdisciplinary collaboration. To fully appreciate its advantages, it is essential to compare it with existing AI-driven scientific knowledge management tools, such as Semantic Scholar, Elicit.org, and IBM Watson Discovery.

1. Semantic Scholar
    - Capabilities: Semantic Scholar is a widely used AI-powered academic search engine that employs natural language processing (NLP) to extract meaningful insights from research papers. It provides citation graphs, influential paper detection, and topic modelling.
    - Limitations: While effective for literature discovery, Semantic Scholar lacks domain-specific contextual reasoning and interdisciplinary synthesis capabilities. It does not dynamically integrate new research insights across fields.
    - How the Proposed System Surpasses it: The multi-agent framework enables agents specialized in different scientific disciplines to collaborate dynamically, ensuring deeper cross-domain integration and more precise contextualization of knowledge.
2. Elicit.org
    - Capabilities: Elicit.org focuses on AI-assisted literature review and evidence synthesis. It helps researchers quickly summarize findings, compare methodologies, and analyze study conclusions.
    - Limitations: Elicit.org primarily functions as a research assistant but does not incorporate probabilistic reasoning models or graph-based knowledge representation to assess conflicting information. It also lacks real-time validation mechanisms for scientific claims.
    - How the Proposed System Surpasses it: The inclusion of probabilistic reasoning models and ontology-based contextualization enhances reliability by evaluating contradictory findings and ensuring a more rigorous synthesis of information.
3. IBM Watson Discovery
    - Capabilities: IBM Watson Discovery leverages AI for enterprise-level document search and text analytics, allowing businesses to extract key insights from large textual datasets.
    - Limitations: Despite its powerful NLP capabilities, Watson Discovery is not designed specifically for scientific knowledge management. It lacks domain-specialized AI agents and mechanisms for validating experimental methodologies or integrating interdisciplinary research findings.
    - How the Proposed System Surpasses it: The proposed multi-agent system is tailored to scientific discourse, integrating domain-specific methodologies and experimental validation protocols to enhance the credibility and relevance of extracted knowledge.

Key Differentiators of the Proposed System

1. Multi-Agent Collaboration: Unlike existing systems, the proposed framework assigns specialized agents to distinct domains, enabling adaptive and cooperative knowledge synthesis.
2. Neuroscience-Inspired Architecture: Drawing insights from memory consolidation, pattern recognition, and associative learning, the AI

agents mimic human knowledge navigation, making information retrieval more intuitive and context-aware.
3. Real-Time Scientific Validation: The system employs hierarchical decision-making, peer-review simulations, and Bayesian optimization models to assess the reliability and reproducibility of scientific claims.
4. Interdisciplinary Integration: By utilizing graph-based knowledge representation, the system bridges gaps between different fields, fostering cross-disciplinary research in ways that current AI systems do not.

By surpassing existing AI-driven scientific knowledge management tools in these areas, the proposed multi-agent system redefines how researchers access, evaluate, and integrate scientific knowledge, ultimately accelerating scientific discovery and innovation.

FUTURE EXPERIMENTAL VALIDATION

Empirical validation will include prototype development with controlled simulations and comparative analysis against existing AI-driven knowledge management systems. The prototype will be tested on diverse datasets to measure its effectiveness in retrieving, organizing, and synthesizing scientific knowledge.

Additionally, controlled studies will assess the AI agents' performance in various domains to ensure robustness across interdisciplinary research. To refine AI decision-making capabilities, human-in-the-loop assessments will be integrated into the evaluation process.

Researchers and domain experts will provide feedback on AI-generated insights, identifying potential biases, inaccuracies, and contextual misinterpretations. This iterative validation approach will ensure alignment with scientific research needs and optimize agent learning over time.

Furthermore, performance benchmarks will be established using standardized evaluation metrics, such as retrieval precision, knowledge integration accuracy, and response time. These benchmarks will be compared with traditional knowledge management tools to quantify the improvements introduced by the proposed system. The validation framework will also include stress tests to examine scalability and adaptability under different research scenarios. Insights from these evaluations will inform system enhancements and future developments, ensuring a reliable and efficient AI-driven scientific knowledge management platform.

Empirical validation will include prototype development with controlled simulations and comparative analysis against existing AI-driven knowledge management systems.

Human-in-the-loop assessments will refine AI decision-making capabilities, ensuring alignment with scientific research needs.

EXPECTED APPLICATIONS AND IMPACT

The AI system will significantly enhance academic research by automating literature retrieval, synthesis, and interdisciplinary research integration, reducing literature review time and improving cross-disciplinary insights.

By leveraging AI-driven methodologies, researchers can streamline the identification of relevant studies, detect emerging trends, and integrate findings from different disciplines with minimal effort.

Additionally, the system's ability to validate and cross-reference scientific claims will contribute to higher research reliability and reproducibility.

In STEM education, the AI system will serve as an interactive learning assistant, providing personalized knowledge retrieval and adaptive learning experiences. It will facilitate dynamic curriculum recommendations, personalized study plans, and intelligent tutoring capabilities tailored to individual student needs. The integration of real-time learning analytics will enable educators to assess student progress more effectively and adjust teaching strategies accordingly.

