The Economy of Reasoning: Incentivizing Epistemic Diversity in Decentralized Scientific Swarms

Tri Minh Nguyen [1] Sherif Abdulkader Tawfik [1]
Truyen tran [1]
[1]*Deakin University*. Correspondence to: Tri Minh Nguyen tri.nguyen1@deakin.edu.au.

1. The Problem: Centralized Bottlenecks in AI-for-Science

This vision paper proposes a novel architectural paradigm for AI-driven scientific discovery. Rather than presenting empirical benchmarks, we articulate a conceptual framework designed to provoke discussion on the future of decentralized, self-organizing research systems.

Scientific discovery is fundamentally a distributed, stochastic enterprise, yet contemporary AI-for-science architectures remain tethered to centralized orchestration. Systems like AutoGen [1], MetaGPT [2], and domain-specific platforms such as AI-CoScientist [3] rely on a central "brain" to decompose problems and delegate subtasks. While effective for well-defined workflows, this topology is fundamentally misaligned with the non-linear nature of scientific inquiry.

We identify three critical bottlenecks that necessitate a paradigm shift:

Context Saturation. The central orchestrator acts as an informational choke point. As complexity scales, the LLM's context window saturates, forcing lossy compression that discards the nuance critical for insight.

Scalability Ceiling. Adding worker agents yields diminishing returns; the orchestrator's fixed cognitive capacity becomes the rate-limiting step, preventing linear throughput gains.

Central Orchestrator Bias. The discovery trajectory is bounded by the planner's initial decomposition. This architectural rigidity precludes serendipitous discovery outside the planner's horizon in which a bias or hallucination at the central node propagates system-wide failure.

2. The Vision: Scientific Discovery Swarms

We propose the **Scientific Discovery Swarm (SDS)**: a decentralized paradigm grounded in Swarm Intelligence that shifts agency from a central planner to self-organizing collectives. The architecture is defined by three principles:

Stigmergic Coordination. Agents coordinate by modifying a shared Knowledge Graph rather than exchanging direct messages. An agent posts a "Hypothesis Node"; available validators sense this state change and initiate critique. This decouples agents entirely, enabling asynchronous, non-blocking parallelism where agents operate at distinct temporal scales.

Emergent Workflow Dynamics. The discovery process is not pre-scripted but emerges dynamically. A reasoning agent generates a hypothesis, triggering a validator to critique it, which subsequently prompts an experiment agent to test it. This chain self-organizes based on information availability, mirroring the web of conjectures and refutations that characterizes real scientific progress.

Horizontal Scalability. While centralized models face logarithmic scaling due to coordination overhead, the swarm model approaches linear scaling. System capacity is limited only by environmental I/O, technically addressable via database sharding, rather than the cognitive context of any single model.

3. Architecture: A Tri-Layered Framework

The SDS is formally defined as $\mathcal{S} = \langle \mathcal{G}, \mathcal{A}, \mathcal{M} \rangle$: a shared Knowledge Graph $\mathcal{G}$, an agent swarm $\mathcal{A}$, and an economic protocol $\mathcal{M}$.

Layer 1: The Stigmergic Environment. A dynamic Knowledge Graph $G = (V, E)$ serves as the shared "pheromone field." Vertices encode Concepts (ontological anchors), Data (immutable experimental ground truth), Hypotheses (probabilistic propositions with confidence scores and economic stakes), and Operational state (agent identities, task manifests). Edges define semantic, epistemic, and transactional relationships.

Layer 2: The Agent Swarm. Specialized agents operate via a dual-system cognitive architecture: System 1 (intuitive LLM-based heuristics) and System 2 (reflexive verification via Chain-of-Verification [4]). Agent types include: *Explorers* (inductive/analogical pattern detection), *Hypothesis Agents* (abductive reasoning for anomaly explanation), *Experiment Agents* (translation to executable predictions), *Validators* (adversarial falsification), *Resource Agents* (infrastructure abstraction), and *Enforcement Agents* (integrity maintenance).

Layer 3: The Economic Protocol. This layer addresses the core challenge: *how do we allocate scarce computational resources without central planning?*

4. The Novel Contribution: An Internal Market for Epistemic Value

The SDS implements an internal market governed by a **RESEARCH utility token**. This economic protocol establishes a closed-loop incentive structure: agents accumulate tokens by successfully verifying novel hypotheses, which subsequently function as liquidity to procure access to high-fidelity verification tools (e.g., DFT simulations, cloud laboratories).

