




NIAID/NH/SPL

Some bacteria are becoming resistant to the antibiotics commonly used to control them.

AI IS TAKING ON ANTIBIOTIC RESISTANCE — HERE'S HOW

A suite of artificial-intelligence tools is helping to speed up the process of discovering new antibiotics. **By Jyoti Madhusoodanan**

When it comes to bacterial infections in the gut, antibiotics are effective yet broad-acting: although they might kill disease-causing species effectively, beneficial microflora can get caught in the crossfire. This indiscriminate effect can be harmful, particularly for people with Crohn's disease or other chronic gastrointestinal conditions. It also increases the likelihood that antibiotic-resistant strains of bacteria will evolve.

In 2023, microbiologist Jonathan Stokes at

McMaster University in Hamilton, Canada, began looking for options that could target pathogens with greater precision. He and his colleagues screened some 10,000 bioactive compounds for antibacterial activity against a strain of *Escherichia coli* that can cause severe gut infections. They filtered the results on the basis of various criteria, including toxicity to bacteria and structural novelty compared with existing antibiotics. "We got super lucky in that we only ended up with one molecule," says Denise Catacutan, the doctoral student in

Stokes's laboratory who led the work¹.

But the team needed to confirm that the promising molecule, named enterololin, was specific to its target pathogen, rather than acting as another broad-spectrum antibiotic. Typically, researchers rely on extensive biochemical screens, RNA sequencing or proteomics to elucidate how molecules such as enterololin disrupt bacterial pathways. This time, the team turned to artificial intelligence.

AI tools can accelerate the process of developing pharmaceuticals and are fast becoming

a crucial component of drug discovery. But using AI to identify an antibiotic's mechanism of action is still uncommon, says Regina Barzilay, a computer scientist at the Massachusetts Institute of Technology (MIT) in Cambridge.

Barzilay's lab developed a tool to fill that gap. DiffDock uses AI to predict how small molecules bind to proteins. It can thereby identify possible protein targets – and potential mechanisms of action for the small molecules. By applying this tool to enterololin, Stokes and his research group “could kind of narrow down our experimental pipeline”, Catacutan says. The team developed bacterial strains with mutations in the genes encoding predicted target proteins, and quickly confirmed DiffDock's predictions.

Barzilay's interest in antibiotics is personal. Her father contracted a bacterial infection of the spine that required complex surgery, and another family member survived an infection that didn't respond to any antibiotics. “We are so used to the idea that antibiotics are there to protect us,” she says. But that protection is precarious. Antibiotic resistance is a pervasive, growing global crisis; estimates suggest that drug-resistant infections could kill at least 39 million people by 2050.

Yet antibiotic development and manufacturing are expensive and rarely profitable, so drug companies are reluctant to invest. Discovering antimicrobials that can be synthesized easily – and cheaply – could help. An increasing number of researchers are turning to AI to address this need. Using machine-learning tools to tackle tasks *in silico*, from identifying new antibiotic candidates to predicting potential mechanisms of action, enables researchers to work faster – and on tighter budgets.

Strong foundation

Barzilay's interest in antibiotic development began in 2018, when she met MIT biomedical engineer James Collins at an institution-wide symposium on the use of AI tools. Barzilay and Collins teamed up to apply these techniques to antibiotic discovery. Stokes, who was a postdoctoral researcher in Collins's lab with expertise in high-throughput screening of small molecules, joined the effort.

The team developed a model based on neural networks (machine-learning architectures inspired by the human brain) that correlated molecular features – for example, bond types, atomic number and electronic charge – with properties such as solubility and microbial growth inhibition.

The researchers trained their model – called Chemprop – on data from 2,300 or so molecules that had been tested for their ability to inhibit the growth of *E. coli*. They then used the model to screen millions of molecules for potential drug candidates, eventually homing in on a kinase inhibitor that the group named halicin. This proved to have a potent effect on



Computer scientist Regina Barzilay is trying to use AI to decipher drugs' mechanisms.

several pathogenic species, including *Mycobacterium tuberculosis* (the causative agent of tuberculosis); drug-resistant *E. coli*; and *Acinetobacter baumannii* (an opportunistic pathogen that can cause infections in hospitalized individuals)². “We were able to create a model that can generalize to totally unseen classes of chemistry,” Barzilay says.

But having training data is only a first step – how it's labelled and classified is equally important, says Molly Bartlett, a chemical informatician at Imperial College London. Bartlett works with the Fleming Initiative, a multi-institutional collaboration headed by Imperial College London and Imperial College Healthcare NHS Trust that aims to combat antimicrobial resistance globally.

“Generative AI gives you this opportunity now to go beyond the sequence space that evolution has explored.”

