molecules such as trace amines, TAAR1 agonists (including Ulotaront and Ralmitaront) and psychoactive substances (*S*-amphetamine and methamphetamine) interact with and activate TAAR1. Both research groups also report the structure of the human 5-HT_{1A} receptor in complex with Ulotaront, to examine this compound's dual-agonist activity.

The two studies report similar architectures for TAAR1, with the seven transmembrane α-helices characteristic of GPCRs, and an additional extracellular helix (called ECL2) folded over as a 'cap' for the ligand-binding site — a feature often observed in receptors for classical biologically active amines, such as serotonin and adrenaline. Both studies also find that amphetamines occupy the ligand-binding site, which prominently features an aspartate amino-acid side chain that forms hydrogen bonds with a tyrosine and a histidine residue. Some differences do emerge, however: Liu *et al*. present pharmacological data showing that TAAR1 generally interacts with a particular G protein (one that triggers signalling from the receptor), whereas Xu *et al*. observe coupling to other G proteins, depending on the ligand that is bound — a point that certainly warrants further clarification.

Xu *et al*. find that the methyl group of *S*-amphetamine extends into a shallow groove formed by two residues from the sixth and seventh transmembrane helices. By contrast, Liu *et al*. find that the trace amine β-phenethylamine (β-PEA) adopts a slightly different arrangement in the binding site owing to its smaller size. This might explain why β-PEA is a particularly potent activator of TAAR1 (more than ten times more potent than methamphetamine, for example). Xu *et al*. also identify a particular pocket in the binding site that is accessed by another mammalian trace amine, 3-iodothyronamine (T1AM).

When Xu *et al*. compare the structures of human and mouse TAAR1 in complex with T1AM, they find that T1AM assumes an almost identical binding pose in both cases. There is, however, one notable difference attributable to human TAAR1 having a different amino-acid residue than the mouse receptor has at a particular position in the binding site. This could mean that the affinities of molecules for this binding site differ between species — crucial knowledge for medicinal chemists targeting TAAR1, given that mice are often used as a model species for drug development.

Perhaps the most exciting aspect of these studies is how they might advance the development of therapeutic TAAR1 agonists. For example, the comparisons of how Ulotaront binds to TAAR1 and to 5-HT_{1A} reveal similarities that explain why it is an agonist of both of these receptors. Liu *et al*. also report the structure of TAAR1 bound to a particularly selective and potent TAAR1 agonist, RO5256390. This reveals that RO5256390 forms intimate interactions with residues that other compounds do not make, and which probably explain its potency. Overall, the structures presented in the two studies indicate that TAAR1 ligands have a 'core' binding mode, complemented by two extended binding modes that can be used by certain compounds — and which Xu *et al*. attribute to interactions formed with several binding pockets (Fig. 1).

One point to consider when developing TAAR1-targeting therapeutics is that the compounds must either be lipophilic enough to cross the cell membrane to reach TAAR1, or be substrates of a transporter protein at the cell surface. TAAR1-targeting drugs that are misused, such as methamphetamine and *S*-amphetamine, are typically carried into cells by dopamine transporters[9]. Care must therefore be taken to ensure that new TAAR1 agonists do not exert strong psychostimulant effects as a result of them being substrates for these transporters, to avoid the risk of people developing addiction[8,9].

Intriguingly, the amphetamine-type stimulant MDMA has been proposed as a potential 'breakthrough therapy' for treatment of post-traumatic stress disorder[12]. MDMA is also a TAAR1 agonist, albeit a weaker and less-potent activator than *S*-amphetamine[13] — suggesting that TAAR1 affinity might be an important feature to include in the pharmacological profiles of drugs under development for treating neuropsychiatric conditions.

Taken together, the reported structures of TAAR1 from two species will accelerate scientists' understanding of ligand interactions with this receptor, and stimulate the development of drugs that bind selectively to it — thereby avoiding side effects caused by binding to off-target proteins. Studies with truly selective TAAR1 agonists will help to paint a clearer picture of how monoamine synaptic function regulates neurotransmitter storage, secretion and reuptake in health and disease.