Novelty Premium Mechanism. The protocol

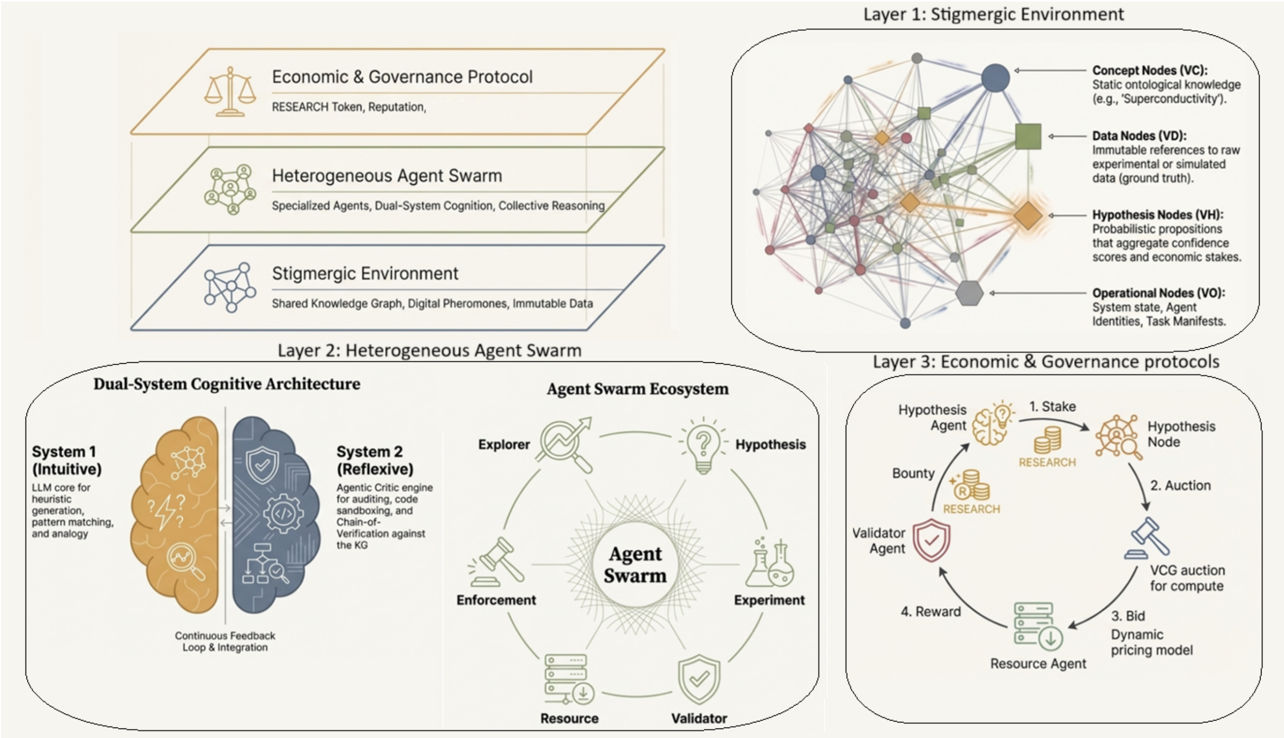

Fig. 1: Overview of the Scientific Discovery Swarm (SDS) Architecture. The framework is organized into a tri-layered structure: Stigmergic environment, Agent swarm, and Economic & Governance protocols

prices hypotheses inversely proportional to the knowledge density of their semantic neighborhood. Hypotheses bridging disparate clusters "structural holes" receive multiplied bounties. As confidence approaches certainty, rewards asymptotically decay, forcing the swarm to migrate resources toward the epistemic frontier rather than converging on local minima.

Cost of Inquiry as Spam Filter. Rational agents propose hypothesis H only if expected utility exceeds zero:

$$E[U] = P(H) \cdot B - (1 - P(H)) \cdot S_{\text{collateral}} \tag{1}$$

where stake $S_{\text{collateral}} = f(R^{-1})$ scales inversely with reputation. Hallucination-prone agents face exponentially higher collateral requirements, effectively self-selecting out of the network.

