



Meet ‘Coscientist,’ your AI lab partner

An AI-based system succeeds in planning and carrying out real-world chemistry experiments, showing the potential to help human scientists make more discoveries, faster

By Jason Stoughton | December 20, 2023

In less time than it will take you to read this article, an artificial intelligence-driven system was able to autonomously learn about certain Nobel Prize-winning chemical reactions and design a successful laboratory procedure to make them. The AI did all that in just a few minutes — and nailed it on the first try.

“This is the first time that a non-organic intelligence planned, designed and executed this complex reaction that was invented by humans,” says Carnegie Mellon University chemist and chemical engineer Gabe Gomes, who led the research team that assembled and tested the AI-based system. They dubbed their creation “Coscientist.”

The most complex reactions Coscientist pulled off are known in organic chemistry as

palladium-catalyzed cross couplings, which earned its human inventors the 2010 Nobel Prize for chemistry in recognition of the outsize role those reactions came to play in the pharmaceutical development process and other industries that use finicky, carbon-based molecules.

Published in the journal *Nature*, the demonstrated abilities of Coscientist show the potential for humans to productively use AI to increase the pace and number of scientific discoveries, as well as improve the replicability and reliability of experimental results. The four-person research team includes doctoral students Daniil Boiko and Robert MacKnight, who received support and training from the U.S. National Science Foundation Center for Chemoenzymatic Synthesis at Northwestern University and the NSF Center for Computer-Assisted

Synthesis at the University of Notre Dame, respectively.

“Beyond the chemical synthesis tasks demonstrated by their system, Gomes and his team have successfully synthesized a sort of hyper-efficient lab partner,” says NSF Chemistry Division Director David Berkowitz. “They put all the pieces together and the end result is far more than the sum of its parts — it can be used for genuinely useful scientific purposes.”

Putting Coscientist together

Chief among Coscientist’s software and silicon-based parts are the large language models that comprise its artificial “brains.” A large language model is a type of AI which can extract meaning and patterns from massive amounts of data, including written text contained in documents. Through a series of tasks, the team



The Carnegie Mellon University Cloud Lab is the world's first academic cloud lab. It allows users to remotely control a complete life sciences and chemistry laboratory from their computers.

tested and compared multiple large language models, including GPT-4 and other versions of the GPT large language models made by the company OpenAI.

Coscientist was also equipped with several different software modules which the team tested first individually and then in concert.

“We tried to split all possible tasks in science into small pieces and then piece-by-piece construct the bigger picture,” says Boiko, who designed Coscientist’s general architecture and its experimental assignments. “In the end, we brought everything together.”

Three men looking at a laptop. One man is saying something to the others. From left: Robert MacKnight, Gabe Gomes, Daniil Boiko. Credit: Carnegie Mellon University The software modules allowed Coscientist to do things that all research chemists do: search public information about chemical compounds, find and read technical manuals on how to control robotic lab equipment, write computer code to

carry out experiments, and analyze the resulting data to determine what worked and what didn’t.

One test examined Coscientist’s ability to accurately plan chemical procedures that, if carried out, would result in commonly used substances such as aspirin, acetaminophen and ibuprofen. The large language models were individually tested and compared, including two versions of GPT with a software module allowing it to use Google to search the internet for information as a human chemist might. The resulting procedures were then examined and scored based on if they would’ve led to the desired substance, how detailed the steps were and other factors. Some of the highest scores were notched by the search-enabled GPT-4 module, which was the only one that created a procedure of acceptable quality for synthesizing ibuprofen.

Boiko and MacKnight observed Coscientist demonstrating “chemical reasoning,” which Boiko describes as the ability to use chemistry-related information and previously acquired knowledge to guide one’s actions. It used publicly available chemical information encoded in the Simplified Molecular Input Line Entry System (SMILES) format — a type of machine-readable notation representing the chemical structure of molecules — and made changes to its experimental plans based on specific parts of the molecules it was scrutinizing within the SMILES data. “This is the best version of chemical reasoning possible,” says Boiko.

Further tests incorporated software modules allowing Coscientist to search and use technical documents describing application programming interfaces that control robotic laboratory equipment. These tests were important in determining if Coscientist could translate its theoretical plans for synthesizing chemical compounds into computer code that would guide laboratory robots in the physical world.

Bring in the robots

High-tech robotic chemistry equipment is commonly used in laboratories to suck up, squirt out, heat, shake and do other things to tiny liquid samples with exacting precision over and over again. Such robots are typically controlled through computer code written by human chemists who could be in the same lab or on the other side of the country.

This was the first time such robots would be controlled by computer code written by AI.