When compiling the initiative's data, Bartlett searched the published literature for examples of high-throughput screening for molecules that can breach the outer cell membrane of pathogenic bacteria and accumulate inside the cells. Bartlett uses RDKit and xTB, computational tools that simulate and analyse molecular structures, to mimic how these molecules might behave when dissolved in water or in a cell's lipid membrane. She also categorizes the chemical features that are responsible for various properties of the molecules, such as their solubility and ability

to enter the bacterial cell. “If the input is not able to represent what makes the property happen, you're not going to be able to get a good prediction,” she explains.

In Stokes's experience, a strong training data set should include at least a subset of available clinical drugs, as well as potential antibiotics that are not currently used in the clinic. Training data should also be physically, chemically and structurally diverse – and should represent powerful antimicrobials, as well as ineffective ones, so that AI models can also learn what not to do. For those interested in building AI tools to find new antibiotics, “80% of your time has to be spent on data acquisition, data processing and data representation,” Stokes advises.

Bartlett, for instance, says that at least 10% of the molecules in her training data need to penetrate the bacterial envelope, so that the model can learn the characteristics that predict accumulation. Some databases she works with contain upwards of 100,000 molecules, but often only 3% of those molecules can enter the type of bacteria that interests her. She also tries to maximize the diversity of chemical structures represented in the training data. Without this breadth, she says, “you're not going to be able to have a predictive model”.

Generative AI tools have made her work easier. Bartlett and Catacutan struggled at first with writing the code necessary to run models, but now use AI tools such as Google's Gemini and OpenAI's ChatGPT to help with troubleshooting. Bartlett will occasionally provide Gemini with a toolkit's explanatory text file, often called a README file, then give the chatbot detailed instructions on what errors to look for in her code. “You still have to know what to ask for, but you don't have to be an expert in the architecture itself of the code,” she says. “It makes it really accessible.”

Peptide power-up

Also on a quest for chemical diversity is César de la Fuente, a synthetic biologist at the University of Pennsylvania in Philadelphia. De la Fuente's work focuses on antimicrobial peptides: naturally occurring short chains of amino acids that could prove to be potent antibiotics. His lab has also pioneered a technique called molecular de-extinction, which aims to ‘resurrect’ molecules that might have useful biological characteristics from extinct organisms.

De la Fuente and his team used a combination of neural networks to create a tool called APEX (antibiotic peptide de-extinction). They used this to screen a database of more than ten million peptides, and identified more than 37,000 that were predicted to have broad-spectrum antimicrobial activity³. About 11,000 of these were derived from the ‘extinctome’ – proteomes from extinct creatures, including, in this case, an ancient magnolia, a giant sloth and a Grant's zebra.

The researchers synthesized and tested 69 of these candidate molecules against bacterial pathogens and found that many of them had an unusual mechanism of action. Rather than targeting a pathogen's outer, rigid cell wall, these compounds acted on the inner cytoplasmic membrane – a strategy that might make them more robust as antibiotics, because bacteria are less likely to have evolved resistance to molecules they have never encountered. “Evolution is this beautiful planetary-scale optimization process,” de la Fuente says. It's “the biggest optimization experiment that we have ever seen”.

Making molecules

Having learnt from extinct molecules, de la Fuente's team took the next step: they developed a generative AI model that can design synthetic molecules that don't – yet – exist in nature⁴. “Generative AI gives you this opportunity now to go beyond the sequence space that evolution has explored, to come up with things that may have properties and functions that are more optimized in certain ways,” de la Fuente says.

The team provides its model – called ApexGO – with a peptide template, and specifies design goals, constraints and rules, such as how closely the designs must adhere to the starting peptide. People in the lab then step in to assess which designs could be viable. They might, for instance, exclude peptides that have too many hydrophobic residues, which would cause clumping in solution, de la Fuente says. The researchers then synthesize promising molecules, which they study in cultured cells and in animal models of infection. The team has so far synthesized and tested around 100 peptides, de la Fuente says. About 86% showed antimicrobial activity against at least one pathogen.

But most antibiotics are small molecules, not peptides. And small molecules designed by generative AI tools can often be challenging to make, Collins says, because they are too unstable, too expensive or simply chemically impossible. AI tools frequently design molecules that cannot actually be made, because the synthetic steps required don't follow real-world rules.

One recent study illustrates the scale of the challenge. Collins and his team trained a generative AI model to learn the features of molecules that inhibited the growth of *Neisseria gonorrhoeae* (the causative agent of gonorrhoea) and *Staphylococcus aureus* (which is often responsible for hospital-acquired infections of the skin, heart and other organs)⁵. The team then applied the model to libraries of chemical fragments, searching for ones that might be antimicrobial. Finally, the researchers used two generative AI approaches to design hypothetical chemical structures based on those pieces.

