



Munich Personal RePEc Archive

**Why Machines Will Not Replace
Entrepreneurs. On the Inevitable
Limitations of Artificial Intelligence in
Economic Life.**

Van Den Hauwe, Ludwig

10 August 2023

Online at <https://mpra.ub.uni-muenchen.de/118307/>
MPRA Paper No. 118307, posted 16 Aug 2023 00:15 UTC

Why Machines Will Not Replace Entrepreneurs. On the Inevitable Limitations of Artificial Intelligence in Economic Life.

Paper prepared for the 7th Annual Madrid Conference on Austrian Economics on 26th and 27th of October 2023

By Ludwig Van Den Hauwe, PhD

Why Machines Will Not Replace Entrepreneurs. On the Inevitable Limitations of Artificial Intelligence in Economic Life.

By: Ludwig Van Den Hauwe, PhD

Abstract

This paper critically explores some supposed implications of the development of artificial intelligence (AI), particularly also machine learning (ML), for how we conceive of the role of entrepreneurship in the economy. The question of the impact of AI and ML is examined by hypothesizing a decentralized market-based system and raising the question of whether entrepreneurs will someday likely be replaced by machines. The answer turns out to be highly skeptical. Not only does the materialist worldview behind the ambitions of much AI research cast serious doubts upon the chances of success of any attempts to emulate entrepreneurship algorithmically with the help of computers, the very possibility of artificial general intelligence (AGI) can also be ruled out on purely scientific grounds. The paper concludes that human entrepreneurs will remain the driving force of the market.

The field of artificial (general) intelligence has made no progress because there is an unsolved philosophical problem at its heart: we do not understand how creativity works.

David Deutsch

1. Introduction

The recent hype cycle surrounding the development of artificial intelligence (AI), especially machine learning (ML), has led economists to reexamine some traditional lessons of economic policy, particularly regarding the question whether AI and ML might circumvent the problems of central planning in view of the way in which societies create and use knowledge. (Hayek 1945) It turns out that the lessons we know about what constitutes good and bad economic policies are likely to remain largely unchanged. (Fernández-Villaverde 2020) As this author reminds, “(t)he objections to central planning are not that solving the associated optimization problem is extremely complex, which it is and increasingly so in an economy with a maddening explosion of products, or that we need to gather the data and process it sufficiently fast. If that were the case, ML could perhaps solve the problem, if not now, then in a few more iterations of Moore’s law. The objections to central planning are that the information one needs to undertake it is dispersed and, in the absence of a market system, agents will never have the incentives to reveal it or even to create new information through entrepreneurial and innovative activity.” (12)

Along similar lines Boettke et al. (2023) argue that, despite the prospect of what King and Petty (2021) refer to as “technosocialism,” technological advances in computation cannot replace the competitive discovery process that takes place within the context of the market. To the extent technosocialism represents a restatement of the case for market socialism, which incorrectly framed the “solution” to economic calculation under socialism as one of computing data, rather than the discovery of context-specific knowledge that only emerges through the exchange of property rights, the arguments put forth by Austrian economists regarding the impossibility of economic calculation under socialism remain just as relevant today.

This paper explores the impact of AI and ML from a somewhat different angle: Hypothesizing a decentralized market economy, the focus is specifically on the impossibility of (strong or general) AI itself by raising and answering the simple question: Will machines ever be capable of fulfilling the entrepreneurial function thus rendering human entrepreneurs obsolete? Before I present a general impossibility argument, which relies to some extent on criticisms formulated from the perspective of scientific frameworks and disciplines other than praxeology, some tenets of Austrian entrepreneurship theory are summarized, the distinction between narrow AI and AGI is clarified, some intuitively illuminating examples of the limits of AI with respect to entrepreneurship are provided and a characterization of human-level intelligence is attempted.

