

# Predictable response:

## Finding optimal drugs and doses using artificial intelligence

By Shraddha Chakradhar

Antkevych/Alamy Stock Vector, Oleksi/Atanasiev/Alamy Stock Vector

On 25 January 2015, Chih-Ming Ho, a former aerospace engineer, waited anxiously in his office at the University of California, Los Angeles, for an e-mail that would potentially alter the fate of a group of patients at the local liver-transplant center. Ho was collaborating with physicians there to help them to determine the right dose of immunosuppressants to give to eight severely ill individuals. Immunosuppressants help to prevent the body from rejecting transplanted organs. However, giving too much immunosuppressant medication can cause damage to the nervous system and kidney tissue. An added challenge of treating transplant recipients is that they are often on multiple drugs. These can range from antibacterials and antifungals to drugs that help to manage their liver disease. To make matters more complicated, as patients recover from transplant surgery, their physiology

fluctuates daily. Ho, who had broadened his research into biomedicine over the preceding decade, thought that he had found an innovative way of calculating the exact immunosuppressant dosage needed.

The e-mail that he had been waiting for was the first of a series of messages that would arrive over the next ten days, detailing how much and how many medicines each patient was receiving, along with blood-test results showing how quickly the various drugs were being metabolized. Ho's collaborators—including his son, Dean Ho, a biomedical engineer also at UCLA—hoped that by using data collected over the ten-day period, they could then predict how much of the drug each patient would need going forward, and thereby avoid unwanted side effects. But rather than attempting these calculations in the conventional way, they used a new tool: artificial intelligence (AI).

At the end of their data-collection period, the Hos and their collaborators were finally ready to start helping the physicians with their dilemma. By using a mathematical equation developed using artificial intelligence, they had calculated the proper dosage to give to the four patients in the experimental arm of the trial. The results were graphs—one for each of these four participants—that showed a smooth 'peak,' clearly delineating the best possible dosage for the individuals. The Hos revealed in having their work validated, recalls Dean, who is co-director of the Weintraub Center for Reconstructive Biotechnology at UCLA. "We knew that artificial intelligence would find the optimum for these patients," he says.

Over the next several weeks, the Hos continued to receive daily reports about the patients in the experimental arm, and they used each day's new data to recalculate the dosage for the following day. It turned out that

patients treated according to the information provided by the Hos were discharged from the Dumont-UCLA Liver Transplant Center about three weeks earlier than those who were dosed on the basis of physician calculations<sup>1</sup>. “To know that we were improving patient care was very gratifying,” Ho adds.


Although computers are commonplace in hospitals and research labs, the use of AI, or human-like decision-making in computers, is relatively new. AI has been used mainly for diagnostics, in which computers are fed hundreds of images of cancerous organs, for instance, to search for the presence of a tumor in a new X-ray image. The boom of ‘big data’ in medicine, combined with more sophisticated methods of training computers to process information in the way humans do but more efficiently than us, has enabled major pharmaceutical companies in recent years to use AI to comb through drug libraries for promising candidates faster than in previous years. However, the use of AI to help predict drug responses—whether to aid in further drug development or to inform clinical treatment—is only just emerging.

As combination therapies against cancer and infectious diseases become more commonplace, researchers and clinicians alike are still struggling to determine the right cocktail of drugs. “We take all of the guesswork out,” Ho says. Now, researchers such as him are developing algorithms that computers can use to quickly generate information about how a particular patient is likely to respond to therapy. Nigam Shah, a bioinformatician at the Center for Biomedical Informatics Research at Stanford University in California, says that the Hos’ efforts “are a novel approach to model the drug dose and efficacy responses from a patient.”

### Machine over mind

Artificial intelligence can be achieved in many ways, the most common perhaps being machine learning. In this approach, a computer ‘learns’ to recognize patterns among scores of data, so that when a new data point is presented, the computer knows which pattern best fits the new information. Other methods of AI include neural networks—a type of machine learning wherein artificial networks are designed to resemble the way neurons in a human brain are connected, so as to simulate thinking.

