avoidance of noxious heat[7,11]. Therefore, the protein has not been implicated in heat sensing in mammals — until now.

Activation of this channel by heat in rattlesnakes or fruit flies involves specialized sequences of amino acids known as ankyrin repeats at the amino terminus[12]. These repetitive stretches are different in mammalian TRPA1, making the channel insensitive to heat under normal conditions[12]. By contrast, activation of TRPA1 in mice by inflammatory molecules is dependent on the oxidation state of three cysteine amino-acid residues located adjacent to the ankyrin repeats[13]. Vandewauw and colleagues therefore proposed that oxidation of these residues might be sufficient to activate the heat-sensing capabilities of mouse TRPA1.

To trigger oxidation of TRPA1 *in vitro*, the authors applied hydrogen peroxide to sensory neurons that lacked TRPV1 and TRPM3 but expressed TRPA1. They found that hydrogen peroxide treatment did sensitize TRPA1 in such a way that heat could now excite the neurons. This suggests that mouse TRPA1 must be oxidized to respond to heat. However, the authors do not provide a physiological mechanism for how TRPA1 might be oxidized *in vivo* to participate in heat sensing. This will no doubt be an area for future investigation.

Vandewauw and colleagues also performed electrophysiological experiments, which revealed that the three TRP channels they studied are co-expressed in more than one class of somatosensory nerve fibre — including at least two types of pain-sensing neuron that encode distinct sensations, such as dull pain or sharp pain. The expression of these channels in different neuronal subtypes might give rise to complex circuitry and crosstalk that aids the rapid avoidance of noxious heat. In addition, it could add a further layer of redundancy to an important protective response.

Finally, the group showed that, although the triple-mutant mice do not exhibit protective heat-avoidance reflexes, they do have the same preference as control mice for innocuous warm temperatures (30 °C) over noxious high temperatures (45 °C) when presented with a choice. This finding suggests that the sensing of painful heat might not have a substantial role in determining a preference for pleasurable temperatures over painful ones, warranting further investigation of warmth sensing — an area of active debate. A previous study has shown that whereas heat-responsive spinal neurons, which receive input from heat-sensitive somatosensory neurons, respond to absolute temperatures, cold-responsive spinal neurons, which receive input from cold-sensitive somatosensory neurons, respond to relative changes in temperature[14]. Together, these data imply that the main role of peripheral heat-sensing neurons, such as those described by Vandewauw *et al.*, might be

to mediate avoidance of high temperatures, rather than to set an animal's preferred temperature range.

Circuit-tracing and *in vivo* imaging will no doubt reveal the contributions of different subsets of peripheral neurons to the various representations of cold, cool, warm and hot temperatures in the central nervous system. Decoding these circuits will help to unravel the nuances of sensory perception and to elucidate the basis for diverse temperature preference between species or individual organisms. ■

Rose Z. Hill and **Diana M. Bautista** *are in the Department of Molecular and Cell Biology, University of California, Berkeley, California 94720, USA.* **D.M.B.** *is also at the Helen Wills Neuroscience Institute, University of California, Berkeley.*

e-mails: rzhill@berkeley.edu; dbautista@berkeley.edu

1. Caterina, M. J. *et al. Nature* **389,** 816–824 (1997).
2. Patapoutian, A., Peier, A. M., Story, G. M. & Viswanath, V. *Nature Rev. Neurosci.* **4,** 529–539 (2003).
3. Vriens, J. *et al. Neuron* **70,** 482–494 (2011).
4. Tan, C. H. & McNaughton, P. A. *Nature* **536,** 460–463 (2016).
5. Vandewauw, I. *et al. Nature* **555,** 662–666 (2018).
6. Caterina, M. J. *et al. Science* **288,** 306–313 (2000).
7. Bautista, D. M. *et al. Cell* **124,** 1269–1282 (2006).
8. Wilson, S. R. *et al. Nature Neurosci.* **14,** 595–602 (2011).
9. Gracheva, E. O. *et al. Nature* **464,** 1006–1011 (2010).
10. Neely, G. G. *et al. PLoS ONE* **6,** e24343 (2011).
11. Kwan, K. Y. *et al. Neuron* **50,** 277–289 (2006).
12. Cordero-Morales, J. F., Gracheva, E. O. & Julius, D. *Proc. Natl Acad. Sci. USA* **108,** E1184–E1191 (2011).
13. Macpherson, L. J. *et al. Nature* **445,** 541–545 (2007).
14. Ran, C., Hoon, M. A. & Chen, X. *Nature Neurosci.* **19,** 1201–1209 (2016).

