

Between the Philosophy of Science and Machine Learning

David Corfield

University of Kent

Kevin Korb:

Machine learning should be regarded as "experimental philosophy of science".

Kevin Korb:

Machine learning should be regarded as "experimental philosophy of science".

Why should we expect philosophy to think up principles to be tested in Machine learning?

Philosophy has a tendency to lose control of its offspring, such as the group-theoretic basis to perception (from Kant to Helmholtz & Poincaré to Cassirer to Mallat?).

Kevin Korb:

Machine learning should be regarded as "experimental philosophy of science".

But for many,

History of science is a testing ground for philosophy of science.

Kevin Korb (revised):

Machine learning should be regarded as "experimental philosophy of science".

History of science is a testing ground for philosophy of science.

Machine learning could be a testing ground for formal philosophy of science.

Or maybe the association is looser.

Did Popper have the seed of a fundamental idea in statistical learning theory?

- VC-dimension: largest number of points which can be shattered
- Popper-dimension: smallest number of points which can't be shattered.

(see [Falsificationism and Statistical Learning Theory](#))

Popper close, but not quite there? Or was he on a different agenda?

- Induction is a myth
- Arrival at hypotheses is a matter of psychology
- As a matter of fact they tend to come from myth and metaphysics
- We don't just go about observing things in the world, we only do so in the context of a test.
- The testing of a theory is a matter of using *modus tollens*.
- We tend to have a bias to look for confirmation, but this is a childish weakness, like neurotics expecting the same patterns.
- Scientists (and adults) have had to learn to look to falsify their theories.

Perhaps we might say the 'metaphysics' of Popper inspired the 'science' of Vapnik, just as Robert Boyle drew inspiration from the atomism of Epicurus.

From a metaphysics of learning to a science of learning.

And the debt paid back to philosophy? As with Harman and Kulkarni?

A better account of the ML-Philosophy relationship might describe it as a '**dynamic interaction**'.
(Jon Williamson).

E.g., Judea Pearl meets David Lewis in the analysis of counterfactuals. A counterfactual is true if in the world closest to ours where the antecedent holds, so does the consequent.

Pearl doesn't buy the existence of possible worlds, but takes on the idea of minimal changes.

Rivals of Popper, the Vienna Circle and then the Logical Empiricists, such as Carnap, developed a probability based confirmation theory.

Bayesianism as the magic ingredient for learning theory?

Whistler 2006 conclusion: frequentists seem more inventive when it came to integrating Background knowledge: cluster assumption, invariance under small changes, robustness under adversarial deletion, structural correspondence (pivot words). Only the first two had been Bayesianised. (See my chapter in '[Dataset Shift in Machine Learning](#)')

Popper attacked Bayesianism

- Laws have zero prior by principle of indifference
- Good science looks for the unlikely not the likely

But there followed a raft of attacks on both.

Kuhn and Lakatos

Science is not in any important way algorithmic. It's about conceptual change. Not only the change of language, but changes to the way we should ask questions of the world. Bayesianism presupposes a fixed theoretical framework.

Popper just looked at the revolutionary part, badly. No-one, not even great scientists, ever throws away a theory that easily. Falsificationism is falsified itself by the history of science.

For Kuhn, we don't just change our theories, we change what we take to be worth answering. We can lose explanatory power. E.g., phlogiston theory gave a reason why metals are similar in a way that oxygen theory did not.

Kuhn replaces one uniformity with another. Copernican revolution is a pattern for all other revolutions.

Feyerabend

Galileo must resort to propaganda to have the new science heard. He has to gain the upper hand in describing what an observation is, i.e., that it concerns only relative motion.

Would anyone in the scientific revolution have benefited from a ML algorithm? Aristotle, Ptolemy, Galileo? Function estimation for Copernicus? Kepler had plenty of good data from Tycho Brahe. How would it arrive at the relative positions of points on different ellipses? (cf. BACON and the attempt to learn the power law between T and R)

The idea of breaking with circular motion was so profound that Galileo rejected it. He had so much invested in circular inertia (hence sketches of the moon with an erroneous circular crater on the midline?).

Perhaps philosophers of science invested too much in early episodes of science. Perhaps modern science is different and offers the opportunity for machine learning.

But further worries...

Ernan McMullin on metaphor and creative incompleteness (Jacob Bronowski too). (Occurs in Machine Learning, e.g., Mallat “Pushing the energy into low frequencies”.) Equivalent formulations of the same theory suggest very different extensions, e.g., Hamiltonian and Lagrangian reformulation of Newtonian dynamics.

Plenty of work done on varieties of scientific models, their role in between theory and data:

“Probing models, phenomenological models, computational models, developmental models, explanatory models, impoverished models, testing models, idealized models, theoretical models, scale models, heuristic models, caricature models, didactic models, fantasy models, toy models, imaginary models, mathematical models, substitute models, iconic models, formal models, analogue models and instrumental models are but some of the notions that are used to categorize models.”

<http://plato.stanford.edu/entries/models-science/>

'Models as Mediators' (Morgan and Morrison). Even here, Hartmann and the narrative of the model. Humans seem to want a story of what's going on.

Polanyi's tacit knowledge and the scientist's passionate anticipation of “an indeterminate range of yet unknown (and perhaps yet inconceivable) true implications”.

“The kind of knowledge I have of my body by dwelling in it is the paradigm of knowing particulars subsidiarily with a bearing on the comprehensive entity formed by them. Hence when I rely on my awareness of particulars for attending to a whole I handle things as I handle my body. In this sense I know comprehensive entities by indwelling their functional parts, as if they were parts of my body. Such is my conception of knowing by indwelling.”

