A Novel Idea Generation Path Recognition Framework Integrating Graph Reasoning and Multi-Agent Collaboration

Authors: Liang Guoqiang, Lin Gege, Zhang Zhihao, Zhang Shuo, Zhang Shuo

Date: 2025-10-13T00:00:00+00:00

Abstract

How new ideas emerge has long been a topic of human concern. Against the backdrop of the transformation from traditional scientific discovery logic to the "double helix" driven logic of AI4Science and Science4AI, utilizing AI Agents to integrate massive, heterogeneous scientific knowledge and to discover and simulate the generation process of human novel ideas holds significant importance in evolving scientific research discovery from a starting point of scientists' experience and intuition to one based on big data and AI algorithms, enabling scientists to shift from laborious information screening to higher-order creative thinking, and actively breaking disciplinary barriers to foster a large number of interdisciplinary scientific discoveries. This paper proposes a framework that integrates graph reasoning with multi-agent collaboration (Graph-reasoning And Multi-agent Pathfinding, GAMP). The framework first extracts triples from paper abstracts through prompt engineering and employs Neo4j for storage, forming a large-scale scientific knowledge graph as a structured knowledge base. Then, it designs a collaborative system comprising multiple AI Agents with different functions (such as domain expert Agent, path exploration Agent, innovation evaluation Agent, etc.), with each Agent driven by large models and endowed with capabilities in semantic understanding, entity extraction, and path finding. For instance, in the domain expert Agent, through knowledge bases and prompt engineering, the Agent focuses on rationality assessment at the gene, protein, and signaling pathway levels; in the path exploration Agent, different path search methods such as breadth-first algorithms, genetic algorithms, and large model-guided search are employed, enabling the Agent to focus on discovering the most novel paths. These Agents conduct collaborative exploration and reasoning on the graph, generating and evaluating generation paths for new ideas by simulating the "brainstorming" and "hypothesis-verification" cycles of scientific teams. Taking the achievement awarded the 2021 Nobel Prize in Physiology or Medicine as an example, we collected literature related to temperature and tactile receptors from the Web of Science Core Collection database, Scopus, and PubMed between January 1, 1995, and December 31, 2005, as a case study, and constructed a three-layer "problem-solution-effect" knowledge network for empirical research. The innovations of this paper are: first, connecting symbolic and connectionist approaches to fully leverage graph-structured reasoning and the powerful semantic understanding capabilities of large models; second, designing a structured multi-agent collaboration protocol with clear division of labor to simulate real scientific research teams. The limitations lie in: the formal representation of "paths for new ideas," deep semantic understanding, and the evaluation of breakthrough potential of new ideas require further deepening and refinement.

Full Text

A Novel Idea Generation Path Identification Framework Integrating Graph Reasoning and Multi-Agent Collaboration

Liang Guoqiang, Lin Gege, Zhang Zhihao, Zhang Shuo
(School of Economics and Management, Beijing University of Technology, Beijing, 100124)

How new ideas emerge has long been a central question in human inquiry. Against the backdrop of a transformative shift from traditional scientific discovery logic to a "double helix" driven by AI4Science and Science4AI, leveraging AI Agents to integrate massive, heterogeneous scientific knowledge and simulate the generation of human novel ideas holds profound significance. This approach promises to shift scientists from laborious information screening to higher-order creative thinking, actively breaking disciplinary barriers and catalyzing interdisciplinary discoveries. This paper proposes a framework that fuses graph reasoning with multi-Agent collaboration (GAMP). The framework first extracts triples from paper abstracts through prompt engineering and stores them using Neo4j to form a large-scale scientific knowledge graph as a structured knowledge substrate. It then designs a collaborative system comprising multiple AI Agents with distinct functions—such as domain expert Agents, path exploration Agents, and innovation evaluation Agents—each powered by large language models and endowed with capabilities in semantic understanding, entity extraction, and pathfinding. For instance, domain expert Agents, through knowledge bases and prompt engineering, focus on assessing the rationality of genes, proteins, and signaling pathways; path exploration Agents employ breadth-first algorithms, genetic algorithms, and LLM-guided search to discover novel paths. These Agents conduct collaborative exploration and reasoning on the graph, simulating scientific teams' "brainstorming" and "hypothesis-verification" cycles to generate and evaluate paths for novel idea emergence. Using the 2021 Nobel Prize-winning discovery in Physiology or Medicine as a case study, we collected literature related to temperature and tactile receptors from the Web of Science Core Collection, Scopus, and PubMed databases between January 1, 1995, and December 31, 2005, constructing a three-layer "problem-solution-effect" knowledge network for empirical validation. The innovations are twofold: first, bridging symbolism and connectionism by harnessing both graph-structured reasoning and the powerful semantic understanding of large models; second, designing a structured multi-Agent collaboration protocol with clear role division that mirrors real research teams. Limitations include the need for deeper refinement in formally representing "paths of novel ideas," achieving deep semantic understanding, and evaluating the breakthrough potential of new ideas.