Resource Allocation via VCG Auctions. Finite computational resources are arbitrated via Vickrey-Clarke-Groves auctions [5]. Agents bid based on expected epistemic utility, and a discovery prediction market allows trading binary shares on hypothesis validity, creating a "computational invisible hand" that directs attention toward credible discoveries.

5. Scientific Rigor by Design

Two protocols enforce methodological integrity:

Falsifiability Protocol. Following Popperian demarcation [6], hypothesis nodes are rejected unless accompanied by a Falsification Contract: $\exists E : \text{Result}(E) \to \neg H$. Smart contracts enforce this condition programmatically.

Generative Adversarial Collaboration. Competing theories trigger structured debates where proponent agents must agree on a "crucial experiment" design, arbitrated by neutral validators. Results are binding, resolving conflicts and pruning the Knowledge Graph.

6. Conclusion: A Call for Architectural Rethinking

This vision paper argues that the future of AI-for-Science demands a fundamental architectural rethinking, not incremental improvements to centralized pipelines, but a paradigm shift toward self-organizing ecosystems. The Scientific Discovery Swarm framework we propose enables: (1) linear horizontal scaling decoupled from central cognitive limits, (2) serendipitous discovery through emergent collective intelligence, and (3) autonomous resource allocation via market mechanisms rather than central planning.

We acknowledge that realizing this vision requires substantial future work: empirical validation of stigmergic coordination efficacy, calibration of economic incentive parameters, and integration with real-world scientific infrastructure. We present this framework not as a finished system, but as a provocation, an invitation to the community to explore whether the decentralized, bottom-up architectures that transformed internet infrastructure might similarly transform how machines conduct science.

We posit that the future of scientific AI lies not in ever-larger monolithic models, but in decentralized swarms where inductive, abductive, and analogical reasoning emerge.

References

[1] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Awadallah, Ryen W White, Doug Burger, and Chi Wang. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. 2023.

[2] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework. *12th International Conference on Learning Representations, ICLR 2024*, 8 2023.

[3] Juraj Gottweis et al. Towards an AI co-scientist. 2 2025.

[4] Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. Chain-of-Verification Reduces Hallucination in Large Language Models. *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 3563–3578, 2024.

[5] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8, 3 1961.

[6] David B Resnik. A pragmatic approach to the demarcation problem. *Studies in History and Philosophy of Science Part A*, 31(2):249–267, 6 2000.

[7] G. Beni. The concept of cellular robotic system. pages 57–62, 1 1988.

[8] Guy Theraulaz and Eric Bonabeau. A brief history of stigmergy. *Artificial life*, 5(2):97–116, 1999.

[9] Marco Dorigo and Luca Maria Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1):53–66, 1997.

[10] Mohammad Reza Bonyadi and Zbigniew Michalewicz. Particle Swarm Optimization for Single Objective Continuous Space Problems: A Review. *Evolutionary Computation*.

[11] Johannes Schneider. Generative to Agentic AI: Survey, Conceptualization, and Challenges. 2025.

[12] Jiaqi Wei, Yuejin Yang, Xiang Zhang, Yuhan Chen, Xiang Zhuang, Zhangyang Gao, Dongzhan Zhou, Guangshuai Wang, Zhiqiang Gao, Juntai Cao, Zijie Qiu, Ming Hu, Chenglong Ma, Shixiang Tang, Junjun He, Chunfeng Song, Xuming He, Qiang Zhang, Chenyu You, Shuangjia Zheng, Ning Ding, Wanli Ouyang, Nanqing Dong, Yu Cheng, Siqi Sun, Lei Bai, and Bowen Zhou. From AI for Science to Agentic Science: A Survey on Autonomous Scientific Discovery. 2025.

[13] Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. ChatDev: Communicative Agents for Software Development. *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 1:15174–15186, 7 2023.

[14] Chi Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate. *12th International Conference on Learning Representations, ICLR 2024*, 8 2023.

[15] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. *EMNLP 2024 - 2024 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 17889–17904, 5 2023.

[16] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors. In *ICLR*, 2024.

[17] Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen Schmidhuber. GPTSwarm: Language Agents as Optimizable Graphs. 2024.

[18] Dawei Gao, Zitao Li, Xuchen Pan, Weirui Kuang, Zhijian Ma, Bingchen Qian, Fei Wei, Wenhao Zhang, Yuexiang Xie, Daoyuan Chen, Liuyi Yao, Hongyi Peng, Zeyu Zhang, Lin Zhu, Chen Cheng, Hongzhu Shi, Yaliang Li, Bolin Ding, and Jingren Zhou. AgentScope: A Flexible yet Robust Multi-Agent Platform. 2 2024.

