



Science in the age of machine learning

Abbas Ourmazd

The rise of machine learning is moving research away from tightly controlled, theory-guided experiments towards an approach based on data-driven searches. Abbas Ourmazd describes how this change might profoundly affect our understanding and practice of physics.

Traditionally in physics, the more tightly controlled an experiment is, the better. Tight experimental control minimizes uninteresting effects, suppresses effects we are not looking for ('nuisance parameters'), constrains the questions asked ('Higgs or no Higgs?'), directs where to look ('Expect a bump there!') and determines what is learned ('The Standard Model still works'). The dividing line between interesting and uninteresting is drawn by recourse to theory — the theory being tested.

In contrast, the modern tech industry is based on the premise that 'interesting' happens unchoreographed and that one learns most by unobtrusive observation without preconceived notions. Unobtrusive observation shifts the focus from tightly controlled, theory-guided experiment to data-driven search. The key questions are now entirely different. How to learn from data without preconceived notions? How to deal with noise and quantify uncertainty? How much data is enough? And perhaps most importantly, in the absence of a theory-driven search, what form does 'scientific understanding' take?

Many would say the problem of learning from data has already been solved by machine learning (ML). In supervised ML, an algorithm is trained with data until it has learned to provide the right answers for the training dataset. This acquired ability to answer correctly is validated with previously unseen ('test') data. Many flavours of ML exist, with deep learning perhaps the most famous. For a growing number of tasks, the performance of ML algorithms is on par with, or supersedes, that of humans. Regardless of accuracy, if you want to find out what a billion people think about a topic on any particular day, ML is the only game in town. This means ML is here to stay.

But questions abound. These include whether an algorithm trained on one dataset can be used to produce reliable answers about a different dataset, whether a particular algorithm is robust against noise or attempts to deceive it, what the basis is for the answers an algorithm provides and whether these answers are free of bias. Putting aside these important questions, I concentrate here on the nature of scientific understanding in an ML-dominated world, and what that means for the new types of experiment one could undertake.

Understanding from machine learning

In the physical sciences, we aspire to encapsulate understanding in mathematical equations, the solutions of which predict the outcome of experiment. The Schrödinger equation and Einstein's field equations are celebrated examples. Such equations provide insight into how the physical world works. Quantum mechanics rests, ultimately, on the uncertainty principle, and general relativity on massive bodies warping spacetime. These concepts represent a deep understanding of the natural world.

The outcome of deep learning is radically different. In essence, deep learning uses multilayered networks to attach labels to objects, for instance, to identify your spouse in a previously unseen snapshot. But the basis for this identification is difficult to discern. The patterns used for recognition are often so abstract that it is hard to understand why they were used. We recognize — that is, attach labels to — individuals all day long, but nobody would argue that we know how the brain works. Providing the right answers does not, in itself, constitute scientific understanding. This problem becomes more acute as one increases the depth of the algorithm, that is, the number of interconnected layers in deep learning. In a real sense, the deeper the learning, the shallower the understanding.

There are many schools of thought on the most appropriate way forward. Some work to derive a mathematical description of deep learning^{1,2}. Others construct analytical pipelines that can inject physical insight into the learning cycle³. Yet others pursue mathematically rigorous approaches that reflect the nature of the world we live in, more precisely, the nature of operations possible in the ambient space, such as rotations and shifts⁴. Although a knowledge of the ambient space can be learned by having robots play with real objects at leisure⁵, it is possible — and far more efficient — to incorporate the nature of allowed operations in the space producing the data into the mathematical fabric of algorithms⁴. In this way, one might prevent a STOP sign lying on the ground from being confused with a 30 mph speed limit⁶!

Being an optimist, I think we will discover means to extract scientific understanding from ML, simply because of the sheer magnitude of the effort directed at

artificial intelligence in general, and ML in particular. My personal bet is that geometric ML, in which information is extracted from the intrinsic geometry of the data, will play an important role, if only because the information is encoded in the metric of the data manifold⁴, and general relativity already shows how scientific understanding can be formulated in such terms. In this scenario, ML identifies the curved hypersurface (manifold) defined by the data, and knowing the metric of the manifold makes it possible to generate intelligible answers by navigating on, or projecting onto, this manifold^{4,7}. Regardless of the specific approach, we are confronted with a different definition of what constitutes scientific understanding.