Keywords: Scientific Discovery; Graph Reasoning; Large Language Models; AI Agent
Classification Code:
Funding: Beijing Natural Science Foundation General Project "Research on the Emergence Mechanism and Identification Methods of Research Frontiers from a Co-evolution Perspective" (Project No. 9232002); National Natural Science Foundation Youth Projects "Research on Large Model-Empowered Personalized Trading Recommendation Methods for High-Value Patents and Applications" (Project No. 72404020) and "Identification of Disruptive Low-Carbon Technologies and Dynamic Selection of Innovation Paths" (Project No. 72304023).

Scientific breakthroughs and the generation of novel ideas have long depended on individual scientists' intuition and experience, as well as collaboration within research teams—a process often characterized by high costs, long cycles, and serendipity. As scientific knowledge enters an explosive growth phase and disciplinary barriers become increasingly rigid, traditional literature review and brainstorming models struggle to comprehensively grasp cross-domain knowledge associations, potentially missing major innovation opportunities. This challenge concerns not only science and technology themselves but also imposes higher demands on research management efficiency and resource allocation optimization.

In recent years, artificial intelligence—particularly large language models (LLM) and knowledge graph (KG) technologies—has offered new possibilities for addressing this challenge. AI4Science aims to leverage AI to solve core problems in scientific discovery, while Science4AI focuses on how scientific practice can 反哺 (feed back into) AI theory, forming a double helix of synergistic development. Against this backdrop, this paper explores an intermediate path: constructing a computational framework that simulates the cognitive collaboration process of interdisciplinary research teams to computationally identify and evaluate the generation paths of scientific novel ideas.

This paper proposes a novel idea generation path identification framework integrating graph reasoning and multi-Agent collaboration (GAMP). The core concept of GAMP is to deeply fuse symbolism (structured knowledge represented by knowledge graphs) with connectionism (semantic understanding represented by LLMs). Through a role-defined, orderly multi-Agent system, it conducts directed, heuristic exploration on a vast scientific knowledge graph to automatically generate, evaluate, and screen potentially breakthrough scientific hypothesis paths.

2 Literature Review

The proposed GAMP framework stands at the intersection of multiple rapidly evolving research domains. To clearly position this study' s contributions, this section systematically reviews the state-of-the-art in scientific knowledge graph construction and application, graph reasoning algorithms, LLM applications in science, and multi-Agent systems, while deeply analyzing existing limitations to provide theoretical justification for the necessity and innovation of the GAMP framework.

2.1 Construction and Application of Scientific Knowledge Graphs (SKG)

Scientific knowledge graphs serve as carriers of structured scientific knowledge and constitute core infrastructure for computational scientific discovery. Traditional SKGs are built by extracting entities (e.g., concepts, methods, materials) and relationships (e.g., "used for," "inhibits," "causes") from large-scale scientific literature (papers, patents). In recent years, construction methods have evolved from predefined template matching to deep semantic understanding and extraction using large language models. For example, Shi et al. utilized event knowledge graph techniques and LLMs to construct a large-scale scientific experimental knowledge graph for organic solar cells containing tens of thousands of nodes and relationships, effectively supporting experimental protocol recommendation and evolutionary analysis. At the application level, SKGs have become important tools for tech intelligence analysis, key technology identification, and disciplinary knowledge evolution analysis. Cao et al. improved the PageRank algorithm by constructing a "science-technology" knowledge topic complex network, enabling fine-grained identification of key core technologies in the CNC machine tool domain.