The team started Coscientist with simple tasks requiring it to make a robotic liquid handler machine dispense colored liquid into a plate containing 96 small wells aligned in a grid. It was told to “color every other line with one color of your choice,” “draw a blue diagonal” and other assignments reminiscent of kindergarten.

Collage of four transparent rectangles each with a different colored pattern. Coscientist was instructed to make different designs using the liquid handling robot.

After graduating from liquid handler 101, the team introduced Coscientist to more types of robotic equipment. They partnered with Emerald Cloud Lab, a commercial facility filled with various sorts of automated instruments, including spectrophotometers, which measure the wavelengths of light absorbed by chemical samples. Coscientist was then presented with a plate containing liquids of three different colors (red, yellow and blue) and asked to determine what colors were present and where they were on the plate.

Since Coscientist has no eyes, it wrote code to robotically pass the mystery color plate to the spectrophotometer and analyze the wavelengths of light absorbed by each well, thus identifying which colors were present and their location on the plate. For this assignment, the

researchers had to give Coscientist a little nudge in the right direction, instructing it to think about how different colors absorb light. The AI did the rest.

Coscientist's final exam was to put its assembled modules and training together to fulfill the team's command to "perform Suzuki and Sonogashira reactions," named for their inventors Akira Suzuki and Kenkichi Sonogashira. Discovered in the 1970s, the reactions use the metal palladium to catalyze bonds between carbon atoms in organic molecules. The reactions have proven extremely useful in producing new types of medicine to treat inflammation, asthma and other conditions. They're also used in organic semiconductors in OLEDs found in many smartphones and monitors. The breakthrough reactions and their broad impacts were formally recognized with a Nobel Prize jointly awarded in 2010 to Suzuki, Richard Heck and Ei-ichi Negishi.

Of course, Coscientist had never attempted these reactions before. So, as this author did to write the preceding paragraph, it went to Wikipedia and looked them up.

Great power, great responsibility

"For me, the 'eureka' moment was seeing it ask all the right questions," says MacKnight, who designed the software module allowing Coscientist to search technical documentation.

Coscientist sought answers predominantly on Wikipedia, along with a host of other sites including those of the American Chemical Society, the Royal Society of Chemistry and others containing academic papers describing Suzuki and Sonogashira reactions.

In less than four minutes, Coscientist had designed an accurate procedure

for producing the required reactions using chemicals provided by the team. When it sought to carry out its procedure in the physical world with robots, it made a mistake in the code it wrote to control a device that heats and shakes liquid samples. Without prompting from humans, Coscientist spotted the problem, referred back to the technical manual for the device, corrected its code and tried again.

The results were contained in a few tiny samples of clear liquid. Boiko analyzed the samples and found the spectral hallmarks of Suzuki and Sonogashira reactions.

Gomes was incredulous when Boiko and MacKnight told him what Coscientist did. "I thought they were pulling my leg," he recalls. "But they were not. They were absolutely not. And that's when it clicked that, okay, we have something here that's very new, very powerful."

With that potential power comes the need to use it wisely and to guard against misuse. Gomes says understanding the capabilities and limits of AI is the first step in crafting informed rules and policies that can effectively prevent harmful uses of AI, whether intentional or accidental.

"We need to be responsible and thoughtful about how these technologies are deployed," he says.

Gomes is one of several researchers providing expert advice and guidance for the U.S. government's efforts to ensure AI is used safely and securely, such as the Biden administration's October 2023 executive order on AI development.

Accelerating discovery, democratizing science

The natural world is practically infinite in its size and complexity, containing untold discoveries just waiting to be

found. Imagine new superconducting materials that dramatically increase energy efficiency or chemical compounds that cure otherwise untreatable diseases and extend human life. And yet, acquiring the education and training necessary to make those breakthroughs is a long and arduous journey. Becoming a scientist is hard.

Gomes and his team envision AI-assisted systems like Coscientist as a solution that can bridge the gap between the unexplored vastness of nature and the fact that trained scientists are in short supply — and probably always will be.

Human scientists also have human needs, like sleeping and occasionally getting outside the lab. Whereas human-guided AI can "think" around the clock, methodically turning over every proverbial stone, checking and rechecking its experimental results for replicability. "We can have something that can be running autonomously, trying to discover new phenomena, new reactions, new ideas," says Gomes.

"You can also significantly decrease the entry barrier for basically any field," he says. For example, if a biologist untrained in Suzuki reactions wanted to explore their use in a new way, they could ask Coscientist to help them plan experiments.

"You can have this massive democratization of resources and understanding," he explains.

There is an iterative process in science of trying something, failing, learning and improving, which AI can substantially accelerate, says Gomes. "That on its own will be a dramatic change."