AI differs from conventional computer modeling in what information it requires to achieve a final result. With standard modeling, biologists need to know the relationship among the different data sets that are being mapped, according to Shah. Trying to track how



Tag team: Chih-Ming Ho (left) and Dean Ho (right).

mosquito populations affect the prevalence of malaria, for instance, requires knowing that there is a relationship between mosquitoes and malaria. But complex reactions such as drug response depend on several factors, including age, weight, genetics, proteins, disease type and many others that biologists might not even realize are important for predicting drug response. So, teasing out the exact relationships among all of the factors to determine how a drug will work becomes much more complicated. By contrast, AI doesn’t require prior knowledge of the relationships between these different biological factors. “You throw up your hands and say, ‘I’m not even going to try,’” Shah says.

The human mind may be able to use conventional statistical modeling to map out tens of variables, but a computer that uses artificial intelligence can sift through millions of variables over an ever-faster timeframe.


The agnostic approach to variables offered by AI has enabled researchers such as the Hos to quickly and accurately make predictions about drug responses. Chih-Ming Ho applied concepts from his days as an aerospace engineer, when he researched turbulence during flights, to simplify the process of determining drug response. One such concept—that of a complex system, where many different parts make it difficult to model—states that even if there are disturbances to the system, the way the system responds to any given disturbance is distillable into a fundamental equation that can be reliably modeled. The human body is one such complex system, Chih-Ming Ho recalls thinking nearly a decade ago. Given this, the

response to a disturbance to the human body like a transplant surgery or swallowing a pill should be able to be mapped using a simple equation, he says.

By this time, Ho had made the switch from studying turbulence to studying microfluidics, and then another to focus on deriving predictable reactions to drugs. He conducted laboratory experiments in which six different drugs in ten different dosages were applied in various combinations to human cell lines infected with herpes. The results—how many infected cells each drug killed while also preserving healthy cells—were fed to a computer after each experiment so that it could learn to rank the drugs and understand the pattern behind the ranking. After each experiment, the computer would eliminate possibilities of drug-dose combinations until, after a dozen or so iterations, the most successful cocktail emerged. The results were mapped onto a smooth curve in which the most elevated point represented the best option; the farther away the result was from the peak, the less optimum the drug-dose combination.

After repeating the process with cells exposed to other pathogens such as the bacterium that causes tuberculosis as well as cancer cells, another pattern emerged. No matter the disease, the drugs or the cell lines, the smooth curve could always be represented by a simple algebraic equation, as Ho had predicted, and it always took between 10 and 20 iterations to arrive at the result<sup>2</sup>. The team needed only to change out values in the equation that were specific to each patient

Dean Ho and Chih-Ming Ho



Chih-Ming Ho

**Taking flight:** Chih-Ming Ho, when he was an aerospace engineer, studying turbulence.

and calculated based on the dosages that each patient had received. The result would be a personalized map of drug responses. “I was really shocked,” Chih-Ming Ho says. “I couldn’t believe that it could be so simple.”

Since then, Ho and his son have tested this program, called Parabolic Personalized Dosing (PPD), in more than 30 different disease settings, and in more than 60 patients. These include the liver-transplant recipients in the 2015 trial, individuals with HIV infection and participants in an ongoing trial who have hematologic cancers. The Hos are currently in the process of getting their HIV trial results published. On their work predicting drug response so far, Dean Ho says, “We have zero misses.” He adds that their PPD program has thus far always managed to map out the best possible drug and dose combinations.

Whereas the UCLA group is focused on tailoring treatment regimens of approved drugs, a group at Virginia Polytechnic University (Virginia Tech), led by nutritional immunologist Josep Bassaganya-Riera, is using elements of machine learning in the hope of predicting people’s response to experimental drugs.

In a paper published in May, Bassaganya-Riera and his colleagues described how their combined modeling algorithm enabled them to identify the clinical responses that they might expect to see in patients with infections caused by a spore-forming type of bacteria called *Clostridium difficile*<sup>3</sup>. Current treatment for *C. difficile* includes antibiotics, which wipe out even beneficial bacteria—and patients are often averse to other treatments, such as fecal transplants. Previous work in

mice demonstrated that a prototype anti-inflammatory molecule called NSC61610 binds to a gut enzyme known as lanthionine synthetase C-like 2 (LANCL2) to dampen the inflammation caused by *C. difficile*, but also preserves the beneficial bacteria in the gut. However, the researchers were unsure whether this newly identified molecule would be an improvement on currently available therapies.


