Chemical synthesis

Machine learning made easy for optimal reactions

Jason E. Hein

An accessible machine-learning tool has been developed that can accelerate the optimization of a wide range of synthetic reactions — and reveals how cognitive bias might have undermined optimization by humans. **See p.89**

The optimization of reactions used to synthesize target compounds is pivotal to chemical research and discovery, whether in developing a route for manufacturing a life-saving medicine[1] or unlocking the potential of a new material[2]. But reaction optimization requires iterative experiments to balance the often conflicting effects of numerous coupled variables, and frequently involves finding the sweet spot among thousands of possible sets of experimental conditions. Expert synthetic chemists currently navigate this expansive experimental void using simplified model reactions, heuristic approaches and intuition derived from observation of experimental data[3]. On page 89, Shields *et al.*[4] report machine-learning software that can optimize diverse classes of reaction with fewer iterations, on average, than are needed by humans.

Machine learning has emerged as a useful tool for various aspects of chemical synthesis, because it is ideally suited to extrapolating predictive models that are used to solve synthetic problems by recognizing patterns in multidimensional data sets[5]. However, chemists need to learn new skills to correctly deploy machine learning in their research, thus limiting the widespread adoption of this approach. Shields *et al.* address this problem by reporting an open-source software toolkit that can be easily adopted by chemists.

A range of machine-learning methods are now available, and the first task when developing any new application is to choose the most appropriate method. The choice depends on the type of data (numbers, pictures and so on), the number of data points available to train the system, and the desired output[6]. Wrong choices can lead to false correlations being made during training and ineffective predictive models.

To train their model, Shields and colleagues selected a method that uses a machine-learning approach called Bayesian optimization. Bayesian-optimization algorithms have proved exceptionally effective in other applications, but the authors are among the first to develop a reaction-optimization toolkit that uses this approach. Their open-source software contains all the components necessary for researchers to carry out Bayesian reaction optimization for systems that have any number of experimental variables.

The toolkit first uses a simple workflow to carry out a quantum-mechanical calculation that encodes the reaction of interest in a machine-readable format involving what are known as chemical descriptors[7]. Reaction parameters that can be represented as a continuous series of numbers, such as temperature and concentration, are already in a form that can be interpreted by the algorithm. However, categorized reaction parameters, such as the identity of the solvent or catalyst, need to be provided by the chemist using one of several commonly applied molecular notations.

Each molecule in the reaction is then decomposed by the toolkit into a subset of numerical values that describe the molecule's inherent chemical properties (molecular weight, charge density, bond strengths and so on), which can be interpreted by the algorithm[8]. Some of the biggest pitfalls in the application of machine-learning methods to chemical systems arise in the execution of this decomposition process. After multiple trials, Shields and co-workers arrived at a balanced approach that can be generalized for a variety of reactions involving many diverse chemicals.

The second part of the workflow is the Bayesian-optimization step. As the authors' work highlights, Bayesian algorithms are well suited for reaction optimization because they excel at handling relatively small data sets[9]. Starting from sparse data, the algorithm creates a surrogate model in an attempt to mathematically define how the input variables (reaction parameters) will affect the output target (the reaction yield or another measure of performance).

At first, the model provides a poor approximation of the reaction system, but the algorithm also evaluates what is learnt when new reaction data are acquired to test the effects of the variables. The algorithm therefore suggests a new experiment for chemists to run, providing specific values for the reaction variables. Once the data from that experiment are available, they are added to the algorithm,