2.2 Advances in Graph Reasoning Algorithms for Path Discovery

Graph reasoning algorithms aim to mine potentially meaningful paths from knowledge graphs, representing the core technology for identifying scientific breakthrough paths. Early methods primarily relied on random walk or meta-path-based graph traversal algorithms, which were efficient but heavily dependent on predefined path patterns and lacked flexibility. Subsequently, knowledge graph embedding (KGE) methods mapped entities and relationships into low-dimensional vector spaces for link prediction through vector operations, but these methods suffered from poor interpretability and could not generate clear paths. In recent years, interpretable path-based reasoning has become a research hotspot. For instance, the KGExplainer framework provides verifiable explanations for knowledge graph completion predictions by exploring multiple collaborative reasoning paths, demonstrating advantages in biomedicine. Graph reinforcement learning (GRL) combines graph neural networks with reinforcement learning, enabling agents to learn exploration strategies on graph structures to discover optimal paths, offering a new paradigm for handling scientific knowledge associations in non-Euclidean spaces.

2.3 Application and Adaptation of Large Language Models (LLM) in Scientific Research

LLMs have revolutionized scientific knowledge processing with their powerful natural language understanding and generation capabilities. Domain-specific models (e.g., HuatuoGPT, BenTsao) demonstrate reliability in multi-round medical dialogue and diagnostic assistance through instruction fine-tuning. In terms of reasoning paradigms, Chain-of-Thought (CoT) and Retrieval-Augmented Generation (RAG) are widely used to enhance LLM logicality and factual accuracy. Particularly noteworthy is the emerging paradigm of iterative reasoning through LLM-KG interaction. For example, the DoG (Debate on Graph) framework reduces long-path interference by introducing multi-role LLM teams (e.g., problem simplification experts, commentators) for iterative debate and reasoning on KGs. The FiDeLiS framework combines Path-RAG and Deductive Verification Beam Search (DVBS) to simultaneously improve factual accuracy and efficiency in question answering.

2.4 Collaboration Paradigms and Efficiency Optimization in Multi-Agent Systems

Multi-Agent systems provide a distributed approach to solving complex problems through division of labor among multiple agents, which has been revitalized through recent integration with LLMs. Early multi-Agent systems focused on designing communication protocols and collaboration mechanisms. Today, LLM-driven agents have become a research hotspot. The Multi-Agent Debate (MAD) paradigm enables multiple LLM agents to collaborate through "roundtable debate," effectively improving decision quality but suffering from high computational overhead and latency. To enhance efficiency, new collaboration paradigms have been proposed. The MARS (Multi-Agent Review System) framework mimics the academic review process with role divisions of "author-reviewer-meta-reviewer," reducing token consumption and inference time by approximately 50% while maintaining reasoning quality by minimizing frequent inter-agent communication. Furthermore, the Federation of Agents (FoA) framework proposes a semantic-aware communication architecture that enables dynamic capability matching and task decomposition through versioned capability vectors (VCVs), laying the foundation for collaborative work in large-scale heterogeneous agent federations.

In summary, while significant progress has been made in each domain, limitations persist. SKGs provide structured knowledge foundations but lack semantic understanding and flexible reasoning capabilities. Graph reasoning algorithms excel at discovering structural patterns in graphs but lack deep semantic understanding. LLMs possess powerful semantic understanding and generation abilities but cannot guarantee factuality or conduct structured exploration. Multi-Agent systems provide collaboration paradigms for complex problem-solving, but general architectures cannot be directly applied to scientific discovery scenarios requiring high rigor. A framework that deeply integrates graph reasoning, LLMs, and multi-Agent collaboration to simulate real research teams for identifying scientific breakthrough paths remains in its early stages.

