



Harry Collins (2011)

Simulating What?

Journal of Artificial Societies and Social Simulation 14 (4) 16

<<http://jasss.soc.surrey.ac.uk/14/4/16.html>>

Received: 14-Jul-2011 Accepted: 15-Aug-2011 Published: 31-Oct-2011

Abstract

Any attempt to simulate science has first to say what science is. This involves asking three questions: 1) The Scope Question: What bit of science is the target? It is immensely confusing (as the history of these debates shows), if one simulates some little aspect of science, as in the case of BACON, and then claims that one has built a machine that can 'do science'. 2) The Micro-World Question: Is the criterion of success the reproduction of human science—with all the same findings turning up—or the simulation of something that is believed to be a scientific process with results that pertain only to the world of the simulation which do not correspond to the outcome of human science as we know it? If the latter it will be important to be sure that one is not merely developing a 'micro-world'—a world so tidied up for the purposes of simulation that it does not bear on human science. 3) The Chess Question: Even if the idea to reach the same results as has been reached by human science, does it have to be by 'the same' means in order to count as a simulation of human science? I call it the 'chess question' because Deep Blue does not play in the same way as human grand masters but is still better at winning.

Keywords:

Science, Language, Demarcation, Micro-World, BACON, Chess

Simulating What

- 1.1 If one is going to simulate science one has to know what science is. Well, we do know what science is—we must know because we are always talking about science and scientists, criticising them, admiring them, trying to emulate them, trying not to emulate them, and so forth. On the other hand, no-one has ever come up with a satisfactory 'recipe' for science—a description that would allow one to say "by referring to how science is made I can see that 'this' is science and 'that' is not science." The problem of finding the recipe for science is often referred to as the 'problem of demarcation' and, as far as I know, it has not been solved in a way that provides a set of defining rules.
- 1.2 There is no paradox here. One can know what a thing is without being able to define it in a crisp way. The most famous example of this point is Wittgenstein's (1953) discussion of 'game'. We all pretty well know what is a game and what is not a game but we cannot provide a set of defining criteria. We know we know what a game is because we know when something is not a game. We know that war is not a game even though we know there are such things as war-games—we know the difference between war-games and war. We understand that fictional games, such as that portrayed in the film 'Roller Ball', are not quite games. We know that Roman gladiatorial contests were not quite games either. We know how to use the phrase 'play the game'. And so on. Still, the idea of 'game' resists any attempt to provide a set of defining criteria which will include all games and exclude all non-games.
- 1.3 'Science' appears to be like 'game'. Every attempt by philosophers to demarcate it from other activities has failed. The attempt to define it as, 'that which takes meaningfulness to be verification by observation', failed because almost nothing could be verified—the criterion was too strong. Defining it as, 'that which takes scientific laws as those which can define the conditions for their own falsification', failed because it is impossible to be sure that something has been falsified. Defining 'scientifically proven' as 'demonstrable by repeated experiment' turns out to be unsatisfactory because it is impossible to know for sure when an experiment has been properly repeated (Collins 1985). One might say that the whole burden of the subject known as 'Sociology of Scientific Knowledge' has been to show that anything anyone thinks science is, it isn't—at least, there are instances of what we call science that do not fit with any closed definition of science.
- 1.4 Wittgenstein's solution in the case of 'game' was the idea of 'family resemblance'. There were rough criteria defining sets of

games which had an overlap but not congruence with the criteria defining other sets of games, and so on, until all sets of games were exhausted. This can be visualised as an overlapping set of circles or ovals. In Figure 1 the ellipse at the left hand end might represent the set of games like soccer and that at the right hand end might be the set of games like dwile-flonking.