This article was published online on 14 March 2018.

CHEMISTRY

AI designs organic syntheses

Software that devises effective schemes for synthetic chemistry has depended on the input of rules from researchers. A system is now reported in which an artificial-intelligence program learns the rules for itself. SEE ARTICLE P.604

DEREK LOWE

Synthetic organic chemistry is the science of building desired chemical structures from simpler parts. The knowledge and experience of researchers has always been the key to combining chemical reactions into successful synthetic schemes. But on page 604, Segler *et al.*[1] report that an artificial-intelligence program can design routes for synthesizing compounds that — at least, on paper — seem just as good as those produced by humans.

Organic chemists often work by thinking backwards as much as they do forwards when designing a synthetic route. The concept of retrosynthesis[2], introduced by E. J. Corey in the 1960s, and for which he was awarded the Nobel Prize in Chemistry in 1990, codified the way in which many chemists think (Fig. 1). When looking at a target molecule, they ask: "What could this have been made from? Which bonds could have been formed, and which atoms or chemical groups could have been added or transformed?" Then, the process starts again, as researchers try to determine the reactions that could have led to the precursor molecule. The aim is to work back to easily available starting compounds, while balancing the

factors that make a good synthesis — including the number of steps involved, the probable product yields of those steps, and the ease of use of the chemistry involved. Organic chemists deal constantly with such questions, for example when making compounds for testing in drug-discovery programmes.

Since the birth of synthetic organic chemistry in the mid-nineteenth century, a huge number of synthetic organic reactions have been reported across a literature that gets larger every hour. Before the 1980s, many chemists kept collections of handwritten, cross-referenced index cards containing useful reactions from the literature, to guide the design of synthetic pathways. These aide-memoires moved naturally on to digital databases as computer technology became widespread.

These days, chemists review the various methods for turning chemical group X into chemical group Y by drawing the molecular structures of interest using a computer program and then performing an online search for relevant reactions. This almost invariably produces a long list, from which researchers must select the most appropriate reaction for their needs, according to their knowledge and experience. Stringing such reactions into a useful
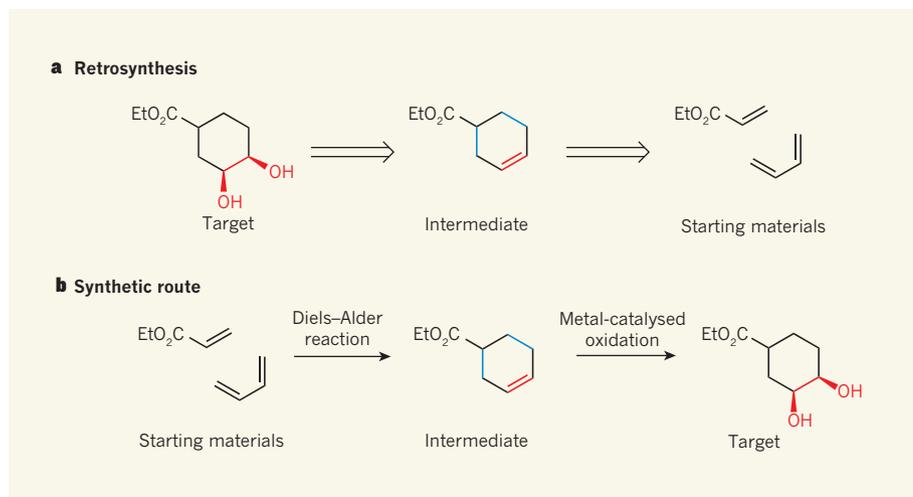
Figure 1 | Retrosynthetic analysis. Chemists use a strategy known as retrosynthesis to design schemes for constructing organic molecules. **a**, The first step is to work out a theoretical pathway by which a target molecule is sequentially converted into intermediates that can be easily prepared or purchased. In this example, the red bonds and atoms in the target can be derived from the red double bond in the intermediate molecule. The red double bond and the blue bonds in the intermediate can, in turn, be formed in a single step from two commercially available starting materials (Et, ethyl group). **b**, The second step is to choose reactions that will convert the starting materials into the target molecule: here, the Diels–Alder reaction forms the intermediate, which undergoes metal-catalysed oxidation to form the target. Segler et al.[1] report that an artificial-intelligence program has devised plausible synthetic routes for molecules by inferring design rules on its own.

synthesis has been thought of as a problem that only humans can solve.