3.1 Overall Framework

The GAMP framework aims to automate the identification of promising scientific breakthrough paths by simulating a virtual, highly specialized interdisciplinary research team conducting collaborative exploration and reasoning on a structured scientific knowledge graph. The framework comprises three core layers: the Data Layer, Knowledge Layer, and Agent Collaboration Layer, which interact through clearly defined interfaces to complete the entire process from raw data to innovative path output (see Figure 1

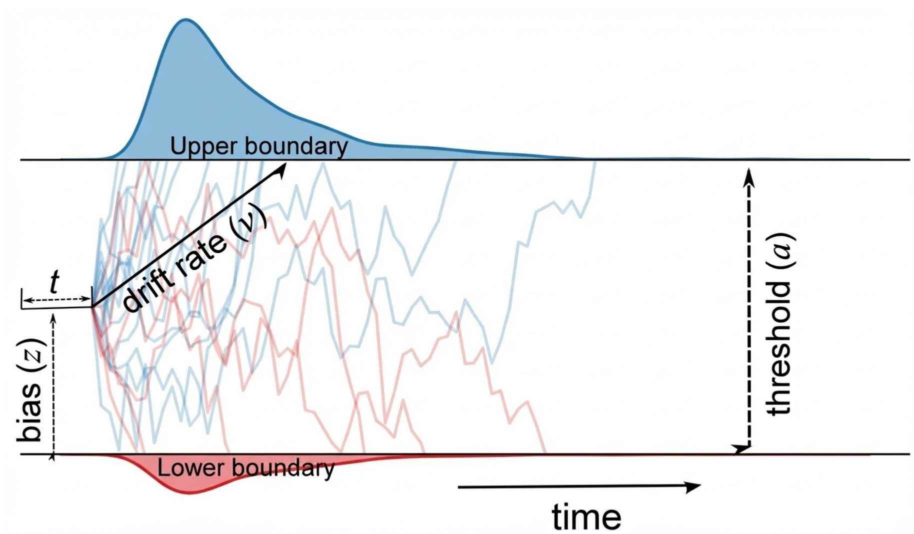

Figure 1: Figure 1

).

Figure 1 GAMP Overall Framework

Data Layer: As the foundation, this layer integrates multi-source heterogeneous scientific data, including academic literature databases (e.g., Web of Science, PubMed), patent databases, and specialized domain databases. It performs data collection, cleaning, and preprocessing to provide raw material for knowledge construction.

Knowledge Layer: The cornerstone of the framework, its core is the scientific knowledge graph. This paper employs an innovative "problem-solution-effect" three-layer semantic model to represent scientific knowledge structurally.

This layer transforms unstructured text into machine-understandable, machine-reasonable semantic networks, providing a single, trustworthy source of facts for upper-layer Agent reasoning.

Agent Collaboration Layer: The "brain" and engine of the framework, consisting of a multi-Agent system where each Agent is driven by a large language model and assigned specific roles and tasks. Agents communicate and collaborate asynchronously through a shared workspace, simulating the "hypothesis-proposing, peer-review, refinement" cycle of real research teams.

3.2 Construction of the "Problem-Solution-Effect" Three-Layer Scientific Knowledge Graph

To precisely characterize the intrinsic logic of scientific discovery, we constructed a "problem-solution-effect" three-layer scientific knowledge graph. This model transcends simple entity-relationship extraction to capture the complete thought chain of "posing questions—designing solutions—verifying effects" in scientific research.

Problem Layer: Nodes represent core scientific questions or challenges that research seeks to address (e.g., "How to identify molecular receptors for noxious heat stimuli?").

Solution Layer: Nodes represent specific methods, techniques, compounds, tools, or theories used to solve problems (e.g., "capsaicin," "gene knockout technology," "calcium imaging").

Effect Layer: Nodes represent results, discoveries, biological functions, or performance metrics produced after implementing solutions (e.g., "activates TRPV1 ion channels," "causes intracellular calcium concentration increase," "produces thermal pain behavioral responses").