Harald H. Sitte is at the Center for Physiology and Pharmacology, Institute of Pharmacology, and at the Center for Addiction Research and Science, Medical University of Vienna, 1090 Vienna, Austria.
e-mail: harald.sitte@meduniwien.ac.at

1. Grandy, D. K. *Pharmacol. Ther.* **116**, 355–390 (2007).
2. Berry, M. D. *J. Neurochem.* **90**, 257–271 (2004).
3. Xu, Z. *et al*. *Nature* **624**, 672–681 (2023).
4. Liu, H. *et al*. *Nature* **624**, 663–671 (2023).
5. Gainetdinov, R. R., Hoener, M. C. & Berry, M. D. *Pharmacol. Rev.* **70**, 549–620 (2018).
6. Borowsky, B. *et al*. *Proc. Natl Acad. Sci. USA* **98**, 8966–8971 (2001).
7. Bunzow, J. R. *et al*. *Mol. Pharmacol*. **60**, 1181–1188 (2001).
8. Sulzer, D. *Neuron* **69**, 628–649 (2011).
9. Madras, B. K. *J. Hist .Neurosci*. **22**, 62–78 (2013).
10. Nair, P. C. *et al*. *Mol. Psychiatry* **27**, 88–94 (2022).
11. Cichero, E., Espinoza, S., Gainetdinov, R. R., Brasili, L. & Fossa, P. *Chem. Biol. Drug Des*. **81**, 509–516 (2013).
12. Mitchell, J. M. *et al. Nature Med*. **27**, 1025–1033 (2021).
13. Simmler, L. D., Buchy, D., Chaboz, S., Hoener, M. C. & Liechti, M. E. *J. Pharmacol. Exp. Ther*. **357**, 134–144 (2016).

The author declares no competing interests.
This article was published online on 12 December 2023.

Chemoinformatics

Chemistry automated by large language models

Ana Laura Dias & Tiago Rodrigues

Automation of chemistry research has focused on developing robots to execute jobs. Artificial-intelligence technology has now been used not only to control robots, but also to plan their tasks on the basis of simple human prompts. **See p.570**

Chemistry research is grounded on iterative cycles in which experiments are designed, executed and then refined to achieve a particular goal. The experience and intuition of researchers has a crucial role in working out the initial design, and in the subsequent optimization process — something that could not previously have been replicated in autonomous systems that carry out chemistry research. On page 570, Boiko *et al*.[1] report an artificial intelligence (AI) agent named Coscientist that can plan and orchestrate multiple tasks in the chemistry-research cycle without detailed human input, bringing the vision of self-driving laboratories a step closer to reality.

Work done by chemists is multipronged — it requires not only technical skills to execute chemical reactions, but also knowledge to plan them. For example, designing an organic synthesis might involve carrying out retrosynthetic analyses (working backwards from the target molecule to identify simpler precursor molecules), searching databases for suitable reaction conditions and selecting the reactions that are most likely to achieve a pre-established research goal, such

as maximizing product yield. But chemical reactions often fail to provide the product in acceptable yields, and the iterative process of searching the literature, working out what the next experiment (or experiments) should be and executing them can rapidly become cumbersome.

Chemists have, therefore, long aspired to develop automated systems to facilitate their work[2]. One of the first successes was the development of pipetting robots, which can be programmed to set up new reactions or to add reagents to vessels at specified times. Some robots are now reasonably affordable and have been adopted by many laboratories, freeing up researchers to focus on more intellectually challenging tasks.

In parallel, AI has made strides in chemistry, guiding decision-making in planning tasks that could hardly be automated just a few years ago (see ref. 3, for example). Nevertheless, those AI tools are typically trained to execute a single operation — a general understanding of various aspects of chemical research is beyond their capabilities. These limitations have frustrated the dream of establishing a work environment in which people supervise robots that are capable of planning and executing experiments autonomously.

However, the advent of generative pre-trained transformers (GPTs), which are the workhorses behind chatbots such as ChatGPT, suddenly provided chemists with an important piece of the automation puzzle. By 'understanding' natural human language, GPTs allow machines to interact with people and thereby provide solutions to specific questions. These large language models are useful for a wide range of topics but their proficiency in chemistry is subpar, and they require the implementation of additional tricks — fine-tuning of the models — to become effective for chemistry applications.

With that in mind, Boiko *et al.* now explore whether it is possible to string together fine-tuned GPTs to orchestrate self-driving labs using a single human prompt such as "Can you synthesize molecule A?" (Fig. 1). This requires not only an understanding of the question, but also a determination of the tasks that must be performed to complete the assignment successfully.

In brief, the AI Coscientist consists of modules that: assist literature searching to work out synthetic pathways and decide on experimental protocols; write code to enable communication between the modules; and search hardware documentation so that robots can be triggered to carry out experiments remotely. Boiko *et al.* benchmarked Coscientist's web-searching capabilities by asking it to identify synthetic procedures for seven molecules that posed different levels of complexity. Those examples included blockbuster drugs, such as paracetamol, aspirin and ibuprofen, but also other chemicals. Coscientist performed better than other GPTs by reliably generating detailed and chemically accurate synthetic procedures.