Figure 1. The idea of family resemblance

- 1.5 If one examines Figure 1 carefully, however, one will note that it has been drawn so as to gloss over an ambiguity. It is hard to see whether the dwile-flonking ellipse—one extreme end of the family—has any overlap at all with the soccer ellipse—the other extreme end. The trouble is that if one allows that there is no overlap at all then the family resemblance relationship could go on for ever absorbing every activity in the universe through overlapping criteria. For example, 'paintball' is a game that involves hitting people with projectiles and execution by firing squad also involves hitting people with projectiles so execution by firing squad has a family resemblance to paintball and by extension, dwile-flonking has a family resemblance to nuclear war and therefore to the formation of the universe (which also involved sub-atomic reactions). It seems that to avoid this problem there must be some area in common among all the ellipses in the family resemblance diagram so as to prevent the ellipses extending ever. But if one has an area common to all the ellipses then one has defining criteria—they are found in the common area—so one is returned to one's starting point.
- 1.6 If one was being charitable one would say that this ambiguity was cynically exploited by Steve Fuller in his evidence for the defence at the trial of the Dover, Pennsylvania, School Board when, in 2005, they were prosecuted for teaching intelligent design in science lessons (<http://www.talkorigins.org/faqs/dover/day15am2.html> accessed 15 Aug 2011). (If one was being less charitable one would say that Fuller was confused.) Fuller argued that since science and religion had overlaps in Newton's time one could not say that science was incompatible with religion nor that it might not come to have overlaps with religion once more and that, therefore, it was not clear that in teaching intelligent design, even if it was essentially a religious doctrine, the School Board was obviously doing something unscientific. The paradox, which ought to have been obvious, is brought out by the example of the fact that alchemy was also part of science in Newton's time so the argument could justify the contemporary teaching of alchemy or, by extension, pretty well anything else—such as that the world is entirely water. Fuller could equally argue that gladiatorial combat to the death is a game since the Romans seemed to treat it as such. The idea of family resemblance has to be kept under constraint if it is to be any use but the trouble is that it is hard to formulate the proper constraints—the idea doesn't really solve the problem of defining criteria, it just makes it a bit easier in our minds about a certain class of problems of definition.



Some previous attempts to simulate science

- 2.1 Given that it is so hard to say what science is, it is a bit of surprise that there have already been attempts to simulate it. I cannot claim to be any kind of expert on these matters but two such attempts have come to my attention. One was carried forward for many years by the late David Gooding and colleagues (e.g. Gooding 1990). They attempted to reconstruct, in the form of a computer program consisting of general rules, the actions of Michael Faraday. They used as their source their unprecedentedly detailed understanding of Faraday's laboratory notebooks. I do not know if what they did should be accounted a successful simulation of Faraday but, if it was, it was not a simulation of science but the simulation of the actions of one particular scientist. This is a project that is far less subject to the problem of demarcation.
- 2.2 The other attempt that I know something about is Herb Simon's 'general problem solver' and its particular manifestation 'BACON'.^[1] BACON, it was said, was a computer program that had, among other things, deduced Kepler's Laws of planetary motion from the data. At best this would mean that the general problem solver was not so general but was capable of simulating a subset of scientific activities belonging to the physical sciences (and geometry). And it was. But Simon and his colleagues did not seem to realise how narrow that subset of activities was. They had, in effect, reduced 'science' to a computational problem: give a computer program a set of clean data, and get it to generate a set of mathematical laws that fit. This, undoubtedly, is a part of science—a part that human scientists find extremely hard and that is generally recognised as symptomatic of human genius—but it is no more the whole of science than factor analysis is the whole of science. Factor analysis is another part of science that is highly demanding for humans and that computers can do, and do better. The same, of course, goes for the program 'Mathematica', which nowadays solves equations, integrates functions, and so on, which tasks were once accounted to be one of the hallmarks of scientific virtuosity. But, then, even rapid, fault-free, arithmetic, such as can nowadays be done by the cheapest pocket calculator, was also once thought of as symptomatic of great scientific promise in a boy or girl.
- 2.3 BACON was not simulating what Kepler and all the equivalents of Kepler do, only that calculating bit that sits in the middle. The first part of anything Kepler-like is the selection of data from noise—a matter of judgment. Only once you know what is the data and what is the noise can you let a BACON-like program free on the data. The iconic example of the noise/data problem is

Millikan's oil drop experiment, Millikan choosing to ignore those apparent data points that did not fit the notion that the charge on the electron was integral (Holton 1978). The last part of anything Kepler-like is the judgment about what is going to be credible to your contemporaries—another quintessentially social judgment which, of course, feeds back to the initial data-noise separation. Millikan must have judged he could 'get away' with data-filtering decisions the final outcome of which confirmed the integral charge idea. His rival Eherenfest, it seems, had data that justified his non-integral charge view (as did Millikan), but Millikan thought, correctly, that data assembled according to the integral charge view would win the day. We don't know what a BACON-like machine given all Millikan's and all Eherenfest's data would have come up with.