In their latest study, Bassaganya-Riera and colleagues identified the cellular and mechanistic characteristics that make up clinical responses to current treatments for *C. difficile* infection; these include changes in the T cell population and the composition of the microbiome. The authors then tested NSC61610 in mice to chronicle the physiological attributes that determine the response to this experimental drug in the animals. The combination algorithm matched the characteristics determined from this mouse experiment to the human equivalent in patients treated with currently available drugs to predict how a patient might respond to the new drug. The study found, for instance, that NSC61610 outperformed antibiotics and antitoxin antibodies when it came to preserving commensal bacteria, which signified that the molecule merited further testing—this time, in patients. “My high-level view is that this type of modeling allows us to simulate experiments that would take several months or years in a matter of three or four days,” says Bassaganya-Riera, who is director of the Nutritional Immunology and Molecular Medicine Laboratory at the Biocomplexity Institute of Virginia Tech and senior author of the study.

But without machine learning, he says, the team would have been forced to select only two or three characteristics to compare to determine drug response. Machine learning allowed them to include up to 52 parameters. “Every person is going to have their own certain set of parameters, and we need to understand what that unique mix of characteristics means, rather than analyzing each individual trait. Machine learning helps us to do that,” he says. Bassaganya-Riera’s company, Landos Biopharma, plans to initiate a phase 1 clinical trial in the second half of 2018 to test its candidate LANCL2-activating drug, BT-11, in individuals with Crohn’s disease.

### Data mining

Mining pre-existing databases of clinical data using artificial intelligence can also offer important clues about how drugs interact with each other. At the University of Kansas Medical Center, bioinformatician Mei Liu and her team reported finding adverse events from drug–drug interactions after mining through data in the US Food and Drug Administration’s Adverse Event Reporting System, or FAERS<sup>4</sup>. Liu and her team modified a pre-established AI algorithm known as association-rule discovery or mining (AR) to create another algorithm that allowed computers to identify drugs that had not only associations with negative symptoms, but also causative relationships with these adverse symptoms. The algorithm assigns values between 0 and 1 to causal relationships it identifies. The closer to 1 the value is, the stronger the likelihood that the drug combination causes a given symptom, Liu says. The team looked through more than 7,700 drugs and nearly 11,600 adverse events. “Machine learning allows us to explore such a huge search space,” Liu says, adding that the alternative would be extremely tedious.

When compared to the conventional AR method, the team’s technique suggested twice the number of causative relationships among the samples. She adds that the team informally validated its findings by asking physicians to check 100 of the algorithm’s results to see whether the adverse events identified by the model were also observed in the clinic, such as when a blood thinner called warfarin interacts with aspirin to cause bleeding. The system still needs fine-tuning before it can be used reliably to make clinical decisions, Liu says. But the next step for her team is to mine through adverse events and prescription information contained in electronic medical records. “There is already so much information in electronic medical records about what doctors have prescribed and what kinds of side effects they’ve observed,” Liu says. And the granularity



**Mind-like machines:** Machine learning is enabling the prediction of drug responses.

of that information—the exact dosage of a particular drug, for instance—will allow for more-accurate algorithms, Liu adds.

Meanwhile, at the University of California, San Francisco, neuroscientist Adam Ferguson is using artificial intelligence to better understand spinal-cord and traumatic brain injuries. He started by reviewing three rat studies from other teams that had aimed to investigate which of three possible therapies for traumatic brain injury (TBI) might be best to pursue. Those groups had been unable to draw meaningful conclusions, however, because of the overwhelming number of data points available to them for testing. “We can collect a bunch of data, but it very rapidly grows beyond human comprehension,” Ferguson says.

In a study published in February, Ferguson and his colleagues used the data that had been generated from the previous three studies, and applied AI techniques to explore which treatment might work best for TBI in rodent models<sup>5</sup>. One challenge they had to overcome was the common practice of assessing brain damage on scales, in which a number along a given scale represents a symptom or the severity of a symptom. Ferguson’s team used an artificial neural network to perform a calculation that converts these scaled measurements into simple numbers, such that it identifies the relationship between the extent of damage and all the biological factors that may determine drug response. “When we did this, we had this resolution for these drug effects that would be hard to determine without artificial intelligence,” Ferguson says.