Machine learning in practice

What are the practical implications of an ML-dominated world for science? Of course, ML will be used to operate complex, difficult-to-control scientific instruments, such as high repetition-rate X-ray free-electron lasers. ML will also help with the ‘data deluge’ expected at these facilities; how else would we deal with petabyte datasets? But sorting and storing data do not, in themselves, lead to increased understanding. Estimates suggest that only a fraction — perhaps even a small fraction — of the data generated by major scientific user facilities is ever analysed. If nothing else, ML will help to close this gap.

But cultural issues are important, too. Physicists are good at building hardware. It was not too long ago that the solution to every problem was made of stainless steel. Noisy data? Let’s build a brighter source, or a better detector. Experiments take too long? Let’s build a machine with a higher repetition rate. The timing of a signal not sufficiently precise? Let’s build a machine to measure it more accurately. For many experimentalists, software was used only to the extent needed to operate the hardware. This is no longer the case: the concept of hardware–software codesign is permeating the vernacular of scientific-instrument design. This development is a consequence of ML delivering the means to alleviate many problems algorithmically, without new hardware⁸. The question is who will take responsibility for algorithms? Will each roll their own, as at present, or will major user facilities provide data-analytical platforms? I believe such platforms are now as indispensable as detectors, which are routinely provided by user facilities.

More exciting is the possibility of designing entirely new experimental approaches based on ML tools. For example, rare but rate-controlling states of a system are currently investigated by collecting data as the system is driven through the states of interest. Such non-equilibrium ‘time-resolved’ experiments are notoriously hard to perform. But a sufficiently large dataset of a system in equilibrium contains snapshots of every state, with a probability determined by the Boltzmann factor. Given sufficiently large datasets and appropriate algorithms, it should be possible to capture rare, fleetingly occupied states from snapshots of systems in equilibrium^{9,10}, an altogether simpler proposition.

What these uses of ML have in common is that they allow experiments to be performed without theory dictating where to look or what to find. ML appears poised to end the era of tight experimental control. But perhaps the most radical consequence of ML for science is a fundamental revision of what constitutes scientific understanding.

- Ollivier, Y. Riemannian metrics for neural networks I: feedforward networks. *Inf. Inference* **4**, 108–153 (2015).
- Pascanu, R., Gulcehre, C., Cho, K. & Bengio, Y. How to construct deep recurrent neural networks. In *Proceedings of the Second International Conference on Learning Representations (ICLR, 2014)*.
- Alber, M. et al. Integrating machine learning and multiscale modeling — perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ Digit. Med.* **2**, 115 (2019).
- Giannakis, D., Schwander, P. & Ourmazd, A. The symmetries of image formation by scattering. I. Theoretical framework. *Opt. Express* **20**, 12799–12826 (2012).
- Hutson, M. How researchers are teaching AI to learn like a child. *Science* <https://doi.org/10.1126/science.aau2576> (2018).
- Heaven, D. Why deep-learning AIs are so easy to fool. *Nature* **574**, 163–166 (2019).
- Lafon, S., Keller, Y. & Coifman, R. R. Data fusion and multicue data matching by diffusion maps. *IEEE Trans. Pattern Anal. Mach. Intell.* **28**, 1784–1797 (2006).
- Fung, R. et al. Dynamics from noisy data with extreme timing uncertainty. *Nature* **532**, 471–475 (2016).
- Hosseinzadeh, A. et al. Conformational landscape of a virus by single-particle X-ray scattering. *Nat. Methods* **14**, 877–881 (2017).
- Dashti, A. et al. Trajectories of the ribosome as a Brownian nanomachine. *Proc. Natl Acad. Sci. USA* **111**, 17492–17497 (2014).

Acknowledgements

The author acknowledges valuable discussions with Tony Hey, Larry Jackel, E. Lattman and many UWM colleagues. Any errors are the sole responsibility of the author. This work was supported by the US Department of Energy, Office of Science, Basic Energy Sciences under award DE-SC0002164 (algorithm design and development), and by the US National Science Foundation under awards STC 1231306 (numerical trial models and data analysis) and 1551489 (underlying analytical models).

Competing interests

The author declares no competing interests.