Published in Kristof Nyiri, ed., *Envisioning an Electrifying Future, 11th Budapest Visual Learning Conference*, Budapest, November 2024, 19-26

LLMs and practical knowledge: What is intelligence?

Barry Smith

In November 2022 I published with my co-author Jobst Landgrebe, a German AI expert, computer scientist, philosopher, and biomathematician, a book entitled *Why Machines Will Never Rule the World*.[1] The book argues that so-called 'Artificial General Intelligence' (AGI) will for mathematical reasons never be achieved. In the same month, ChatGPT was unleashed onto the world, and while our book has been well-received, many times reviewed, and many times the subject of entertaining podcasts, we have received little response from AGI proponents.

To our surprise, however, our publisher invited us already in April 2024 to prepare a revised and expanded 2nd edition of the book. Nothing has changed in our arguments since the time Before ChatGPT. However, we have taken the opportunity to respond to the many claims made by self-declared technophilosophers such as David Chalmers and by AI-entrepreneurs such as Sam Altman to the effect that AGI is, if not *already here*, then at any rate `just around the corner'. We have also monitored the reactions to such claims by others, many of them emanating from the admirable Gary Marcus, who focuses not only on AI hype, but also on the devastation that is being wrought by ChatGPT and other large language models (LLMs) in their march through the institutions of, for example, science.[2]

The Bet

In April of this year Elon Musk predicted that an AGI – by which he meant an *artificial intelligence superior to the smartest individual human* – would arrive by the year 2025. In response to this prediction Gary Marcus

[1] Jobst Landgrebe and Barry Smith, *Why Machines Will Never Rule the World*, Abingdon UK: Routledge, 2023.
[2] Gary Marcus, "The exponential enshittification of science", https://garymarcus.substack.com/p/the-exponential-enshittification, March 15, 2024.

offered Musk a $1 million bet to the effect that he would be proved wrong. In specifying the conditions of this bet (which Musk did not take) Marcus lists the following 'tasks that ordinary people can perform' which, he claimed, AI will not be able to perform by the end of 2025.[3]

- Reliably drive a car in a novel location that they haven't previously encountered, even in the face of unusual circumstances like hand-lettered signs, without the assistance of other humans.
- Drive an off-road vehicle, without maps, across streams, around obstacles such as fallen trees, and so on.
- Learn to ride a mountain bike off-road through forest trails.
- Babysit children in an unfamiliar home and keep them safe.
- Tend to the physical and psychological needs of an elderly or infirm person.

This list provides interesting insights concerning the limits of (current) AI systems. First, it reminds us that there is a serious lag on the side of tasks involving physical behavior (thus a lag on the side of robotics) as compared with the in some ways impressive progress being made on the cognitive side by, for example, LLMs. But secondly, and more importantly from our point of view here, the tasks we have listed involve *practical* or *tacit knowledge*, or what is also called 'knowing how', which is to say knowledge of a sort that is captured not by means of sentences or propositions or explicit rules, but rather through the expertise demonstrated in human actions.[4]

Knowing How

There is a rich literature in philosophy and psychology concerning the topic of practical knowledge to which major contributions have been

[3] https://garymarcus.substack.com/p/superhuman-agi-is-not-nigh

[4] Even speaking a language is an example of knowing how. The letter combination 'th' in English represents two phonemes, /ð/ and /θ/: *voiced,* as in 'this'; and *voiceless*, as in 'thing'. English speakers deploy this distinction effortlessly when they speak; but few of them are aware that they are doing so, or of the rules they are thereby following.

made by thinkers such as Scheler and Gehlen, by Ryle – who introduced the idea of a dichotomy between *knowing how* and *knowing that* – and by Polanyi, Merleau-Ponty, and Hubert Dreyfus.[5] Yet even though some 100,000 papers on the topic of AI have appeared on the arXiv.org preprint server since 2022, many of them drawing on philosophical aspects of AI, not one of these papers has addressed the issue of *practical* or *tacit knowledge* or 'knowing how'.

