Accelerating AI Co-Scientists With HPC Infrastructure

Surya Appana [1] Alok Kamatar (Advisor) [2] J. Gregory Pauloski (Advisor) [2] Kyle Chard (Advisor) [2]

[1]University of California, Berkeley [2]University of Chicago

Abstract

We present MOFAI, an agentic AI Scientist coupled with high-performance computing (HPC) resources for the generation and property prediction of metal-organic frameworks (MOFs). MOFAI exhibits autonomous agents to enable tool-calling of linker generation, MOF assembly, molecular dynamics simulation, and deep learning-based property prediction in asynchronous and distributed fashion. MOFAI demonstrates success in leveraging mult-node computation to progressively discover stable MOFs with high CO_2 adsorption by learning from past successful MOFs.

Methods

We developed MOFAI for autonomous discovery and property prediction of metal–organic frameworks (MOFs) using Academy [2]. Academy provides a framework to launch and coordinate stateful agents on distributed research infrastructure. MOFAI is composed of the following agents:

1. **Tool-Calling Agent (LLM)** – interprets input objectives and calls relevant tools.
2. **Reflection Agent (LLM)** – evaluates and returns tool-calling agent's response or re-prompts the agent.
3. **Prediction** – calls MOFTransformer [1] to predicts MOF properties from structural descriptors.
4. **Generation** – generates SMILES strings of candidate linkers.
5. **Assembly** – assembles linkers and inorganic metal nodes into candidate MOFs.
6. **Validation** – performs a LAMMPS molecular dynamics simulation with the UFF4MOF force field to test structural stability under ambient conditions.

Agent Details and Compute Resources

Agent	Description	Code	Resource
Generate linkers	Generate SMILES strings of new linkers	API Calls to LLMs	Cloud
Assemble MOFs	Connect linkers & metal clusters Check bonds & atomic distances	Custom RDKit	1 CPU, multiple nodes 1 CPU, multiple nodes
Validate structure	Check geometry & bonds Test stability	cif2lammps LAMMPS	1 CPU, multiple nodes 1 CPU, multiple nodes
Property Prediction	Estimate CO_2 adsorption (mol/Kg ·Pa)	MOF Transformer	1 NVIDIA RTX 6000 GPU

Distributed Agents and Workflow

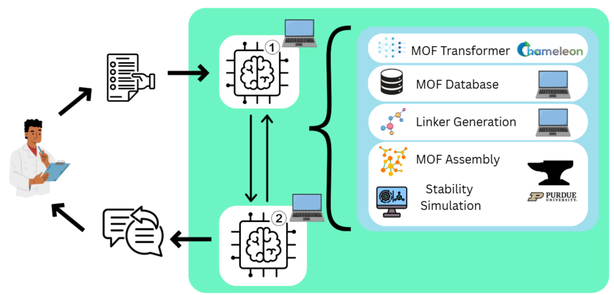

Figure 1. Graphic of AI Co-scientist workflow. The human scientist proposes a question or request. The tool-caller agent interprets the request and decides how to make use of its available tools before providing an initial response. The reflection agent assesses the response and either returns it to the user or prompts the tool-caller agent to provide a more appropriate response, possibly by using the tools differently. The MOFTransformer agent is hosted and run on a Chameleon compute node that includes an NVIDIA RTX 6000 GPU while linker generation, MOF assembly, and simulation agents are run on the Anvil supercomputer. The other agents are run locally.

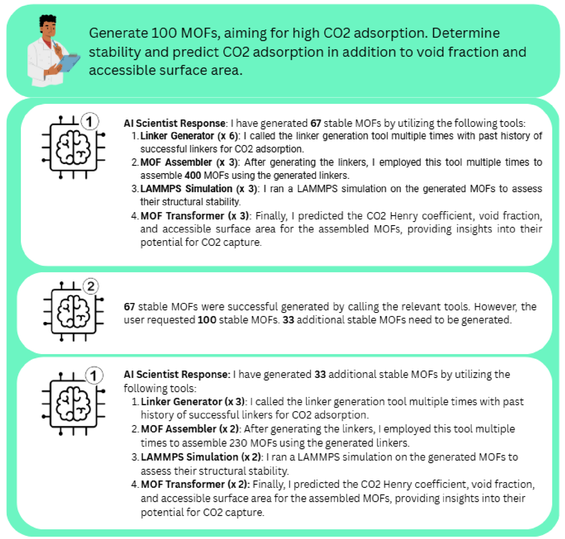

Figure 2. Example human scientist prompt with MOFAI response.

Evaluation: Scaling and Comparing LLMs

- **Scaling Test:** Evaluated MOFAI on Purdue Anvil to discover stable MOFs (<10 % structural strain) optimized for CO_2 adsorption across multiple Anvil compute nodes, fixing o3-mini for linker generation.
- **LLM Comparison:** Assessed MOFAI's ability to generate stable MOFs using different linker-generation LLMs by requesting 300 candidate linkers.

Scaling

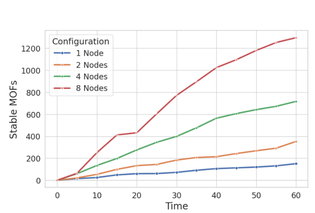

(a) Stable MOFs discovered in a walltime of 1 hour.

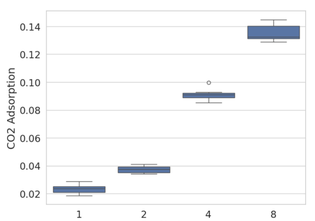

(b) CO2 adsorption of discovered stable MOFs

Figure 3. Stable MOFs discovered and their respect CO2 adsorption across runs provided with different numbers of compute nodes on Anvil.

Comparing LLMs

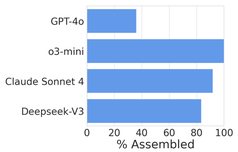

(a) MOF assembly success rate

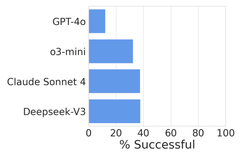

(b) Percentage of assembled MOFs that are stable.

Figure 4. MOF assembly and LAMMPS success across linker generation LLMs when generating 300 linkers.

Acknowledgements

This work is supported in part by the National Science Foundation OAC-2150500 award.

References

[1] Yeonghun Kang, Hyun Park, Berend Smit, and Jihan Kim. Moftransformer: A multi-modal pre-training transformer for universal transfer learning in metal-organic frameworks. *ChemRxiv*, 2022. doi:10.26434/chemrxiv-2022-hcjzc. URL `https://doi.org/10.26434/chemrxiv-2022-hcjzc`.

[2] J. Gregory Pauloski, Yadu Babuji, Ryan Chard, Mansi Sakarvadia, Kyle Chard, and Ian Foster. Empowering scientific workflows with federated agents, 2025. URL `https://arxiv.org/abs/2505.05428`.

[3] Xiaoli Yan, Nathaniel Hudson, Hyun Park, Daniel Grzenda, J. Gregory Pauloski, Marcus Schwarting, Haochen Pan, Hassan Harb, Samuel Foreman, Chris Knight, Tom Gibbs, Kyle Chard, Santanu Chaudhuri, Emad Tajkhorshid, Ian Foster, Mohamad Moosavi, Logan Ward, and E. A. Huerta. Mofa: Discovering materials for carbon capture with a genai- and simulation-based workflow, 2025. URL `https://arxiv.org/abs/2501.10651`.