Check for updates

The perils of machine learning in designing new chemicals and materials

To the Editor — Machine learning is poised to revolutionize practice in chemistry and materials science. Already, machine learning is being used to find new pharmaceutical compounds, including in the fight against the COVID-19 pandemic. This holds great promise for the future, but also great peril. Right now, too little attention is being paid to the downside, as pointed out in a recent Comment by Urbina et al.[1].

It is easy to recognize the benefits of the machine-learning approach to, for example, testing chemicals and materials for toxicity — an area that we work on as a combined team of computer scientists and chemists. First, the need is obvious when you consider that less than 1% of the chemicals registered for commercial use in the United States have undergone toxicity characterization, whether they are used for medicinal purposes or for fracking. Moreover, there are many scientific, ethical, and economic advantages to replacing the animals currently used in toxicity tests with non-animal test systems, and great speed and cost advantages in using computer systems. Second, material and chemical usage has increased to 60 billion tonnes per year during the twentieth century[2], underscoring the advantages of a rapid machine-learning approach for toxicity characterization. Finally, the number of materials and chemicals that can be designed digitally far exceeds the number that have been well characterized. For example, our estimates based on the number of material combinations with six surfaces exceed trillions, while the organic chemicals based on only hexanes exceed 10^{30} (Fig. 1), clearly indicating the vastness of possibilities.

Less obvious are the disadvantages that this new power of machine intelligence brings. Just as the machine-learning analysis of bodily fluids might be used to develop medicines targeted to an individual, it may also be used to engineer viruses or toxins that might infect only certain people depending on their genes — a truly frightening prospect. Similarly, while machine learning might help to invent improved materials, such as biodegradable plastics and longer-lasting batteries, it is also easy to imagine how it might be used to design a tasteless compound that could be used to poison a community's water supply.

The scientific establishment has confronted dual-use problems similar to this before. When gene editing became possible, for example, leading scientists and ethicists called for a moratorium on clinical applications of germline gene editing, which involves inheritable alterations to the DNA of embryos to change the physical and mental capabilities of newborns. But a moratorium also means that we sacrifice some of the potential advantages that such advances might bring to society.

As machine-learning tools are made more broadly available for communities to use when making new compounds and materials, the possibilities of misuse also increase and need to be guarded against. In the current environment, the wide availability of computing capabilities means that anyone with access to the internet can potentially misuse these tools. The time to recognize and address this peril is now.

Recently, our university's Ethics and Society Review panel reached out regarding one of our proposed projects, which involved the use of machine learning to predict the toxicity of chemicals and materials. The panel raised important questions about the ethics and societal consequences of our research. On the one hand, once perfected, this power could be used to scan for unwanted toxic materials — for example, in all the chemicals that are used in fracking fluids to extract oil. On the other hand, it could also be used by malicious actors to search for new ways to poison the ground or water. Specifically, the panel told us we should think about ways to control the distribution of the software, the model, and its output to minimize potential misuse.

After discussions with the panel, we sought the advice of other experts on how to overcome this dual-use problem. We contacted many other scientists and engineers, including those who have deployed open software or databases, to learn from their experiences. To our profound dismay, we found very little guidance, either from federal agency officials or the corporate world. We found plenty of discussions about the ethical use of artificial intelligence algorithms when it comes to protecting privacy or preventing bias, such as the Montreal Declaration for a Responsible Development of Artificial Intelligence, but nothing specifically about the dual-use problem. Maybe there are people out there thinking about this, but they are not easy for us to find before embarking on these kinds of projects. That's a problem.

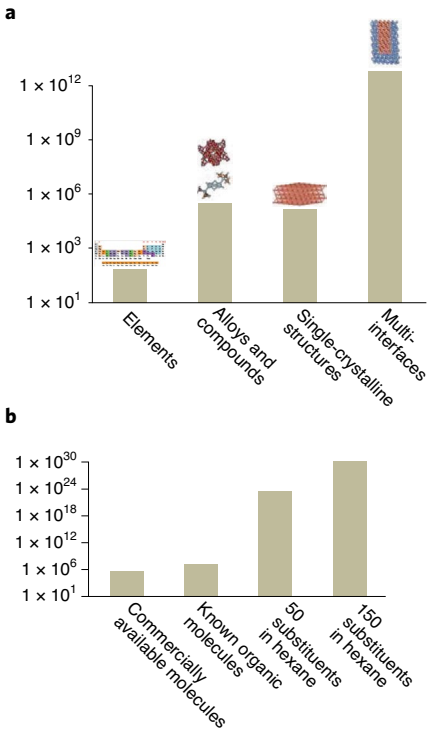

Fig. 1 | Illustration of the number of possible material and chemical combinatorial possibilities. a, 70 elements from the periodic table translate to 343,000 possible three-element alloys and compounds; 150,000 single-crystalline structures can lead to trillions of interfaces. **b**, The number of chemicals that are commercially available is about 350,000 (ref. [3]) and the number of known organic molecules is less than 20 million, whereas the number of possible chemicals by substitution in hexane itself is of the order of 10^{30}. This is a lower bound of the possible chemical combinations in hexane itself, as this does not include stereochemical isomers.

There is an intrinsic conflict between making work public, to encourage adoption and improvement of new codes and databases, and protecting it from abuse and misuse. We call for a conference among experts and interested parties to come up with a workable plan that meets as best as possible

both objectives: wide utility of the new capabilities while ensuring safe deployment. The research community, and those who fund this community, need to tackle this issue before it grows into a serious problem. ❐

Sadasivan Shankar[1,2] ✉ and Richard N. Zare[3] ✉

[1]*SLAC National Accelerator Laboratory, Menlo Park, CA, USA.* [2]*Department of Materials Science and Engineering, Stanford University, Stanford, CA, USA.* [3]*Department of Chemistry, Stanford University, Stanford, CA, USA.* ✉*e-mail: sadas.shankar@stanford.edu; zare@stanford.edu*

Published online: 22 April 2022
https://doi.org/10.1038/s42256-022-00481-9

References

1. Urbina, F. et al. *Nat. Mach. Intell.* **4**, 189–191 (2022).
2. Krausmann, F. et al. *Ecol. Econ.* **68**, 2696–2705 (2009).
3. Pelley, J. Number of chemicals in commerce has been vastly underestimated. *Chemical & Engineering News* (12 February 2020).

Acknowledgements

We would like to acknowledge the Stanford Human AI Institute (HAI) and HAI Ethics Panel for bringing the various issues to our attention; K. Geiser (Massachusetts) and others for sharing their experiences; and V. Shankar (Stanford) for his analysis.

Competing interests

The authors declare no competing interests.