But does it have to be? Could a sufficiently large and well-curated database of chemical transformations be used as the basis for a program that not only finds reactions, but also arranges them into plausible synthetic plans? Such programs have been sought since Corey's work in the 1960s, but (until recently) with little practical success.

Two fundamental problems have frustrated the dream. First, computing hardware simply could not tackle the scale of the challenge. Second, the chemical literature is hard to define in terms that a software program can understand: given reactions would work for the type of compound for which they were claimed to work (most of the time), but only under certain conditions. For example, group X would turn into group Y, unless group Z was present elsewhere in the reactant molecules. When group Z was present, the reaction might still work if group Q was nearby in the same molecule — but only, for instance, when the pH was lower than a certain value, or when the temperature was high enough, or when there was no water present.

There are various ways to overcome this second problem. One is to provide the program with an exhaustively human-curated list of the reactions that can enable a desired chemical transformation, and which takes into account all the limitations and conditions. The program can then combine such reactions into synthetic routes in ways that are broadly similar to those used to evaluate combinations of chess moves. This approach is starting to yield

results (see ref. 3, for example), and several competing commercial software products are available.

Segler et al. have investigated another method: instead of getting researchers to load their expertise into a machine, is it possible to design a program that learns by itself what researchers know? This concept has already produced startling results, with programs that can learn to play games such as Go on their own[4], rather than being trained using lists of human strategies.

> "The AI–generated synthetic routes were evaluated by trained organic chemists in a blind test for plausibility."

The authors devised a computational process that starts by automatically extracting chemical transformations from a large commercial database, being careful to include only reactions that have been reported several times. Their system accepts these well-precedented reactions as 'allowed moves' in organic synthesis. When the system is asked to devise a synthetic route to a target molecule, it works backwards from the target as would a human, picking out the most promising precursor molecules according to the design rules that it has learnt, and then seeing how feasible it is to synthesize those. The authors combined three artificial neural networks with a random Monte Carlo tree search — a type of search algorithm used by computers in certain decision-making processes — to narrow down the most

promising synthetic routes, without getting stuck too quickly on a particular path.

Importantly, the routes that emerged were evaluated not only by the program's scoring system, but also by trained organic chemists in a blind test for plausibility. When the chemists were asked to assess machine-generated synthetic pathways for target molecules alongside routes reported in the literature, they expressed no preference for the routes that had been shown to work by their fellow researchers. In other words, they found the chemistry suggested by the program to be as reasonable as the syntheses proposed by researchers.

This does not necessarily mean that all machine-suggested routes will work in the laboratory — but, as organic chemists know to their sorrow, many routes designed by humans fail there, too. Further development of the program could include such 'reductions to practice', to determine whether the machine-proposed routes are better (or, at least, no worse) than those devised by people. A study[3] this year that assessed a more conventional, hand-curated retrosynthesis program is notable for its inclusion of such a laboratory test. Achieving routes that are 'no worse' than those of researchers is a clear victory for Segler and colleagues' program, which arrives at pathways in considerably less time, and with much greater coverage of the literature, than a person could manage.

If such programs fulfil their promise, and there is little reason to think that they won't, synthetic chemists will find that a mainstay of their work starts to disappear. Technological innovations have had similar effects in the past, but usually by automating physical 'grunt work' that is missed by no one. Disconcertingly, developments in artificial intelligence encroach on the thinking part of the job. There will always be complex, unusual and unprecedented structures that such software cannot handle, but the task of solving more-routine synthetic questions will be taken out of the hands of researchers.

The idea that intellectual tasks can be categorized as automatable grunt work will probably be insulting to many chemists, and will certainly feel like a threat. But the use of artificial intelligence will actually free up time in which they can think more about the higher-level questions of which molecules should be made and why, rather than focusing on the details of how to make molecules. Not all researchers will welcome this shift. But it seems to be coming, regardless. ∎

Derek Lowe is a science writer based in Boston, Massachusetts.
e-mail: derekb.lowe@gmail.com

1. Segler, M. H. S., Preuss, M. & Waller, M. P. Nature **555**, 604–610 (2018).
2. Corey, E. J. Angew. Chem. Int. Edn **30**, 455–465 (1991).
3. Klucznik, T. et al. Chem **4**, 522–532 (2018).
4. Silver, D. et al. Nature **550**, 354–359 (2017).