[19] Yuan Li, Lichao Sun, and Yixuan Zhang. MetaAgents: Large Language Model Based Agents for Decision-Making on Teaming. *Proceedings of the ACM on Human-Computer Interaction*, 9(2), 8 2025.

[20] Talha Zeeshan, Abhishek Kumar, Susanna Pirttikangas, and Sasu Tarkoma. Large Language Model Based Multi-Agent System Augmented Complex Event Processing Pipeline for Internet of Multimedia Things. 2025.

[21] Alpár Jüttner and Péter Madarasi. VF2++—An improved subgraph isomorphism algorithm. *Discrete Applied Mathematics*, 242:69–81, 6 2018.

[22] Svante Janson and Gregory B. Sorkin. VCG Auction Mechanism Cost Expectations and Variances. pages 1–29, 10 2013.

[23] Robert Tarjan. Depth-First Search and Linear Graph Algorithms. *SIAM Journal on Computing*, 1(2):146–160, 6 1972.

Appendix A. Related Work: Swarm Intelligence and Agentic Science

1.1 Foundations of Swarm Intelligence

Swarm Intelligence (SI) shifts from centralized processing to decentralized, self-organizing systems modeled on biological collectives [7]. Defined by robustness, flexibility, and decentralized coordination, SI relies heavily on *stigmergy* [8] which is indirect communication via environmental modification to decouple agents spatiotemporally. While 1990s algorithms like Ant Colony Optimization [9] and Particle Swarm Optimization [10] successfully applied these principles to combinatorial and continuous optimization, they lacked the semantic reasoning required for complex scientific discovery.

1.2 Agentic Science and Centralized Architectures

The evolution from specialized computational tools to autonomous research partners, termed Agentic Science, is a foundational development in the AI-for-Science landscape. These systems are designed to autonomously execute the scientific discovery cycle, encompassing hypothesis generation, experimental design, execution, and iterative refinement [11, 12].

The dominant topology in contemporary Multi-Agent Systems (MAS) is the Centralized Orchestrator, or "Star Architecture." Frameworks such as AutoGen [1], MetaGPT [2], and ChatDev [13] exemplify this structure. While alternative topologies like Ring [14, 15], Graph [16, 17], and Bus Architectures [18, 19] exist, the centralized model remains prevalent [3, 20].

Appendix B. Detailed Architecture Specifications

2.1 The Stigmergic Environment (Knowledge Graph)

The foundational layer is the Shared Knowledge Graph (KG), a dynamic, multi-modal directed graph $G = (V, E)$. The vertex set V is the union of four disjoint subsets $V = V_C \cup V_D \cup V_H \cup V_O$:

- V_C (Concept nodes): Static ontological knowledge (e.g., "Superconductivity") linked to standard ontologies.
- V_D (Data nodes): Immutable references to raw experimental data or simulated outputs (e.g., via IPFS) serving as ground truth.
- V_H (Hypothesis nodes): Probabilistic propositions acting as "Digital Pheromones" by aggregating confidence scores and economic stakes.
- V_O (Operational nodes): System state including Agent Identities (reputation and economic profiles) and Task Manifests (code and resource requirements).

The edge set E defines semantic relationships: Ontological edges (structure), Epistemic edges (causal logic with confidence weights $w \in [0, 1]$), and Transactional edges (economic interactions).

2.2 Agent Swarm Specifications

The reasoning engine comprises specialized agents a_i, each instantiated with a Dual-System Cognitive Architecture. **System 1 (Intuitive)** utilizes an LLM core for heuristic generation and pattern matching. **System 2 (Reflexive)** employs an Agentic Critic engine for auditing via Chain-of-Verification [4].

Explorer Agent (A_{exp}) operates via Inductive and Analogical reasoning. It begins with a **Scan phase**, where System 1 identifies high semantic similarity between unconnected concepts in the KG. This is followed by a **Verify phase**; for inductive patterns, it checks statistical correlations in V_D, while for analogies ($G_{src} \approx G_{tgt}$), it uses algorithms like VF2 [21] for subgraph isomorphism matching. Finally, the **Broadcast phase** instantiates a "Proto-Hypothesis" node.