Solving the demarcation problem

- 3.1 In the book *Rethinking Expertise* (Collins and Evans 2007), the demarcation problem had to be solved. The authors of the book wanted to provide some guidelines for policy-makers on how to choose between competing sets of advice that bore upon technological problems. They wanted to recommend that scientific advice was to be taken to be better than other kinds of advice where 'propositional questions' were involved. This meant that they had to say what science was.
- 3.2 The solution adopted was family-resemblance but with a modification: scientists proper, according to the new demarcation principle, have to be informed by the intention of trying to persuade all other scientists within the relevant domain to their view. A scientist whose aim is wholesale revolution—that is the replacement of existing science with an entirely different set of institutions, laws, cultural practices, etc, is not a scientist. Even the most 'revolutionary' scientists, such as Albert Einstein, Joe Weber (who claimed to have discovered 'impossibly' large fluxes of gravitational radiation and refused to give up the idea, Collins 2004), or any serious parapsychologist, are extremely conservative in respect of scientific institutions: their goal is to be accepted by the existing body of scientists, not to replace them. The proponents of intelligent design are ruled out of science in virtue of the fact that they want a revolution that would make revelatory books of obscure origin far more important in the institutions of science than they currently are and that would provide a cut-off for scientific inquiry (the intelligent designer), whereas science as we know it is characterised by a refusal to accept any end to scientific inquiry other than a scientific solution (also making the many-worlds hypothesis or the anthropic principle very much on the borderline of what counts as scientific explanation). The modification to the idea of family resemblance 'tames' the potential disastrous 'inflation' of any family such that it fills the entire universe—it slows down the spread of the family and keeps the family pretty local.



The simulation of science

- 4.1 The main characteristics of the science family are fairly prosaic. Once it is allowed that one is thinking of the actions and intentions that characterise an overlapping set of practices, rather than of logically immaculate definitional rules, all the old philosophical criteria come back in but in weakened form. Science, often, though not absolutely always, places observational evidence above all else when it comes to generating a description of the world; it always values this above the revelations of books of obscure origin. In the same way, scientists should mostly aspire toward generating theories that carry with them a means for their own falsification. The Mertonian norms: the encouragement of constructive criticism, universalism, freedom of information, and the aspiration toward disinterestedness also help one to recognise when science is happening though they are no longer to be thought of as defining science (Merton 1942). Another new criterion introduced in *Rethinking Expertise* is that the 'locus of legitimate interpretation' of a piece of work, is closer to the producer in the sciences than in the arts and humanities. It follows that the duty of a scientist is to try to be clear and unambiguous—to minimise the scope for alternative interpretations of the work irrespective of the extent to which this can be achieved in practice. Other rules are that the scientists should have integrity. Of course, all these characteristics, except, perhaps the integrity rule, can be violated in certain circumstances and most of them apply, in greater or lesser degree, in institutions other than science. So, if it is science as a whole that is to be simulated, the goal is pretty vague.
- 4.2 But even if the vagueness was not thought to be an obstacle, right at the heart of the project is the following central difficulty which is implicit in much of what has been said. To do science one has to be thoroughly embedded in the human society of other scientists. It has been seen that one has to be embedded in order to know what counts as data and what counts as noise because that decision is tied up with the scientist's judgment about what counts as a credible result. One has to be embedded in order to know what counts as a suitable balance with conservatism in institutional and cultural change when radicalism in cognitive change becomes a necessity. One has to be embedded in order to know when a violation of the normal aspirations of science is acceptable and when it goes too far. And, not so far mentioned but implicit in the sociology of scientific knowledge, one has to be integrated into society to know when to believe a result and when not to believe it.
- 4.3 In this respect, the simulation of science is presented with exactly the same, so far unsolved, problem that faces every artificial intelligence project. It is a problem that is not only unsolved but the solution to which is currently unforeseeable. This is the problem of the simulation of polymorphic actions (Collins and Kusch 1998), or, if one prefers, the problem of the formalisation or absorbing of 'collective tacit knowledge' (Collins 2010). The usually overlooked or ignored problem can be seen even in the smallest and most trivial of AI projects—say creating an automatic spell-checker. The spell-checker of the machine on which I am typing this paper is about to flag-up a word with a jagged red line: jaged. That's right, there is no such word as jaged but jaged is what I want to write because it makes sense in this context. (Actually the spell-checker was even more irritating, replacing each instance of jaged with jagged until I went back and altered it again.) To know enough to leave jaged as jaged one has to

