

Making lethal molecules

by GWYNNE DYER

'We are but one very small company (among) many hundreds of companies using AI software for drug discovery and de novo design. How many of them have... the know-how to find the pockets of chemical space that can be filled with molecules predicted to be orders of magnitude more toxic than VX?'

This is a warning that requires a little explanation.

VX is a variety of 'nerve gas', first synthesized by Britain's chemical warfare lab at Porton Down in the 1950s. Just 8 milligrams (two grains of salt) will kill an adult human being if it comes into contact with their skin. Half a milligram if it is inhaled or swallowed. It's odourless and tasteless, and it comes in a mist so fine that it's practically invisible.

It took Porton Down almost a decade to develop it from the German nerve gases that the British discovered at the end of the Second World War, but with machine learning, you can now come up with something similar in practically no time.

The 'very small company' that issued the above warning, published in the scientific journal *Nature Machine Learning*, is 'Collaborations Pharmaceuticals, Inc.', based in Raleigh, North Carolina. Their business model is discovering new drugs, or more precisely designing them, and they have nothing whatever to do with poisoning people.

Like everybody in the drug discovery business these days, their primary tool is 'machine learning' - not exactly Artificial Intelligence as it was originally conceived, but a sub-category of AI that simply ingests vast quantities of data and searches it for resemblances.

Those resemblances will suggest possible new ('de novo') molecules that may not exist in nature but might be useful in treating disease. So you synthesize them, test them, and once in a while you come up with one that really does fill a gap in the existing pharmaceutical arsenal. But they might serve other purposes, too.

Two years ago, four research scientists working for Collaborations Pharmaceuticals - Fabio Urbina, Filippa Lentzos, Cédric Invernizzi and Sean Ekins - were asked to speak at the biennial conference hosted by the Swiss Institute for the Protection of the Population against Nuclear, Biological and Chemical Threats and Dangers - the Spiez Laboratory, for short.

The Spiez lab is one of five in the world certified by the Organisation for the Prohibition of Chemical Weapons to identify developments in chemistry, biology and enabling technologies that have implications for the Chemical and Biological Weapons Conventions. It asked the four to consider how AI technologies for drug discovery might be misused.

'The thought had never previously struck us,' wrote the researchers. 'We were vaguely aware of security concerns around work with pathogens or toxic chemicals, but that did not relate to us. Our work is rooted in building machine learning models for therapeutic and toxic targets to better assist in the design of new molecules for drug discovery.'

Their company had already designed a 'de novo molecule generator' that produced vast numbers of possible molecules and sorted them for bioactivity (desirable) and for toxicity (undesirable). Then the ones that passed those tests were examined for further good properties.

However, now they were investigating how easy it would be to make molecules with really lethal properties. 'We simply inverted this logic by using the same approach to design molecules de novo, but now guiding the model to reward both toxicity and bioactivity instead.' In fact, they chose VX as the kind of molecule that the AI should be aiming at. They didn't fabricate real molecules, but the AI produced 40,000 deadly new virtual molecules within six hours. Some resembled VX and other existing nerve

poisons, but others occupied a region of molecular property space that was entirely separate from the many thousands of molecules in the (VX/pesticide/organophosphate domain).?

At this point, the US researchers and the Spiez Lab must both have stopped to ask themselves: should we publicize this and risk giving bad people ideas, or keep it quiet and hope that nobody else realizes how easy it would be? The answer, obviously, is that people aren't stupid. The cat will be out of the bag soon no matter what they do now.

Nobody wants to restrict research for new drugs, but it's time to start thinking about how to control access to certain machine-learning models, public databases and chemical synthesis pharmaceutical plants that would be required to create ultra-lethal molecules.

And here's a small consolation. The standard antidote for all the nerve gases is atropine, but it's not very reliable, and you never have the little auto-injector around when you need it. However, in a pinch, Valium will probably do just as well. Really.