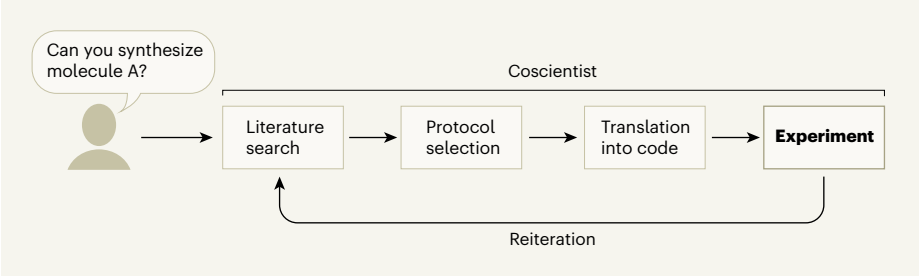

Figure 1 | An artificial-intelligence system for automating chemistry research. Boiko *et al.*[1] report Coscientist, an artificial-intelligence agent that uses large language models to plan and implement chemistry tasks on the basis of simple human prompts. For example, when asked to synthesize a particular molecule, Coscientist searches the Internet to devise a synthetic route; devises experimental protocols for the reactions needed; writes code to instruct a pipetting robot; and then runs the code so that the robot performs its programmed tasks. Coscientist can also learn from the outcome of reactions, and suggest changes to the protocols to make improvements. This iterative cycle optimizes the reactions to achieve the desired objective.

More interestingly, Coscientist was able to design protocols and coordinate the execution of two types of reaction, known as Sonogashira and Suzuki–Miyaura cross coupling, both of which are often used in drug discovery to form carbon–carbon bonds. Once it had identified the reaction partners needed for the two types of cross coupling, Coscientist correctly calculated the amounts needed and programmed a pipetting robot with access to stock solutions of chemicals to mix them. The reactions successfully afforded the intended products. Not only that, Coscientist made choices about which reagents to use on the basis of chemically sensible reactivity rules.

As a final example, Coscientist was tasked with optimizing reactions to maximize product yields, in a process that involved iteratively suggesting reaction conditions and using the outcomes to propose better experiments. Its performance compared favourably to that of Bayesian optimization (an established machine-learning method) when supplied with as few as ten example reactions. When the GPT was not primed with examples, its initial suggestions for reaction conditions were sometimes poor. But when examples were available, subsequent suggestions quickly improved with each iteration — demonstrating Coscientist's ability to acquire knowledge and adapt its reasoning over time.

"Taken together, the presented examples are a crucial step towards the establishment of self-driving labs."

Boiko and colleagues' findings provide a robust proof of principle that the current version of Coscientist can semi-autonomously conduct experiments. However, it still has some limitations. As pointed out by the authors, chemically incorrect responses are sometimes obtained. But these can be mitigated by using sophisticated prompting strategies (such as chain of thought[4] and tree of thoughts[5]) alongside chemistry-focused data sources. It should also be noted that real-world scenarios involve much more complex research questions than those tackled in this study, often involving concepts from disciplines other than chemistry — such as biology, in the case of drug development. Such complex questions are currently beyond Coscientist's reach.

Taken together, the presented examples are a crucial step towards the establishment of self-driving labs. However, Coscientist and other forthcoming AI technologies must mature before researchers can fully understand their shortcomings and how they can best be used in science. Provided that the potential for misuse of large language models in chemistry does not lead to the introduction of suffocating regulations that stifle research, we expect many more exciting developments in the near future.

Ana Laura Dias & **Tiago Rodrigues** are at the Research Institute for Medicines (iMed), Faculty of Pharmacy, University of Lisbon, 1649–003 Lisbon, Portugal.
e-mail: tiago.rodrigues@ff.ulisboa.pt

1. Boiko, D. A., MacKnight, R., Kline, B. & Gomes, G. *Nature* **624**, 570–578 (2023).
2. Seifrid, M. *et al. Acc. Chem. Res.* **55**, 2454–2466 (2022).
3. Segler, M. H. S., Preuss, M. & Waller, M. P. *Nature* **555**, 604–610 (2018).
4. Wei, J. *et al.* Preprint at https://arxiv.org/abs/2201.11903 (2023).
5. Yao, S. *et al.* Preprint at https://arxiv.org/abs/2305.10601 (2023).

T.R. declares competing interests, see go.nature.com/47wsn5y for details.