understand the meaning of this paragraph and that is more than any spell-checker can do. And, of course, if you tell me, 'add jaged to the permissible spellings in the spell-checker', then the spell-checker won't work properly in any future paragraph where I accidentally write jaged meaning to write jagged. But we still haven't got to the heart of the problem. This is that the permissible spelling of words changes not only in local context but over time as language changes in society as a whole. This is why Searle's Chinese Room won't work—it can capture a frozen moment of language at best, not the changing flux of language that responds to changing social events in the wider world (Searle 1980). To work as a human editor, the spell-checker needs to be just as firmly embedded into society as the human editor and we simply do not know how to make it so; we do not even know how we humans do it beyond 'hanging around with other humans'—the mechanism of what goes on when we do hang around, we just don't understand.



Recommendations

- 5.1 I have to repeat my confession that I do not know anything about the field of simulating society so I may be pushing at an already open door. If I am, then that's great. If I am not then I have three recommendations when it comes to simulating science:
1. The 'scope question' has to be asked: In so far as attempts are made to simulate science the authors should be very clear about exactly what bit of science is the target. It is immensely confusing (as the history of these debates shows), if one simulates some little aspect of science, as in the case of BACON, and then claims that one has built a machine that can 'do science'.
 2. The 'micro-world question' has to be asked: Is the criterion of success in the case of simulation is the reproduction of human science—with all the same findings turning up—or the simulation of something that is believed to be a scientific process with results that pertain only to the world of the simulation which do not correspond to the outcome of human science as we know it. If the latter it will be important to be sure that one is not merely developing a 'micro-world'—a world so tidied up for the purposes of simulation that it does not bear on human science.
 3. The 'chess question' has to be asked: Even if the idea to reach the same results as has been reached by human science, does it have to be by 'the same' means in order to count as a simulation of human science. I call it the 'chess question' because Deep Blue does not play in the same way as human grand masters but is still better at winning.^[2]



Notes

¹ For BACON see Langley et al. (1987) and for a critique see Collins (1989).

² An awkward question is posed by Otto Sibus's finding that, unknowingly, the thermal signature of Joule's own body was an element in his pioneering research on the mechanical equivalent of heat—the foundation of thermodynamics.



References

- COLLINS, H M (2010) *Tacit and Explicit Knowledge*. Chicago: University of Chicago Press [doi:10.7208/chicago/9780226113821.001.0001]
- COLLINS, H M (2004) *Gravity's Shadow: The Search for Gravitational Waves*. Chicago: University of Chicago Press [doi:10.7208/chicago/9780226113791.001.0001]
- COLLINS, H M (1989) Computers and the Sociology of Scientific Knowledge, *Social Studies of Science*, 19, pp. 613-24. [doi:10.1177/030631289019004003]
- COLLINS, H M (1992) *Changing Order: Replication and induction in scientific practice* Chicago: University of Chicago Press (first published 1985)
- COLLINS, H M and EVANS, R J (2007) *Rethinking Expertise*. Chicago: University of Chicago Press [doi:10.7208/chicago/9780226113623.001.0001]
- COLLINS, H M and KUSCH, M (1998) *The Shape of Actions: What humans and machines can do*, Cambridge Mass.: MIT Press
- GOODING, D (1990) *Experiment and the Making of Meaning*. Dordrecht: Kluwer [doi:10.1007/978-94-009-0707-2]
- HOLTON, G (1978) *The Scientific Imagination*. Cambridge: Cambridge University Press.
- LANGLEY, P, SIMON, H A, BRADSHAW, G L, AND ZYTKOW, J M (1987) *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, Mass: MIT Press
- MERTON, R K (1942) Science and Technology in a Democratic Order. *Journal of Legal and Political Sociology*. 1 pp. 115-26.

SEARLE, J (1980) Minds, Brains and Programs. *The Behavioural and Brain Sciences*. 3, pp. 417-24 WITTGENSTEIN, Ludwig (1953/2001). *Philosophical Investigations*. Blackwell Publishing