Many scientists still don’t know what qualities to look for in order to help a patient achieve ideal recovery, Ferguson says, which

complicates things. “Rats will get better, but we don’t know what ‘better’ looks like,” he says. “We let the machine-learning algorithm tell us what better is.” The calculations and modeling that the machine performed enabled the researchers to first identify what characteristics would comprise ‘better’—things such as lesion size, motor ability, memory capacity and others. Once the machine calculated this score, the team was then able to measure and rank the possible combinations from the three therapies tested in the previous studies to identify which combination best stacked up to achieve the characteristics scored by the computer as ‘better.’

They found that a mix of an anti-inflammatory agent called minocycline and LM11A-31, a molecule that helps to support and maintain nervous tissue, was the best drug regimen to promote recovery following TBI among the rats included in the original three studies. The team also discovered that waiting until at least a week after injury to initiate physical therapy would enhance the recovery from injury. Ferguson explains that although the three preclinical studies had collected and curated the data correctly, the authors had too many potential variables to look at to make sense of the data. “Machines iterate through that stuff really quickly, and did what it would have taken 1,000 postdocs to complete,” Ferguson says. “It’s a way to carve through the complexity and find something simple.”

#### Future focus

A major challenge for those using artificial intelligence to make sense of massive data collections is ensuring the quality of the data used to predict responses to treatment.

“What’s going to matter is the data and where they came from,” says Jonathan Chen, a physician and medical-informatics researcher at Stanford University. “If you don’t have a good data source, it doesn’t matter what the result is.”

Chen adds that although a distinguishing factor of artificial intelligence is its agnostic way of coming up with predictions, these results may not always translate to actionable therapies. For example, an algorithm found that a chaplain being called to a patient’s room may be a good predictor of patient deaths<sup>6</sup>. Of course, that’s more likely because patients tend to call chaplains when they are close to death. “But it’s not like I can banish chaplains from visiting patients in order to save the patient from dying,” Chen says.

Chen cautions against overhyping artificial intelligence’s use in medicine. At the same time, he and a colleague urged scientists not to dismiss this approach at the first sign of failure<sup>7</sup>. Although genomics was overhyped, for example, the field’s shortcomings haven’t stopped scientists from pursuing the technology and studying its merits. Chen urges similar optimism when it comes to AI. According to Ferguson, much like how a microscope is best used by a knowledgeable pathologist, “you need someone using AI intelligently.”

“It extends scientists’ ability to do research, but doesn’t replace the scientists’ expertise,” he adds.

Back in Los Angeles, Dean Ho says that he is now often approached by patients with autoimmune disorders, cancers or other diseases who heard about his PPD program and are eager for him to apply artificial intelligence to guide their treatment regimens. “A question I’m often asked is, ‘Does this algorithm really represent something as complex as a human body?’” Ho says. Maybe not—but he expresses unabashed confidence in artificial intelligence’s capabilities. At the end of the day, he contends, “When you give a person a drug based on what the algorithm says, it’s going to work on the disease.”

**Shraddha Chakradhar is Nature Medicine’s associate news editor in Boston.**

- 1 Zarrinpar, A. *et al. Sci. Transl. Med.* **8**, 333ra49 (2016).
- 2 Wong, P.K. *et al. Proc. Natl. Acad. Sci. USA* **105**, 5105–5110 (2008).
- 3 Leber, A. *et al. Artif. Intell. Med.* **78**, 1–13 (2017).
- 4 Cai, R. *et al. Artif. Intell. Med.* **76**, 7–15 (2017).
- 5 Haefeli, J. *et al. Sci. Rep.* <http://dx.doi.org/10.1038/srep42474> (2017).
- 6 Choi, P.J., Curlin, F.A. & Cox, C.E. *J. Pain Symptom Manage.* **50**, 501–506 (2015).
- 7 Chen, J.H. & Asch, S.M. *N. Eng. J. Med.* **376**, 2507–2509 (2017).