Philosophy in Review

Hector J. Levesque, "Common Sense, the Turing Test, and the Quest for Real AI."

Sheldon Richmond

Volume 41, Number 1, February 2021

URI: <https://id.erudit.org/iderudit/1076213ar>

DOI: <https://doi.org/10.7202/1076213ar>

[See table of contents](#)

Publisher(s)

University of Victoria

ISSN

1206-5269 (print)

1920-8936 (digital)

[Explore this journal](#)

Cite this review

Richmond, S. (2021). Review of [Hector J. Levesque, "Common Sense, the Turing Test, and the Quest for Real AI."]. *Philosophy in Review*, 41(1), 25–28.
<https://doi.org/10.7202/1076213ar>

Copyright © Sheldon Richmond, 2020

This document is protected by copyright law. Use of the services of Érudit (including reproduction) is subject to its terms and conditions, which can be viewed online.

<https://apropos.erudit.org/en/users/policy-on-use/>

érudit

This article is disseminated and preserved by Érudit.

Érudit is a non-profit inter-university consortium of the Université de Montréal, Université Laval, and the Université du Québec à Montréal. Its mission is to promote and disseminate research.

<https://www.erudit.org/en/>

Hector J. Levesque. *Common Sense, the Turing Test, and the Quest for Real AI.* The MIT Press 2017. 192 pp. (Hardcover ISBN 9780262036047); \$19.95 USD (Paperback ISBN 9780262535205).

Another book about AI, but not a book that is AI jargon replete, nor a book that is AI hype, nor a book full of doom and gloom warning of the impending catastrophe of mass unemployment and the enslavement of humanity. Rather, *Common Sense, the Turing Test, and the Quest for Real AI* (*Common Sense*) addresses a general audience and presents a straightforward argument for AI as science, or at least, as theory, as opposed to AI as technology. Indeed, we find out in the last chapter that Hector Levesque is wary of AI as technology, because according to Levesque, AI technology, is neither reliable nor predictable. Moreover, for Levesque, the strength of ‘Real AI’ is the theory part of AI. But not just any old AI, a specific type of old AI, nicknamed, GOF AI: good old-fashioned AI. However, the current leader in the field of AI, even according to Levesque, is what has been variously dubbed, Neural Nets, Deep Thinking, and Deep Learning. Levesque's goal is to reveal that the new royalty in AI, the Deep School (my term), is both naked and a pretender to the throne. If we want real AI, we have to return GOF AI to the throne, even though GOF AI has limited technological results, unlike the Deep School. The Deep School is heralding innovations in technologies such as self-driving cars, medical diagnostics, drug discovery, and facial recognition. These Deep School technologies are not merely in development, but are in the prototype testing stage by Amazon, Google, Facebook, and IBM. Other new Deep School technologies in the security, military, government, cloak and dagger field are in testing and are kept under wraps.

The question that an intelligent reader who is blunt and direct (I mean human, not one of the apps that one can download for one's personal smart device that summarizes large documents) might ask: does Levesque's book amount to an attempt to revive a dead horse? Is Levesque, as a prominent second generation GOF AI researcher and disciple of John McCarthy, the pioneer of GOF AI, falling sway to the spell of sunken costs, and continuing to invest more effort and time in a now bankrupt project where he already invested an entire career? In an ideal model of science, and Real AI, is supposedly a science, once a scientific theory has been falsified, the intellectually honest scientist rejects the theory and looks for a new theory. Has GOF AI been falsified? If not, has GOF AI, according to the Kuhnian model, been replaced by a new paradigm that has been adopted by the next generation in the early stages and mid-stages of their career? Or, more realistically, which is to say, more materialistically and commercially, have the companies who want to develop systems based on AI, including privately and publicly funded research institutes, found that the Deep School approach is a magnet for venture capital? Furthermore, have developers of the Deep School become tied in with the private and public agencies that are the hunters and gatherers of Big Data? The Deep School quietly claims, unlike GOF AI, to be able to mine Big Data for financially valuable market information, as well as information relevant to security, policing, military, and political interests. Hence, the magnetic charm of the Deep School with its ability to carry on with Big Data, has turned GOF AI into an ageing has-been, only able to find supporting roles, if any. (By the way, the main leader of the Deep School, Geoffrey Hinton, came to Canada from the US because he feared that military interests were taking advantage of his research. But is Hinton somewhat naïve in believing that he can sidestep military interests after the military giant has become aware of potential Deep School products delicious to military hunger?) Does that socio-commercial-materialist model tell the story of how the Deep School has become the winning player?

