Global Literature Reasoning for Autonomous Materials Discovery Agents

Thomas M Pruyn [1 2] **Andrew Yi** [1] **Seyed Mohamad Moosavi** [1 2]

[1]*Chemical Engineering & Applied Chemistry, University of Toronto, Toronto, Ontario M5S 3E5, Canada* [2]*Vector Institute for Artificial Intelligence, Toronto, Ontario M5G 0C6, Canada*. Correspondence to: mohamad.moosavi@utoronto.ca.

1. Introduction

Scientific research fundamentally begins with the literature. Human and AI scientists alike read prior work, synthesize ideas across studies, identify gaps and contradictions, and build new hypotheses on accumulated knowledge. Despite rapid advances in artificial intelligence (AI) for scientific prediction and automation—including in self-driving laboratories (SDLs)—this foundational step of discovery, *reasoning over the scientific literature at scale*, remains largely inaccessible to current AI systems. Most existing literature search approaches, such as PaperQA [1], rely on localized retrieval of a small number of documents, limiting their ability to aggregate evidence, identify global trends, or reconcile conflicting experimental outcomes distributed across thousands of studies.

In this work, we present a literature-informed, autonomous framework for scientific discovery that treats the research literature as a first-class data modality and integrates it with machine learning and experimentation in an iterative, closed-loop system. Rather than focusing on isolated predictions, the framework is designed to support global reasoning across prior studies, causal hypothesis generation, and continuous refinement based on new evidence. The framework consists of two tightly coupled components.

The first component is a **global literature reasoning system** that moves beyond standard retrieval-augmented generation (RAG). Conventional RAG methods retrieve a small number of text passages, constraining reasoning to local context. In contrast, our approach aggregates and hierarchically synthesizes evidence across thousands of studies, enabling large language models (LLMs) to identify distributed patterns, reconcile conflicting experimental reports, and extract emergent design principles that only appear at corpus scale [2]. This capability is enabled by MOF-ChemUnity [3], a literature-integrated database that unifies structured materials data with experimentally grounded insights extracted from the scientific literature.

The second component is an **iterative, multi-tool autonomous discovery agent** that utilizes this global reasoning capability within a scientific workflow. Given a high-level research question, the agent decomposes it into sub-questions, explores them using literature reasoning, and autonomously constructs task-specific machine-learning models to test emerging hypotheses. Each discovery cycle updates the agent's internal state, allowing conclusions to evolve as new computational or experimental evidence becomes available [4]. By combining explainable machine learning with explicit literature-based justification, the agent supports causal reasoning rather than purely correlative inference.

We demonstrate this paradigm in the context of materials discovery, with a particular focus on metal–organic frameworks (MOFs) for carbon capture, utilization, and storage (CCUS), while emphasizing that the approach is broadly applicable to other data-rich scientific domains. The framework is applied to the identification and experimental validation of candidate MOFs for CCUS. Candidate materials proposed by the agent are computationally tested, and the resulting outcomes are fed back as direct design feedback, enabling the system to prioritize real-world performance over abstract metrics. We further outline the extension of this approach toward a fully integrated SDL, in which hypothesis generation, experiment execution, and learning are tightly coupled.

Overall, this work advances a step toward truly autonomous scientific discovery through SDLs—scaling reasoning across the scientific record, closing the loop between literature, computation, and experiment, and enabling both human and AI scientists to navigate bodies of knowledge that far exceed the capacity of any individual researcher or conventional AI system.

Acknowledgments

This work received support from the University of Toronto's Acceleration Consortium through the Canada First Research Excellence Fund under Grant number CFREF-2022-00042 and from the Data Science Institute (DSI) at the University of Toronto. S.M.M. research program receives financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC) through the Discovery Program.

References

[1] Jakub Lála, Odhran O'Donoghue, Aleksandar Shtedritski, Sam Cox, Samuel G Rodriques, and Andrew D White. Paperqa: Retrieval-augmented generative agent for scientific research. *arXiv preprint arXiv:2312.07559*, 2023.

[2] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.

[3] Thomas Michael Pruyn, Amro Aswad, Sartaaj Takrim Khan, Ju Huang, Robert Black, and Seyed Mohamad Moosavi. Mof-chemunity: Literature-informed large language models for metal–organic framework research. *Journal of the American Chemical Society*, 147(47):43474–43486, 2025.

[4] Ludovico Mitchener, Angela Yiu, Benjamin Chang, Mathieu Bourdenx, Tyler Nadolski, Arvis Sulovari, Eric C Landsness, Daniel L Barabasi, Siddharth Narayanan, Nicky Evans, et al. Kosmos: An ai scientist for autonomous discovery. *arXiv preprint arXiv:2511.02824*, 2025.