When extracting entities using large models, the prompt is: "You are a scientific knowledge engineer. Please precisely identify the [research problem], [core methods or substances used], and [most critical research findings or effects] from the following paper abstract. Ensure extracted content comes directly from the text without speculation. Output format: JSON: {" problem" : "“," solution" : "“," effect" : "“}." Extracted entities undergo normalization (eliminating naming ambiguities, e.g., unifying "VR1" and "TRPV1") and linking. The cleaned triples are then stored in Neo4j. Intra-layer and inter-layer connections use rich relationship types, such as "studied through⋯" between problem-solution and "leads to," "inhibits," "enhances" between solution-effect. Figure 2 [FIGURE:2] shows a network diagram after large model extraction, where L1, L2, and L3 correspond to the problem, solution, and effect layers, respectively.

Figure 2 Entity Relationship Diagram Across Different Periods

3.3 Detailed Design of the Multi-Agent System

The core of the GAMP framework lies in its multi-Agent collaboration system. We defined clear roles, responsibilities, and decision-making mechanisms for each Agent, collectively simulating an efficient virtual research team.

Agent Roles and Function Definitions:

Chief Scientist Agent: Acts as team leader and coordinator. Responsible for receiving user queries, decomposing complex problems into subtasks, assigning tasks to other Agents, and making final decisions and path rankings after synthesizing input from all parties.

Domain Expert Agents (Multiple): Each Agent represents a specific discipline (e.g., Molecular Biologist Agent, Physiologist Agent, Chemist Agent). Their core responsibility is to evaluate the scientific rationality and logical coherence of each step in a path from their disciplinary perspective. They undergo role solidification through specific instructions; for example, the Molecular Biologist Agent' s instructions emphasize deep understanding of genes, proteins, and signaling pathways.

Path Exploration Agent: Responsible for active exploration on the SKG. It combines traditional graph algorithms (e.g., breadth-first search for discovering direct associations) with LLM semantic guidance (e.g., LLM predicting "Which ion channels might be functionally complementary to TRPV1?") to escape local optima and discover non-obvious associations.

Innovation Evaluation Agent: Focuses on assessing the breakthrough potential of paths. It scores the novelty and potential impact of paths based on predefined quantitative metrics (e.g., path topological novelty, semantic rarity) and deep LLM semantic understanding.

Fact-Checking Agent: The "gatekeeper" of system reliability. Its task is to ensure all generated inferences and hypotheses can find evidential support in the SKG, strictly suppressing potential LLM "hallucinations" and enhancing overall system credibility.

4 Core Algorithms and Implementation

Each Agent' s decision-making core is a carefully engineered prompt template that solidifies the agent' s role, task, knowledge background, and behavioral constraints to ensure consistent and professional behavior. Below, the Molecular Biologist Agent is used as an example to illustrate its decision prompt template (see Figure 3 [FIGURE:3]). The Path Exploration Agent' s core algorithms primarily draw from existing methods such as breadth-first search and ant colony algorithms, which will not be elaborated here.

Figure 3 Domain Expert Agent Prompt Template

Novelty evaluation primarily employs the formula:

$$Novelty(P) = \frac{1}{1 + \log(freq(P))}$$

where $freq(P)$ is the frequency of path P or its subpaths appearing in historical literature; lower frequency yields higher novelty.

In 2021, research on temperature and tactile receptors was awarded the Nobel Prize in Physiology or Medicine, revealing the neural signal transduction pathways and mechanisms for human temperature sensing, pain, and touch. Receptors within human cells can sensitively perceive high (heat) or low (cold) temperature stimuli from the environment. This temperature sensing mechanism, along with mechanical force-induced responses in touch, is closely related to pain formation, providing new targets for pain therapeutic strategies. This breakthrough offers an excellent opportunity to validate the methodological framework discussed earlier. Its mature development trajectory provides a rich practical foundation for verifying the framework' s rationality.

Considering data authority and completeness, this study selected the Web of Science Core Collection, Scopus, and PubMed databases as data sources. Search terms were based on common English expressions for "temperature" and "tactile receptors" in SCI papers, supplemented by MeSH thesaurus terms for broader and narrower concepts. Potentially ambiguous domains were retained to avoid omissions. The search time span was January 1, 1995, to December 31, 2005, with the retrieval date being March 28, 2024, yielding 3,234 papers. After format conversion, deduplication, and retaining only publication year and abstract, 3,107 valid abstract records were obtained.