All the above is background. There is one more background item, that deserves the attention of those interested in intellectual history, as well as the social history of ideas, science, and the humanities. There is an untold story by both Levesque, prominent defender of GOF AI, and the advocates of Neural Nets (many in the Vector Institute in Toronto, Canada). The story is that GOF AI in its infancy replaced cybernetics in its attempt to, both in theory and in technology, develop machine simulations of brain functioning for creating artificial intelligence machines. Ross Ashby as well as Norbert Wiener, the pioneers of cybernetics, disappeared from the history of artificial intelligence in the late 1950s when John McCarthy launched his research program for developing artificial intelligence at Dartmouth College, and then later at MIT. This program fit into the dominant philosophy of the day: language is the beginning, middle, and boundary of human understanding and intelligence. So, then, artificial intelligence must simulate the workings of language, or more generally, symbolic systems. The brain is a black box regarding how thinking and intelligence works. But language and all symbol systems form the public, transparent, and social framework and functioning of how intelligence and understanding work. However, the tables turned and returned, though tacitly, to the cybernetic program when neuroscience took off and new tools such as the MRI allowed the neuroscientist to pinpoint cognitive and other brain functions. Neural researchers saw neural functioning as multi-layer and parallel: neurons formed and reformed multiple associations in parallel among themselves depending on both prior learning (training) and new sensory input. The Deep School created itself out of GOF AI by developing at first symbolic approaches that they realized were similar to brain functioning as multi-level or multi-layer and parallel. Indeed, without acknowledgement to cybernetics and its founders, Ross Ashby and Norbert Wiener, the Deep School tacitly carried forward the early and once abandoned research of the pioneers of cybernetics: the Deep School took the idea that machines could be developed along the model of brain functioning that not merely simulated but ‘amplified’ (in Ross Ashby’s terminology), the human brain.

Here enters Levesque’s with his argument and rationale: with the turn (and as I see it, also return) to modelling brain function, the Deep School forgets its own starting point as a deviant version of the by-then mainstream theory that human intelligence at core involves the use of symbolic systems. Hence, biological condition for symbolic processing is the brain, the design and functioning of symbol systems for intelligence, artificially, can use other material conditions: electrical, mechanical, biological, physical: machines, bio-materials, buildings and cityscapes. What matters is the form, rather than the matter: biological, physical, social. The form amounts to how we use symbolic systems. In other words, brain function amounts to the material workings of intelligence that operate with the abstractions of abstract symbolic systems: verbal languages for communication; symbols systems for encoding the languages of music, speech, dance, the visual arts; symbol systems for encoding computation and mathematics.

The short of it is that intelligence is at core a function over the abstractions encoded by symbolic systems. The question arises: where does the computer fit in? The computer seems to be at the core of whatever idea we have of AI. If AI amounts to functions over symbol systems, how is it that the computer and AI seem to go hand in hand? Levesque’s answer to this very logical question is: it is not the physical computer in and of itself, but what the computer represents, that goes hand in hand with AI. It is the computable as conceived by Alan Turing: it is how we determine in principle, whether anything—a number, a mathematical conjecture, an algorithm, or intelligence is computable. It is not whether we can build a real computer to determine the computable, though we now can, but whether in principle a model of the computable—a design for a machine to carry out calculations, such as a Turing machine and the Universal Turing machine, can also compute intelligence as well as numbers or digits. Moreover, a differently designed abstract machine such as John von

Neumann's automaton can in principle carry out computations, and so if that machine can as well compute intelligence, it then can function as an AI machine.