“The Holy Grail of integrating automated reasoning across all relevant representations and processes seems far from current reality. This is in no small part due to our continuing ignorance of the creative human thought processes guiding the art of doing science...It will be a fascinating and instructive endeavor – requiring contributions across technology, science, and even philosophy – to develop and understand the full spectrum of such inference systems.” (Mjolsness and DeCoste, *Machine learning for science: state of the art and future prospects*, 2001)

But is all this just too high level? Why mimic the top end of learning, the natural scientist? Perhaps we shouldn't expect ML to be there for the revolutions, more for the normal science. Even so, we need to know more about scientific learning.

Dudley Shapere, critic of historical philosophers projecting from their case studies.

There is a history of scientific learning: Metaphysical; piecemeal domains (falling bodies, salts, electrical phenomena); unifying domains (Whewell consilience), e.g., chemical bonds, static electricity. It could not be foretold that these strategies would work.

Where we were once satisfied with two astronomies (Greek crystal spheres and ptolemaic epicycles), we had to learn (Shapere with McMullin)

- that testable theoretical postulation was worth doing,
- to trust instruments (against Aristotle and adequacy of senses)
- to use mathematics (against Aristotle),
- to use a form of experimental method (against Hobbes, and Aristotle on natural rather than violent motion),
- not to expect the Bible to be a guide to scientific truth (against Newton).

Dudley Shapere (cont.)

We have learned how to learn. There is no ahistorical split between the scientific (forces, fields, genes) and the metascientific (observation, theories, explanation, laws).

E.g., desirable patterns of explanation change: Newton's gravitational force was unsatisfactory to Cartesians (and himself); Hershel's *verae causae*, and natural selection as "the law of Higgledy piggledy"; the *gruppenpest* in physics (coming soon to Machine Learning?); Hilbert's *Theology* (Gordon's description of proof of *Nullstellensatz*).

But is there not a single account of explanation, or of causation? Perhaps not a universal account, but enough local pockets of meaning?

Causality as difference making (counterfactual, agential – ideal intervention, statistical) or productive (continuous transmission of quantity, mechanism). Can absences or omissions be causes?

“In cases of causation between absences, production is unnecessary; in cases of redundant causation, difference-making is unnecessary; and there may be cases that involve neither production nor difference-making” (Schaffer)

Causality in the Sciences, Edited by Phyllis McKay Illari, Federica Russo, and Jon Williamson, OUP, 952 pages.

Shapere again:

“...the problems we face in our inquiries about nature, and the methods with which we deal with those problems, co-evolve with our beliefs about nature.”

Might we learn how to use machines to learn? Might new methodologies emerge allowing for further automation?

As ever, the role of background knowledge is key:

"Even when prior beliefs are mentioned, their role or roles in new inquiry tend to go undiscussed; even for the hypothetico-deductivists, they may have been "verified" or "confirmed," but their roles in a large variety of the most important functions of scientific inquiry are left out of account, indeed receive no mention."

We should be answering the following questions:

- (1) What considerations (or, better, types of considerations, if such types can be found) lead scientists to regard a body of information as a body of information - that is, as constituting a unified subject matter or domain to be examined or dealt with?
- (2) How is description of the items of the domain achieved and modified at sophisticated stages of scientific development?
- (3) What sort of inadequacies, leading to the need for further work, are found in the bodies of information, and what are the grounds for considering these to be inadequacies or problems requiring further research? (Included here are questions not only regarding the generation of scientific problems about domains, but also scientific priorities - the questions of importance of the problems and of the "readiness" of science to deal with them.)

(4) What considerations lead to the generation of specific lines of research, and what are the reasons (or types of reasons) for considering some lines of research to be more promising than others in the attempt to resolve problems about the domain?

(5) What are the reasons for expecting (sometimes to the extent of demanding) that answers of certain sorts, having certain characteristics, be sought for those problems?

(6) What are the reasons (or types of reasons) for accepting a certain solution of a scientific problem regarding a domain as adequate. (Shapere 1984, 277-8).

He reckons that only the last, and to some extent (2), have been seriously examined by philosophers of science.

“Surely these are central features of the scientific enterprise; to ignore them is to ignore the heart of science. But to deal with them successfully it is necessary to recast the way philosophy of science is done, shifting away from the essentialist, atemporal methodology that has too often characterized it, to focussing rather on the ways in which science moves from one stage to another in its inquiries—on its evolution.”

In current science, there has been internalised a huge number of background beliefs. With these in place we can talk of 'the given' (data). But it is not something read simply off the world. Never forget our shaping of the data.

Let's look to see where ML techniques can be used:

Data Mining and Machine Learning in Astronomy, Ball & Brunner, <http://arxiv.org/abs/arXiv:0906.2173>

star/galaxy

galaxy morphology

detection of quasars

“We conclude that, so long as one carefully selects an appropriate algorithm, and is guided by the astronomical problem at hand, data mining can be very much the powerful tool, and not the questionable black box.”

To conclude:

There is no indication that we should expect some universal method, but there's a possibility for local methodologies using ML to emerge. There would need to be a pairing between what aspects of the theory of the domain being studied are accessible and the ML technique.

Some domains won't easily lend themselves, e.g., Mathematics. Nor some disciplines at certain stages.

Perhaps theories could co-evolve with ML methodologies.

There's some potentially useful philosophy literature (modelling, causality, explanation, philosophy of specific sciences)