For industry and applied research, the AI-driven framework will accelerate R&D in pharmaceuticals, engineering, and environmental sciences, optimizing innovation processes. The system will support researchers in formulating hypotheses, designing experiments, and analyzing results with increased precision and efficiency.

In pharmaceuticals, AI-driven drug discovery can expedite the identification of promising compounds, while in engineering, it can enhance simulations and predictive modelling for better system designs.

In policymaking, the system will support evidence-based decision-making by providing real-time scientific insights for addressing global challenges such as climate change and public health crises. AI-driven analysis of policy implications, risk assessments, and scenario modelling will empower policymakers with reliable, data-driven recommendations.

Additionally, the system will facilitate international collaboration by synthesizing research from different geopolitical regions, ensuring that policy decisions are informed by diverse perspectives and the latest scientific advancements. The AI system will significantly enhance academic research by automating literature retrieval, synthesis, and interdisciplinary research integration, reducing literature review time and improving cross-disciplinary insights.

In STEM education, it will serve as an interactive learning assistant, providing personalized knowledge retrieval and adaptive learning experiences. For industry and applied research, the AI-driven framework will accelerate R&D in pharmaceuticals, engineering, and environmental sciences, optimizing innovation processes. In policymaking, the system will support evidence-based decision-making by providing real-time scientific insights for addressing global challenges such as climate change and public health crises.

CONCLUSION AND FUTURE WORK

This study presents a multi-agent AI architecture designed to enhance scientific knowledge management through contextual adaptability and interdisciplinary integration. The framework facilitates automated knowledge retrieval, structured synthesis, and collaborative reasoning across diverse research domains, significantly improving efficiency and accuracy in scientific exploration.

Future research will focus on refining the prototype to enhance its scalability and adaptability in real-world applications. Empirical validation will be expanded to include longitudinal studies assessing the long-term impact of AI-driven knowledge navigation on research productivity and interdisciplinary collaboration. Additionally, further development will incorporate reinforcement learning techniques to enable continuous improvement in AI agent performance and decision-making capabilities.

Another critical aspect of future work will be the seamless integration of real-time research data streams, ensuring that AI agents remain updated with the latest advancements across disciplines. Ethical considerations, including algorithmic transparency, data privacy, and bias mitigation, will also be prioritized to foster trust and responsible AI deployment in scientific knowledge management.

By refining AI-driven multi-agent systems, this research establishes a robust foundation for AI-enhanced knowledge navigation, fostering innovation, accelerating scientific progress, and shaping the future of knowledge discovery across multiple disciplines. This study presents a multi-agent AI architecture designed to enhance scientific knowledge management through contextual adaptability and interdisciplinary integration. Future research will focus on prototype implementation, empirical validation, reinforcement learning enhancements, real-time research data integration, and ethical considerations.

By refining AI-driven multi-agent systems, this research establishes a foundation for AI-enhanced knowledge navigation, fostering innovation and accelerating scientific progress across multiple disciplines.

REFERENCES

Alier, M., Pereira, J., García-Peñalvo, F. J., Casañ, M. J., & Cabré, J. (2025). LAMB: An open-source software framework to create artificial intelligence assistants deployed and integrated into learning management systems. Computer Standards & Interfaces, 92, 103940. https://doi.org/10.1016/j.csi.2024.103940

Baillifard, A., Gabella, M., Lavenex, P. B., & Martarelli, C. S. (2025). Effective learning with a personal AI tutor: A case study. Education and Information Technologies, 30(1), 297–312. https://doi.org/10.1007/s10639-024-12888-5

Bettin, D., Maurer, T., Schlatt, F., & Bettin, S. (2022). The scientific publication score - a new tool for summarizing evidence and data quality criteria of biomedical publications. Journal of Bone and Joint Infection, 7(6), 269–278. https://doi.org/10.5194/jbji-7-269-2022

Eppler, M. J., & Mengis, J. (2004). The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines. *The Information Society*, *20*(5), 325–344. https://doi.org/10.1080/01972240490507974

Li, H., Nourkhiz Mahjoub, H., Chalaki, B., Tadiparthi, V., Lee, K., Moradi Pari, E., Lewis, C., & Sycara, K. (2025). Language Grounded Multi-agent Reinforcement Learning with Human-interpretable Communication. Advances in Neural Information Processing Systems, 37, 87908–87933.

Navigating the Future: Large Language Models (LLM) and the Imperative of Responsible Artificial Intelligence (RAI). (s.d.). SpringerLink. Recuperato 26 febbraio 2025, da https://link.springer.com/collections/egiiehbedc.

Tran, K.-T., Dao, D., Nguyen, M.-D., Pham, Q.-V., O'Sullivan, B., & Nguyen, H. D. (2025). Multi-Agent Collaboration Mechanisms: A Survey of LLMs (arXiv:2501.06322). arXiv. https://doi.org/10.48550/arXiv.2501.06322

Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., … Gui, T. (2025). The rise and potential of large language model based agents: A survey. Science China Information Sciences, 68(2), 121101. https://doi.org/10.1007/s11432-024-4222-0

Zignego, M., Bertirotti, A., Gemelli, P., & Pagani, L. (2023, gennaio 12). Neurodesign: A Game-Changer in Educational Contexts. LearnXDesign 2023. LearnXDesign 2023. https://doi.org/10.21606/drslxd.2024.014