Hypothesis Agent (A_{hyp}) employs abductive reasoning to explain anomalies. It initiates **Abduction**, detecting structural holes or conflicts where $V_D \to \neg V_H$. Proposals undergo **Refinement** via a recursive loop where a Critic generates counter-arguments ($S_{counter}$) for the Generator to resolve (H'). After a **Consistency Check** against immutable data (V_D), the agent performs **Staking**: if internal confidence $C_{int} > \tau$, it locks RESEARCH tokens ($T_{stake} = \alpha \cdot C_{int}$) to signal epistemic value.

Experiment Agent (A_{lab}) functions by first performing **Derivation**, translating abstract hypotheses into concrete Prediction Tuples $P = \langle \text{Config}, \text{Perturbation}, \text{Expected Observable} \rangle$. It then proceeds to **Synthesis**, generating executable scripts π_{code} to test P, pre-screened via static analysis. Finally, the agent initiates an **Auction** by bundling the script into a Task Manifest broadcast to the Resource Auction.

Validator Agent (A_{val}) executes rigorous falsification and statistical auditing. Its workflow begins with **Adversarial Surface Mapping**, where System 1 identifies logical vulnerabilities in high-confidence Hypothesis Nodes (V_H) to generate stress-testing Challenge Vectors (C_{vec}). System 2 performs a **Computational Audit**, calculating metrics (e.g., R^2, p-values) on Data Nodes (V_D). If the hypothesis fails, a Falsification Proof is broadcast, slashing the Proposer's stake; otherwise, "Failed Challenge" metadata is appended, increasing the confidence score.

Resource Agent (A_{res}) serves as an infrastructure abstraction layer, encapsulating heterogeneous tools (Cloud Labs, HPC) into standardized economic services. It monitors the Task Manifest stream, semantically matching abstract requirements to hardware profiles. System 2 employs a **Dynamic Pricing** mechanism via load balancing algorithms (e.g., PID, VCG [5, 22]) responsive to queue depth and energy costs.

Enforcement Agent (A_{enf}) acts as the system's immunological defense, funded by a protocol-level tax. Its System 1 scans transactions and graph updates for behavioral anomalies. Suspicious clusters trigger a **Forensic Graph Audit** using algorithms like Tarjan's [23] to detect Wash Trading. Upon proving violations, A_{enf} executes binding sanctions, calling `slash_stake()` or issuing "Hazard Flags" to freeze dangerous Hypothesis Nodes.

Appendix C. Feasibility and Stability Analysis

A critique of market-based MAS is the potential for chaotic dynamics. We outline stability conditions for the SDS.

3.1 Prevention of Hypothesis Spam

In a centralized system, a hallucinating agent floods the planner. In SDS, the **Cost of Inquiry** acts as a filter. Let $E[U]$ be the expected utility of proposing hypothesis H, where B represents the bounty reward for successful verification and $S_{collateral}$ denotes the tokens staked by the proposer:

$$E[U] = P(H) \cdot B - (1 - P(H)) \cdot S_{collateral}$$

Rational agents will only propose H if $E[U] > 0$. Crucially, the stake is a function of reputation, defined as $S_{collateral} = f(R^{-1})$. As an agent's reliability (R) decays, the protocol demands exponentially higher collateral, effectively banning hallucination-prone agents from the network.

3.2 Market Equilibrium and Scalability

With N agents, will the auction mechanism scale? Vickrey auctions generally scale to thousands of participants efficiently. Furthermore, we do not require a single global auction. The Knowledge Graph is sharded; auctions occur locally within semantic clusters (e.g., "Materials Science" agents bid in a local pool distinct from "Genomics" agents), ensuring $O(1)$ scaling relative to total system size, provided the graph partitions are sparse.

3.3 Incentivizing Novelty via Information Gain

The protocol implements a Novelty Premium Mechanism to mitigate convergence on local minima. The internal pricing algorithm weighs the Return on Investment for a Hypothesis H inversely proportional to the Knowledge Density of its semantic neighborhood. Hypotheses located in "Structural Holes", bridging disparate clusters, receive a β-multiplier on validation bounties. As confidence approaches certainty, rewards asymptotically decay ("Diminishing Returns on Consensus"), forcing the swarm to migrate resources toward the epistemic frontier.