The question now becomes: if intelligence at its core involves symbolic systems, and AI must involve the computable, how are symbolic systems computable? The answer is: Symbolic systems are computable with respect to, firstly, the formal aspects of symbolic systems, propositional logic (chapter 8); and secondly, the semantic aspects of symbolic systems as in Winograd schemas (54 ff.). Let's go back a step and ask along with Levesque, what is the importance of human intelligence? What do humans achieve with intelligence, apart from solving practical problems, that is significantly human? Levesque's answer is: knowledge (chapter 3). Humans use their intelligence, which at its core is the use of symbolic systems, to achieve knowledge. This then is the focus of the book in its defence of GOF AI as opposed to the Deep School: how do we design artificial intelligence machines that in theory or principle can achieve knowledge? The Deep School in its concern with modelling brain function overlooks the purpose of intelligence, and models how the brain learns, or filters and adapts to the rivers of data that flood the brain. Levesque's challenge to the Deep School that has replaced GOF AI as the mainstream in AI is: how can automated machine learning or AML (what Levesque calls the Deep School) explain the development of knowledge that is represented by humans in terms of symbolic systems? In other words, the Deep School, or AML, overlooks the significance of human intelligence; not merely learning how to sift through data and solve piecemeal practical problems; but to develop systems of knowledge represented in symbolic systems. Levesque's critique of the Deep School, is that it shifts the problem from how to create (designs for) machines that artificially produce knowledge and that artificially use common sense, to how to create both designs and working prototype machines that sift through Big Data for solving Big Data puzzles concerning analysis (or what data mining companies prefer to call "analytics"), diagnostics, pattern and face recognition:

the move away from GOF AI is more like a shift in subject matter...researchers in that part of machine learning we have been calling AML [Automated Machine Learning or Deep Learning] might focus on our ability to recognize hand written-digits... Researchers in AML concentrate on showing how the necessary patterns and features can be learned from these samples in an automated way. (62)

The reader might wonder, so what? Changes in models, paradigms, research programs often result in changes in subject matter, in problems, in questions. What's so terrible about AML having shifted the problems and questions, the subject matter of AI? For instance, GOF AI itself depended for its development on a shift in approach to logic, mathematics, and epistemology that had traditionally focused on the problems of empirical or rational foundations. The shift in the problems (in logic, mathematics, and epistemology) was to the development of the new problems of formalization, logical analysis, and formal provability within artificial axiomatic, symbolic or formal systems, such as mathematical set theory (the so-called, 'linguistic turn'). One can provide multiple examples in the sciences and in the humanities: a change in approaches involve a change in subject matter or problems and research questions. I leave it to the reader to pick their own favourite example, for instance those interested in art history who no longer tell The Story of Art (apologies to E.H. Gombrich), but tell stories of artists, museums, schools, cultures.

However, to be fair, Levesque does respond in specific to the question of why making knowledge is central to AI, as opposed to developing AI for Big Data diving: the program of AI is really to understand how people are intelligent, and the purpose of artificial intelligence, is to explain intelligence by developing designs for abstract machines that function intelligently. Humans function

intelligently through the use of common sense and knowledge: Since it is a fact of life that humans need knowledge, then designing knowledge-based AI (GOF AI) + ‘is really the only game in town.’ (126)

Sheldon Richmond, Independent Scholar

*I acknowledge Dr. Patrick Gamez, Assistant Professor of Philosophy, Department of Art, Languages, and Philosophy, Missouri University of Science and Technology for his incisive comments on a draft of this review (though he bears no responsibility for any remaining defects)