Defining 'intellectual task'

The proposed bet between Musk and Marcus concerned the issue of the possibility of AGI, which we can define as: *the capacity to understand or learn any intellectual task that a human can*. What, now, is the meaning of the phrase 'intellectual task' in this definition? For some (very few) such tasks – language production and interpretation, playing championship-level chess and Go, and many other games – 'narrow AI' has already demonstrated the sought-for capacity on the part of the machine. We show in our book that narrow AI will bring us further along many similar dimensions by providing support for coding in areas such as office work, industrial automation, missile defense, and many more.

But narrow AI has failed to achieve similar successes in many areas where humans engage in 'intellectual tasks'. These include not only the tasks on Marcus' list but also, for example tasks such as: resolving a dispute; gathering information from participants and witnesses after a major traffic accident, reading and interpreting medical imaging scans to diagnose and treat injuries and diseases, managing a company, or commanding a special forces squadron, as well as all tasks performed by humans involving some sort of creativity.[6] These examples bring to light a whole world of `intellectual tasks' where the narrow AI approach thus

[5] I focus in what follows on the work of Polanyi; a broader treatment, comprehending also the contributions of Ryle, Merleau-Ponty and Dreyfus, is found in my "Knowing How and Knowing That", in J. C. Nyíri and Barry Smith (eds.), *Practical Knowledge: Outlines of a Theory of Traditions and Skills*, London/Sydney/New York: Croom Helm, 1988, 1-16. The topic is treated also in the book referred to in footnote 1 above.

[6] We pay little attention to creativity in our book, but the account provided by Mersch in this conference (http://www.hunfi.hu/nyiri/EEE/Mersch_paper.pdf) is fully in keeping with our arguments there.

far has been (and we believe will forever be) unable to achieve any sort of success.

The OpenAI Charter, which defines AGI as 'a highly autonomous system that outperforms humans at most economically valuable work' makes matters worse by ignoring the different sorts of economically valuable *physical* work that involves the deployment of practical knowledge.

AGI is impossible: the problem of complex systems

The characteristic feature of all the mentioned tasks is that the machine called upon to address them would have to predict how *complex systems* will behave. But the central thesis of our book is that a prediction of this sort is impossible.[7] This is because complex systems have themodynamical properties – which we document at length – which preclude the collection of data pertaining to system behaviour in such a way that the data collected will manifest a statistical distribution[8] that is *representative* of that system's behaviour in the future.

> Examples of non-complex systems[9] are: the solar system, your phone, your car, the Chernobyl power station in its regular state.
>
> Examples of complex systems are: every organism, every family, the New York Stock Exchange, the earth's climate system, the earth's water system, the Chernobyl power station in the period beginning with the freak power surge caused by the system's operators on 26 April 1986.

The behaviour of complex systems is such that data that is representative of such behaviour cannot be collected because this behaviour *has no regular distribution*. Thus it is impossible to create any sort of stochastic AI model of such behaviour, since models of this sort are at core mathematical algorithms for predicting specific sorts of outputs

---

[7] More precisely, we demonstrate that it is impossible to predict the behaviour of a complex system in a way that would allow us to engineer a system – for example an AI system – that would simulate or emulate its behaviour.

[8] Illustrated for example by the familiar Bell curve.

[9] These are referred to in our book under the heading of 'logic systems'.

from specific sorts of inputs,[10] predictions which can be made only in those cases where the model has been configured with data that is representative of the relation between the system's inputs and outputs.

Every conversation between human beings is an example of a complex system, because representative data are impossible to obtain for the relations which hold between successive utterances in a conversation. This is why, when we find ourselves talking to a chatbot on the phone, we immediately start searching for ways to be put through to a human being. Somfai, in her contribution to this conference[11], shows that even the study of medieval manuscripts involves engaging with a variety of complex systems.