Feeding the abstract data into the GAMP framework revealed multiple paths for novel idea generation in this field from 1995 to 2003 (see Table 1).

Table 1 Top 5 New Idea Generation Paths Based on GAMP Framework

Path	Key Nodes	Historical Truth?	Novelty Score	Interpretation
【Problem】Heat pain mechanism → 【Solution】Capsaicin → 【Effect】Activates TRPV1 ion channels	TRPV1 identified as heat pain receptor	Yes (Hit)	0.92	Historically validated
【Problem】Cold sensation mechanism → 【Solution】Menthol → 【Effect】Activates TRPM8 ion channels	TRPM8 identified as cold receptor	Yes (Important discovery)	-	-

Path	Key Nodes	Historical Truth?	Novelty Score	Interpretation
【Problem】Heat hyperalgesia → 【Solution】Inflammatory factors → 【Effect】Enhances TRPV1 function	Explains inflammatory heat pain	No (Forward-looking)	-	-

Path	Key Nodes	Historical Truth?	Novelty Score	Interpretation
【Problem】Nociceptive stimulus gating → 【Solution】Capsaicin analogs → 【Effect】TRPV1 isoforms discovered	Predicts TRPV1 functional diversity	No (Forward-looking)	-	-

Path	Key Nodes	Historical Truth?	Novelty Score	Interpretation
【Problem】Heat sensation → 【Solution】Capsaicin resistance study → 【Effect】Potential novel heat receptor discovered	Suggests existence of other heat receptors	No (Forward-looking)	-	-

Experimental results demonstrate that the GAMP framework can not only effectively trace historically significant scientific breakthrough paths, showing excellent identification accuracy (high hit rate and ranking), but more importantly, it can generate highly heuristic and forward-looking research hypotheses based on historical knowledge states. However, this study has not yet conducted ablation experiments, and many detailed issues require deeper investigation. For instance, framework performance is highly dependent on underlying SKG quality, as biases or omissions in historical data directly affect results. The framework is more adept at combinatorial innovation within existing knowledge systems, while its ability to identify ideas that completely overturn existing cognition requires further validation. Additionally, the evaluation metrics for novel ideas in this paper adopt only a single formula, which is relatively crude, leaving substantial room for refinement.

In conclusion, this paper aims to provide peers with a heuristic framework for identifying novel idea generation paths that can significantly accelerate scientific discovery. Due to time constraints, only partial results have been compiled and reported for academic exchange.

[1] Lu, C., Lu, C., Lange, R.T., Foerster, J.N., Clune, J., & Ha, D. (2024). The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery. ArXiv, abs/2408.06292.
[2] Kitano, H. (2021). Nobel Turing Challenge: creating the engine for scientific discovery. NPJ Systems Biology and Applications, 7.
[3] Tang, J., Xia, L., Li, Z., & Huang, C. (2025). AI-Researcher: Autonomous Scientific Innovation. ArXiv, abs/2505.18705.
[4] Hu, Mingxue et al. "A Survey of Scientific Large Language Models: From Data Foundations to Agent Frontiers." ArXiv abs/2508.21148 (2025): n. pag.
[5] King, R.D. (2012). Opportunities for Automated Workflow and Planning with Robot Scientists.
[6] Ghafarollahi, A., & Buehler, M.J. (2024). SciAgents: Automating scientific discovery through multi-agent intelligent graph reasoning. ArXiv, abs/2409.05556.
[7] Ghareeb, A., Chang, B., Mitchener, L., Yiu, A., Szostkiewicz, C.J., Laurent, J.M., Razzak, M.T., White, A.D., Hinks, M.M., & Rodriques, S.G. (2025). Robin: A multi-agent system for automating scientific discovery. ArXiv, abs/2505.13400.
[8] Savage, N. (2012). Automating scientific discovery. Communications of the ACM, 55, 9-11.
[9] Gower, A.H., Korovin, K., Brunnsåker, D., Kronström, F., Reder, G.K., Tiukova, I.A., Reiserer, R.S., Wikswo, J.P., & King, R.D. (2024). The Use of AI-Robotic Systems for Scientific Discovery. ArXiv, abs/2406.17835.