In this paper we will not take any definite stance on the mind-body problem, nor do we have any intention to here solve the problem of whether materialism is or is not a defensible or adequate philosophical or scientific worldview. But as will be elaborated further it is important to understand that, according to the present state of knowledge, even from a monist materialistic viewpoint according to which mental processes are physical processes the impossibility of AGI is an undeniable fact due to severe limitations on our ability to model complex systems mathematically. There is no need to invoke any mind-body discontinuity or to reject scientific materialism to demonstrate the impossibility of artificial (general) intelligence.

2. The essential nature of entrepreneurship

Austrian economists can pride themselves with having a theory of entrepreneurship, or in any case an economic theory that *includes* entrepreneurship. An examination of the literature reveals, however, that several conceptions of entrepreneurship have been developed within the broad field of Austrian economics not all of which are equally relevant from the perspective adopted here.

Hayek (1945) notes how entrepreneurs adapt to events they neither have nor need to have knowledge about by responding to price changes. Similarly, Kirzner (1973) argues that entrepreneurship as alertness to opportunity contributes to equilibrating the economy. In more recent theory development, Foss and Klein (2012) argue, alongside Knight (1921), that entrepreneurship is about exercising judgment by establishing business firms within which they can conduct controlled experiments.

One could argue that from the perspective developed by these authors, much of modern Austrian theorizing on entrepreneurship somewhat misses the mark by (1) treating entrepreneurship as an important component in but not the *driving force* of the market process and (2) conceptualizing the entrepreneur as primarily *a responsive agent*.

Although Hayek and Kirzner conceive of the entrepreneur as acting within a market process, they both subscribe to the view of entrepreneurship as responsive to given circumstances. They take the boundaries of the market process as given and attempt maximizing, or at any rate improving, adjustments of production for profit. Both explain entrepreneurship as a force that equilibrates and improves on the overall market, but neither conception of entrepreneurship explains the driving force of the process. (Per L. Bylund 2022b)

The more recently developed judgment-based approach (Foss & Klein 2012) is complementary to Hayek's and Kirzner's arguments by focusing on what affords the entrepreneur the decision-making power and ability to make adjustments and act on opportunities. It focuses on the entrepreneur as an active owner-decision-maker, a capital owner who bears the uncertainty of production.

In the present context entrepreneurship will to the contrary be conceptualized along lines developed on the one hand by Ludwig von Mises, who was very clear and explicit about the importance of entrepreneurship as *the driving force of the market process*, and on the other hand by Jesús Huerta de Soto, who has emphasized *the essentially creative nature and spiritual aspects of entrepreneurship*.

Jesús Huerta de Soto has distinctly highlighted the essentially creative nature and spiritual dimension of entrepreneurship. According to Prof. Huerta de Soto “(t)he exercise of entrepreneurship does not require any means. That is to say, entrepreneurship does not entail any costs and is therefore fundamentally creative. This creative aspect of entrepreneurship is embodied in its production of a type of profit which, in a sense, arises out of nothing, and which we shall therefore refer to as pure entrepreneurial profit. To derive entrepreneurial profit one needs no prior means, but only to exercise entrepreneurship well.” (Huerta de Soto 2008, 21)

All human action thus has an essentially creative component, and no basis exists for distinguishing between entrepreneurial creativity in the economic realm and creativity in other human spheres (artistic, social, and so on). The essence of creativity is the same in all areas, and the concept and characteristics of entrepreneurship, both of which we are analyzing, apply to all human action, regardless of the type. (Huerta de Soto 2010, 42)

Moreover, “(t)he fact that entrepreneurship is distinctly creative and that therefore pure entrepreneurial profits arise from nothing can lead us to the following theological digression: if we accept for the sake of argument that a Supreme Being exists, one who created all things from nothing, then when we suppose entrepreneurship to be an ex nihilo creation of pure entrepreneurial profits, it seems clear that man resembles God precisely when man exercises pure entrepreneurship! This means that man, more than homo sapiens, is homo agens or homo empresario, and that more than when he thinks, he resembles God when he acts, that is, when he conceives and discovers new ends and means. We could even construct an entire theory of happiness, a theory which would suggest that man is happiest when he resembles his Creator. In other words, the cause of the greatest happiness in man would be to recognize and reach his objectives (which implies action and the exercise of entrepreneurship).” (ibid. 42)

