AI & the Discovery of Molecules through Autonomous Laboratories

Connor Coley

We need new molecules. Medicines for diseases lacking treatments and for future pandemics, catalysts to decarbonize chemical industries and enable plastic recycling, stabilizers for battery longevity and energy storage, and agricultural agents to protect crops from blights and invasive species. Finding such molecules is a search. We *design* candidates based on what we know, then *evaluate* them in the lab, then learn and repeat. Artificial intelligence has accelerated design, but the bottleneck to practical impact is not a lack of ideas. It is the hard, tactile work of making and testing new physical matter.

As much as we hope to simulate the behavior of molecules in the digital world, every proposal eventually collides with the physical world. And in the physical world, things rarely go according to plan. Many molecules, including those with exquisitely complex structures found in nature, can take months or years to synthesize. When finally made, top-scoring molecules designed by AI often underperform.[1] Even as our ability to design, score, and rank structures improve, physical experiments remain the final arbiter. The practical question is not merely whether AI can propose molecules but whether these proposals can be efficiently translated into reality.

This "execution gap" is the central challenge of autonomous chemical laboratories – facilities that seamlessly close the loop between hypothesis and experiment – and has shaped how I view opportunities in AI for science for the past decade. What good is it to design molecules we cannot practically test? Chemistry offers a vivid case study, but the challenge echoes across biology and other scientific fields. The community has taken notice of this question and is slowly shifting focus from *design* to *execution*. Excitement has grown around the concept of autonomous laboratories with enough optimism to fuel hundred-million-dollar seed rounds.

The impulse behind these efforts is older than AI. Autonomous laboratories are part of a long tradition of scientists searching for new scientific tools that improve our capacity to measure, reason, and design.

My path to the use of AI in chemistry began in chemical engineering. Chemical engineering is often associated with the petroleum industry but is fundamentally about using quantitative reasoning to understand and

https://doi.org/10.1162/DAED.a.994

optimize physical processes from atomic to global scales. Optimizing chemical transformations for the pharmaceutical industry using design of experiments and statistical techniques progressed into cheminformatics and machine learning.

During my graduate work, I split my days between code and hardware by writing algorithms for planning synthesis recipes of new molecules while constructing physical platforms that could, in principle, execute those plans. The culmination of one large collaborative effort was a demonstration of manufacturing drug molecules through AI-designed reactions carried out by a robotic platform. It worked, but admittedly only within narrow boundaries. What looks simple on a whiteboard becomes far less so at the lab bench. The gap between the level of detail in today's AI-generated recipes and what is needed to make them real recalls the classic children's exercise of writing instructions for a robot to make a peanut butter and jelly sandwich.

Autonomous laboratories are, in essence, *embodied AI* systems that link design to evaluation. Our physical laboratories are the hands (and eyes) of our AI designers. Our best autonomous laboratories are currently fenced-in explorers that can move quickly and reliably inside a well-defined search space, where the scientific hypotheses they are allowed to generate are limited to what they are equipped to test. Similar approaches have already borne fruit in fully digital domains where AI systems learn within clearly defined boundaries and evaluate their own progress through simulation, as seen in efforts from game play to Google's foray into "empirical software."[2]

The current and perennial challenge is to expand the playing field for AI design tools. The new molecules we need to meet pressing societal challenges are unlikely to be within the capabilities of today's push-button systems. Progress will depend on connecting three fronts: better planning of detailed synthetic routes and experimental procedures, richer perception through sensing and real-time data interpretation, and orchestration systems that translate protocols into instrument commands, coordinate queues, and adapt to failures. The physical world adds complications – sticky powders, viscous liquids, sensitive intermediates – that will require advances in robotic dexterity and reliability. Even successful reactions can produce ambiguous results and, as the A-Lab study shows, understanding what those results mean is a challenge of its own.[3] The success of autonomous laboratories cannot be measured by the number of experiments performed or terabytes of information acquired per day. It must be measured by how effectively it accelerates learning through tight design cycles, faster loops, and fewer dead ends. These problems and pursuits lack the clean success metrics and benchmarking datasets that propelled progress in computer vision and protein structure prediction, yet they are essential to progress.[4]

The long arc of this field points toward a new style of scientific infrastructure: robust, broadly capable autonomous laboratories that link physical and digital worlds. And suppose we close the execution gap – What then? What happens when

laboratories can learn from every experiment, plan new ones on their own, and operate with the reliability of the best human scientists? In chemistry and medicine, such systems could map unexplored regions of chemical space, uncover new reaction pathways, and generate hypotheses that no individual researcher could have enumerated. They might shorten the distance between molecular insight and therapeutic impact, allowing us to test ideas that would otherwise be dismissed as too slow, too expensive, or too uncertain to pursue. Requesting a physical experiment might not feel so different from requesting a long-running simulation. More capital intensive and still with environmental impact considerations, yes, but no less convenient. Human scientists would be freed to focus on defining the goals, priors, and ethics of scientific inquiry.

Investing in this future is essential if we are to discover new medicines and materials at the pace the world requires. Opportunities beyond efficiency lie in using these tools to ask better questions, not merely to find faster answers.

ABOUT THE AUTHOR

Connor Coley is the Class of 1957 Career Development Professor and Associate Professor in the Department of Chemical Engineering and the Department of Electrical Engineering and Computer Science at MIT. He leads a research group that integrates chemistry and artificial intelligence to develop data-driven models that describe how molecules behave, interact, and react for the purposes of molecular design, synthesis, and analysis. He has published in such journals as *Nature*, *Science*, and *Nature Catalysis*.

ENDNOTES

[1] We will set aside the additional complexities that arise when progressing from the "in vitro" lab to, for example, clinical trials.

[2] Lizzie Dorfman and Michael Brenner, "Accelerating Scientific Discovery with AI-Powered Empirical Software," Google Research (blog), September 9, 2025, https://research.google/blog/accelerating-scientific-discovery-with-ai-powered-empirical-software.

[3] Nathan J. Szymanski, Bernardus Rendy, Yuxing Fei, et al., "An Autonomous Laboratory for the Accelerated Synthesis of Inorganic Materials," *Nature* 624 (7990) (2023): 86–91; and Josh Leeman, Yuhan Liu, Joseph Stiles, et al., "Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis," *PRX Energy* 3 (2024): 011002.

[4] See Jia Deng, Wei Dong, Richard Socher, et al., "ImageNet: A Large-Scale Hierarchical Image Database," 2009 *IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society, 2009); and Helen M. Berman, John Westbrook, Zukang Feng, et al., "The Protein Data Bank," *Nucleic Acids Research* 28 (1) (2000): 235–242.