In the first of these approaches, fragments were modified repeatedly to build larger and more varied molecules, which were then scored for antibacterial potential. In the second approach, fragments were expanded atom-by-atom until they assembled into new chemical structures. The process yielded several million compounds with predicted specific antibacterial activity against *N. gonorrhoeae* and *S. aureus*. With the aid of another AI tool that helps predict which molecules can feasibly be synthesized, and after filtering on the basis of factors such as toxicity and antimicrobial potency, the researchers whittled down their long list to fewer than 5,000 molecules, only 90 of which seemed feasible.

But most of these molecules turned out to be impossible to produce, Collins says. The researchers eventually found a firm that could synthesize 22 of the compounds, six of which showed antibacterial activity. But as for the rest? Generative AI had provided an immense list of possibilities, but the number that they could actually test was “frustratingly small”, Collins says.

To further refine this process, Stokes and his colleagues built SyntheMol-RL, a generative algorithm trained to select molecular building

“There isn't a lot of financial motivation or financial resources to be screening massive libraries.”

blocks and chemical reactions and combine them *in silico* in a range of ways to create a structure. The algorithm works in conjunction with another AI tool that predicts the properties of molecules that SyntheMol-RL builds and offers feedback⁶. “SyntheMol is like a student, and the molecular-property predictor is like a teacher,” Stokes explains. “SyntheMol is learning to pick the correct building blocks to put together to make an antibiotic.”

To Stokes, another element of improving AI-based antimicrobial development is making the tools more accessible to researchers globally. In 2025, Autumn Arnold, a doctoral student in his lab, developed a free web-based tool that uses a suite of machine-learning models to identify potential antibiotics against several multidrug-resistant bacteria.

Making the tool accessible reduces the burden of needing to physically screen hundreds of molecules, Arnold says. “There isn't a lot of financial motivation or financial resources to be screening massive libraries,” she adds. “Something like machine-learning-based virtual screening is really helpful, so I wanted to make that technology accessible to antibiotic-discovery experts.”

Stokes's team does not track users or store their data. But statistics suggest that many

are from low- and middle-income countries, Arnold says, where antibiotic resistance imposes an especially large burden on health.

From possible to practical

For now, AI-augmented antibiotics are still in the early, preclinical phases of development, and none have so far yielded a new drug that helps patients. “The problem that we have today is we still don't really have a new antibiotic,” Barzilay says. The goal of AI-based drug design is not to have the perfect method, she adds, but to find working solutions to the antibiotic-resistance crisis. “To me, the art is really in taking the tools we currently have, which are already doing quite a bit, and translating them into something which is useful in clinic.”

Collins, de la Fuente, Stokes and others are working with various start-up companies to bring their molecules closer to clinical utility. Similarly, the Fleming Initiative has partnered with pharmaceutical giant GSK in the United Kingdom to help move its results into practice.

Diverse cross-disciplinary teams that can combine the expertise of biologists, chemists and computer scientists are essential to the success of AI-directed antibiotic discovery, says molecular microbiologist Andrew Edwards at Imperial College London, who also works with the Fleming Initiative. “No one person understands every aspect of the project.”

To de la Fuente, the future of cross-disciplinary research also includes finding effective ways for people and AI-based tools to work together. His lab frequently holds ‘human-machine meetings’, he says, in which the team integrates feedback from chemists, microbiologists, engineers and others with recommendations made by their various AI tools. The models process huge amounts of data and reveal patterns; humans can use that input to generate hypotheses that can be tested in the real world. “I think that's quite predictive of how science is going to go,” de la Fuente says.

But practicality is paramount, Stokes says: the goal is not an effective AI tool, but an effective antibiotic that is quick and inexpensive to produce at scale. It's an approach that Stokes describes as “blue collar”. “I don't care about inventing fancy models that look cool in a computer,” he says. “If we can't translate what we're doing in the computer to a wet lab experiment that can then progress towards clinical utility, we don't do it.”

Jyoti Madhusoodanan is a science writer in Portland, Oregon.

1. Catacutan, D. B. et al. *Nature Microbiol.* **10**, 2808–2822 (2025).
2. Stokes, J. M. et al. *Cell* **180**, 688–702 (2020).
3. Wan, F., Torres, M. D. T., Peng, J. & de la Fuente-Nunez, C. *Nature Biomed. Eng.* **8**, 854–871 (2024).
4. Torres, M. D. T. et al. *Nature Mach. Intell.* **8**, 841–856 (2026).
5. Krishnan, A. et al. *Cell* **188**, 5962–5979 (2025).
6. Swanson, K. et al. *Mol. Syst. Biol.* <https://doi.org/10.1038/s44320-026-00206-9> (2026).