The phenomenon of entrepreneurship according to this view exhibits a spiritual non-material dimension. The non-material dimension of human entrepreneurship is also highlighted by Sautet (2022) who argues from within an Aristotelian framework that alertness, the central concept in Kirzner's theory of the entrepreneurial function, can be understood as a potentiality or propensity with a very specific meaning: it emanates from the human intellect, which, through its immateriality, is capable of introducing novelty in the subjectively perceived world by the agent doing the acting. Austrian economics thus assumes, most of the time implicitly, an open-ended world and a human mind or intellect that, as in the hylomorphic tradition of the human soul known to Aristotelian scholars, is itself open-ended, immaterial, and capable of sheer creation.

This is an important and significant conclusion that, given the undeniably materialist worldview underlying the field of AI research, already casts some serious doubts upon the chances of success of any attempts to emulate entrepreneurship algorithmically with the help of machines. But as will be noted further, the impossibility of emulating entrepreneurial creativity with the help of machines does not strictly depend or rely upon entrepreneurship being an immaterial rather than a material phenomenon. The argument depends upon entrepreneurial creativity being a capability of the complex dynamical system which is the mind-body-environment continuum and the impossibility of adequately modelling this system mathematically.

3. Narrow AI versus AGI

Computers have transformed almost every aspect of life in modern technology-based societies. They have transformed health care, law enforcement, scientific research, commerce, in many cases in ways which have involved the use of purpose-built AI software. However, all successful uses of AI are examples of *narrow AI*. Examples include facial recognition, disease prediction, advanced manufacturing, spam filters, marketing content recommendations, approximate text translation etc.

In each case the software works by converting data sampled in a given area into vectors or matrices; the latter are then used to obtain a model to fulfill the task at hand. The benefits can be significant but there are also limits. AI can never deal with new types of data—

exhibiting patterns not present in its sample data—without some sort of retraining directly or indirectly involving inputs from human beings. AI does not have the natural intelligence even of an arthropod.

Artificial general intelligence (AGI), in contrast to narrow AI, can be defined as an AI that has a level of intelligence that is either equivalent to or greater than that of human beings or is able to cope with problems that arise in the world that surrounds human beings with a degree of adequacy at least like that of human beings. In 1980, philosopher John Searle introduced a distinction between weak AI—the idea that machines could act as if they were intelligent—and strong AI—the assertion that machines that do so are consciously thinking (not just simulating thinking). Over time the definition of strong AI shifted to refer to what is also called “human-level AI” or “general AI”—programs that can solve an arbitrarily wide variety of tasks, including novel ones, and do so as well as a human. (Searle 1980)

For general AI, the goal is to create a computable model of the behaviour of important aspects of the human mind-body continuum (or perhaps better: of the human mind-body-environment continuum), thereby enabling an emulation of intelligent human behaviour. But the mind-body continuum is a complex system (it is indeed a complex system of complex systems, at many levels). Thus, if our ability to create mathematical models of complex systems is severely limited, then so also is our ability to create the computable models that would be needed to create general AI.

The No Free Lunch (NFL) theorem, which was formulated and proven in the fields of search and optimisation, states that if the problem space in which an optimum is to be found must be modelled as a probability density function, then the computational cost of finding the optimum averaged over all problems in the space is the same for any solution method. (Wolpert et al. 1997) It follows that there cannot be any optimisation procedure that is globally superior to all others—a procedure can be superior only with regard to some specific problem class.

The theorem applies in particular to complex system emanations yielding data which correspond to unique (non-repeatable) multivariate distributions at each step. Indeed, for

data of this sort, per the NFL theorem, it is not only that we cannot find a globally superior optimisation method. We cannot obtain an adequate (requirement-fulfilling) predictive model of any sort.