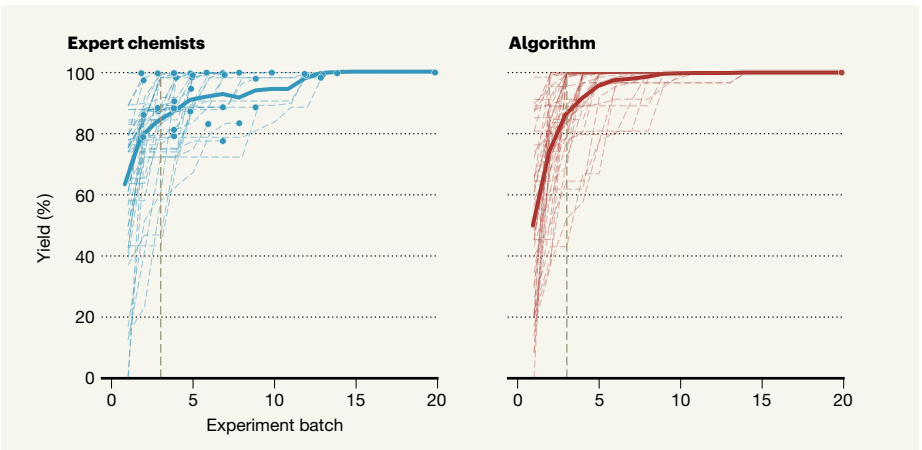

Figure 1 | Humans versus machine learning for reaction optimization. Shields *et al.*[4] have developed a machine-learning algorithm that optimizes the outcome of chemical reactions, and tested it in an optimization game. The authors selected a reaction, and defined five reaction variables that could be altered. They limited players to a fixed set of possibilities for each variable, and measured the reaction outcomes for all 1,728 possible combinations of variables. They then asked 50 expert chemists to carry out a virtual optimization: participants selected five combinations of variables and were shown the experimental outcomes, and could then select a new batch of five combinations to try to achieve the best possible reaction yield, up to a maximum of 20 batches (thin dashed lines indicate best yields per batch for each player; thick solid line indicates the mean average of the best yields). The algorithm also played the game 50 times, but started with random batches of variables. The experts made better initial choices, but the algorithm outperformed the players, on average, after the third batch of experiments (vertical grey dashed line). Moreover, the experts often did not achieve the optimal yield because they gave up too soon (blue dots indicate the end of each player's game), whereas the algorithm always achieved greater than 99% yields using its full allotment of batches.

which updates the model. The cycle then continues until the reaction performance meets the specified target, or the reagents are exhausted.

Shields *et al.* successfully applied this workflow to three reaction classes, in which the algorithm varied multiple reaction parameters, including the temperature, solvent and ligand (the molecule that binds to the metal centre of the catalyst). In each case, their algorithm successfully identified optimal reaction conditions using approximately 50 test experiments from a pool of up to 312,500 possible combinations of variables.

After tuning the algorithm using published data sets, the authors statistically evaluated its performance using an optimization game, in which the algorithm competed with expert chemists. The authors selected a reaction that would be optimized in the game, and then defined five reaction variables that could be altered, limiting the players to a fixed set of possibilities for each variable: 12 ligands, 4 bases, 4 solvents, 3 temperatures and 3 concentrations. The researchers then experimentally tested and measured the outcomes of all 1,728 possible combinations of variables.

Next, Shields and colleagues gave 50 expert chemists up to 100 attempts to carry out a virtual optimization: participants selected 5 combinations of variables, and were then shown the experimental outcomes, after which they could select a new batch of 5 combinations to try to achieve the maximum possible reaction yield. Likewise, the algorithm played the game 50 times, but each time starting with random experimental values. The experts made better initial choices, but the algorithm outperformed the players, on average, after the third batch of experiments (Fig. 1).

More notably, the algorithm consistently arrived at greater than 99% yields, which was possible only by using an unusual ligand that was not known to work well for the type of reaction targeted in the game. Overall, the game provided a crucial lesson in cognitive bias: most chemists ended the game early, having used only mainstream reagents, without realizing that they could have further improved the yields by making more-adventurous choices. Shields *et al.* have thus developed an accessible tool that could be used by non-experts to optimize a wide range of reactions. Importantly, the workflow carries researchers through the encoding and optimization processes for multiple chemistries without requiring changes to the code.

This empowering tool is poised to provide chemists with an alternative approach for reaction optimization, unlocking the many benefits of machine learning. Not only will it accelerate the pace of research in chemical synthesis, but it will also help to increase the range of reaction variables tested. Moreover, the tool eliminates the need for chemists to triage potential experimental variables to mitigate time and material costs. I expect this technology to be widely adopted, enabling rapid optimization of reaction conditions and discoveries off the beaten path.