Science

One prime example of a complex system (or rather: of a system of complex systems of complex systems) is the human enterprise we call 'science'. It was Polanyi who showed that science is a domain that does not consist of rule-governed activities that can be described in propositional form. Indeed, far from being a purely rational enterprise of cognition and calculation, science involves of necessity a non-formalisable, non-mechanisable, characteristically human phenomenon which he refers to under the headings of 'personal' and 'tacit' knowledge. He uses the former to bring out the scientist's commitment to an as yet unknown discovery that forms the horizon of his activities,[12] a horizon that is determined by his skills or know how and to what he calls 'tacit knowledge'.[13]

Imagine, therefore, a scientist who has the first glimmering of a new discovery. The discovery will one day, if all goes well, be expressed in some propositional form. But in the beginning it might consist merely in the fact that the scientist noticed some subtle mismatch in the way two streams of data were lining up. The scientist's skills may draw on the

---

[10] 'Predicting' is used in such contexts to include also 'generating'. Thus when a chatbot generates a response to an utterance then this is analysed from the mathematical point of view as a prediction.
[11] http://www.hunfi.hu/nyiri/EEE/Somfai_paper.pdf
[12] Michael Polanyi, *Personal Knowledge: Towards a Post-Critical Philosophy*, Chicago: University of Chicago Press, 1958.
[13] Michae Polanyi, *The Tacit Dimension*, London: Routledge and Kegan Paul, 1967.

propositional content which he has learned from lectures at the beginning of his career. But then these skills will have matured. What is propositional will recede into the background and be replaced by an un- or semi-conscious application of judgment and expertise – judgment, for example, about who in the field has results that can be trusted; expertise of the sort needed to recognise an anomaly in the pattern generated by some new apparatus.

Polanyi, in fact, sees the scientific enterprise itself as resting on a deep-rooted and fundamentally non-utilitarian fascination with order or pattern, a fascination that is present already in the baby's pleasure in experimenting with coloured blocks or with the melodies of language, and which is manifested particularly clearly in the drive of the pure mathematician to discover the properties of abstract mathematical structures for their own sake.

This personal dimension of science is not capable of being rendered explicit and codified into rules, since the higher forms of human activity are always such that the rules for their performance are not and cannot be fully known to the performer. This implies the indispensability, where such activities are cultivated, of personal contact between master and pupil, of learning by doing.

Language and tools

As the carpenter should focus not on his tools but on the object worked, so, Polanyi argues, the novice scientist must be brought to a state where he need pay only subsidiary attention to the theories or interpretative frameworks which he is called upon to employ in his work. He must, in Polanyi's own words, learn to 'dwell within them', to allow theoretical tools, languages, disciplines, to serve as natural extensions of his psyche in much the way that the blind man's stick serves as an extension of his body in walking. Theories, languages and interpretative frameworks are then not abstract objects fixed in some Platonic realm, but rather social formations tied to their contingent factual realisations in the practices nurtured at any given stage by the community of scientists.

The technical terms of a science as these are conceived by Polanyi thus have meanings which are the residues of established usage and they

will change and mutate with the gradual evolution of this usage within the larger context of scientific practice and will at any given stage be only partially determinate. Each scientist's individual grasp of the science will itself change and mutate as he learns to 'see' the objects with which it deals. Thus Polanyi points to the way in which, when novice radiologists are attending lectures on how to interpret radiographs, what they see is to a large extent dependent on what they hear the expert say, whereby the meaningfulness of the latter is itself at the same time dependent on the novices' gradually developing capacity to see appropriate structures in the radiographs before them. And as Polanyi points out, it is here not so much individual words that are important, but rather the general structures to which these words relate and which they may indeed have helped to crystallise.

AI winter ahoy

Polanyi sees what might be called discursive or theoretical intelligence as resting necessarily on a seedbed of practical knowledge and perceptual judgment. This means that AI models based on a purely rational and discursive conception of human knowledge will be incapable of coming close to simulating those achievements of human beings which involve the taking account of a wealth of interdependent contextual clues in spontaneously adaptive behaviour.

Today, every university – and I mean *every* university; every university *in the world* – is putting together its plans for a new AI building. Researchers everywhere, and in whatever field, are assembling plans for using AI (or at least for bringing about the appearance of using AI) in their work. Many AI engineers and AI entrepreneurs are devising new AI models and founding new AI companies with the goal, they say, of helping to advance the progress of scientific research, thereby ignorant of the fact that all of the features identified by Polanyi as playing an essential role in the scientific endeavour will defeat the AI models which result from their efforts.