Jason E. Hein is in the Department of Chemistry, University of British Columbia, Vancouver V6T 1Z1, Canada.
e-mail: jhein@chem.ubc.ca

1. Hardy, M. A. *et al. ACS Central Sci.* **6**, 1017–1030 (2020).
2. Tabor, D. P. *et al. Nature Rev. Mater.* **3**, 5–20 (2018).
3. Ley, S. V. *Angew. Chem. Int. Edn* **57**, 5182–5183 (2018).
4. Shields, B. J. *et al. Nature* **590**, 89–96 (2021).
5. Strieth-Kalthoff, F., Sandfort, F., Segler, M. H. S. & Glorius, F. *Chem. Soc. Rev.* **49**, 6154–6168 (2020).
6. de Almeida, A. F., Moreira, R. & Rodrigues, T. *Nature Rev. Chem.* **3**, 589–604 (2019).
7. Reid, J. P. & Sigman, M. S. *Nature* **571**, 343–348 (2019).
8. Zahrt, A. F. *et al. Science* **363**, eaau5631 (2019).
9. Reker, D., Hoyt, E. A., Bernardes, G. J. L. & Rodrigues, T. *Cell Rep. Phys. Sci.* **1**, 100247 (2020).

Medical research

Food for thought about the immune drivers of gut pain

Stuart M. Brierley

Debilitating gut pain is common, but the underlying cause is often unclear. It emerges that gut infection triggers localized immune responses that cause normally innocuous foods to be perceived as harmful, leading to persistent pain. **See p.151**

Pain evolved to alert us to and protect us from actual or potential tissue damage. There are three common forms: nociceptive pain, which is associated with the detection of damaging stimuli; inflammatory pain, associated with inflammation or infection; and chronic pain, which is a maladaptive, long-term form of pain[1]. On page 151, Aguilera-Lizarraga *et al.*[2] report evidence from mice and humans indicating a previously unknown mechanism that contributes to chronic gut pain.

When an individual has an obvious injury, such as a broken arm, we can readily appreciate that they are in pain. But if the injury site can't be easily pinpointed, it can be difficult to determine the origin of the pain. This is a common problem for people who have pain in their internal organs, with some forms of such pain being easier to diagnose than others. For instance, people who have inflammatory bowel disease might have easy-to-spot indicators of disease, such as gastrointestinal bleeding, inflammation of the intestinal lining, or the presence of distinctive biomarker molecules in faecal or blood samples[3]. However, people who have irritable bowel syndrome (IBS), a condition that affects 11% of the global population[4], lack such clear hallmarks of illness, and no obvious cause explains their chronic abdominal pain and concurrent constipation or diarrhoea.

> **"A bacterial gut infection can profoundly change local immune responses."**

Previous studies indicate that IBS is more common in women than in men[4], with IBS symptoms being triggered by factors such as stress[4], gastroenteritis (a disease caused by the ingestion of contaminated food or water)[4], alterations in gut microorganisms[5], and changes in communication between the gut and brain[6]. Aguilera-Lizarraga *et al.* now show that a bacterial gut infection can profoundly change local immune responses in the gut, resulting in certain foods being perceived as harmful, and thereby causing persistent gut pain (Fig. 1).

In healthy individuals, a process called oral tolerance results in the immune system 'ignoring' orally consumed substances[7,8]. An exception to this tolerance occurs for substances perceived by our bodies as being dangerous, such as harmful (pathogenic) bacteria, parasites and viruses. Our bodies identify foreign invaders by detecting molecular fragments called antigens, which provide a type of 'barcode' enabling our immune systems to specifically identify the intruders. Our immune systems can tag these antigens by producing antibodies that recognize them, enabling the pathogen to be quickly targeted if it appears again. Our defences should focus only on the 'bad guys', and leave the innocent bystanders alone. However, Aguilera-Lizarraga and colleagues hypothesized that a failure in oral tolerance might result in an indiscriminate targeting of both friend and foe.

To determine how this proposed breakdown in tolerance might occur in mice, the team used a model system that harnessed the