The theorem helps us to understand why general problem solvers cannot be found for many real-world problems and why such problems need to be restricted to cases in which special solvers can provide a solution. These are exactly the cases where AI—more precisely: narrow AI—works. If intelligence is a problem-solving algorithm, then it can only be understood with respect to a specific problem. (also Chollet 2017)

What sometimes happens, however, is that such approximative special solutions—which work only for a subset of cases within a given field—are associated with claims of general applicability. Solutions of this sort will inevitably result in failures when they are applied to cases outside the restricted set. Recent cases of driver casualties in self-driving cars confronted with sensor input deviating from the training distribution are just one example of this phenomenon.

It is thus not contested that narrow AI can support or even outsmart humans including entrepreneurs at specific tasks. The tremendous successes of artificial intelligence along certain narrow lanes, such as text translation or image recognition, are not denied. Obviously, the exercise of entrepreneurship requires a broad spectrum of (not only cognitive) abilities and mimicking or emulating it computationally would certainly require AGI.

4. A few intuitive examples of what entrepreneurs can do but machines cannot.

Before attempting a theoretical characterization of what human intelligence is and what its emulation in the form of AGI would have to amount to, some easily comprehensible examples of things human entrepreneurs can do but computers cannot are here listed. They all illustrate the gulf that separates human intelligence from presently available machine intelligence:

Extreme generalization

Deep learning achieves local generalization via interpolation on a learned approximation of the data manifold. Interpolation can help to make sense of things that are very close to what one has seen before. But remarkably, humans deal with extreme novelty all the time, and they do just fine. They don't need to be trained in advance in countless examples of every situation they'll ever have to encounter. Humans are capable of extreme generalization, which is enabled by cognitive mechanisms other than interpolation: abstraction, symbolic models of the world, reasoning, logic, common sense, innate priors about the world—what we generally call reason. (Chollet 2021, 130)

Abductive reasoning.

Larson (2021, 275), pointing out that “no one has the slightest clue how to build an artificial general intelligence”, distinguishes three different types of inference: deduction, which is explored by classic symbolic AI; induction, which he classifies as the province of modern stochastic AI; and a third type which, following the American pragmatist philosopher Peirce, he calls abduction. Peirce's term is nowadays used in different contexts as another word for “hypothesis formation” or also just plain “guessing”. It is abduction, Larson argues, which is at the core of human intelligence, and thus engineering a counterpart of abduction—a combination of intuition and guessing—would be needed for human-level AI. His book provides a thorough and convincing account of why this is so. But attempts to engineer the types of abductive inference characteristic of human reasoning have in every case failed to reach even first base.¹

Making use of tacit knowledge

One of the characteristics of entrepreneurship highlighted by Prof. Huerta de Soto is that it involves tacit knowledge which cannot be articulated. (ibid. 22-4) One possibility in the AI debate is indeed that we have general intelligence, but that we can't actually write down what it is—program it, that is—because in important respects it's a black box to ourselves.(Larson 2021) Michael Polanyi argued that articulations necessarily leave out

“tacit” components of intelligence—aspects of thinking that can’t be precisely described by writing down symbols. Intelligence is only partly captured by the symbols we write down—the uses of language that he called “articulations.” Polanyi was anticipating many of the headaches AI systems have caused for AI designers, for reasons stemming from the incompleteness of articulations.

Causal understanding

Judea Pearl, while not excluding the possibility of creating an AGI, emphasizes that the currently fashionable stochastics-based “opaque learning machines” (Pearl 2020) lack an important feature of human-level intelligence in that they cannot answer questions related to causality and thus they cannot develop understanding about how things work.

Learning and self-improvement

Understanding the concept of learning is essential for understanding what drives the market process. (Harper 1996) Could computers learn in this sense? AI systems do not learn in the sense that animals and humans do. To use the term “learning” when speaking of the mechanics of stochastic AI is inappropriate because the optimization algorithms used to train neural networks do not learn in anything like the sense in which vertebrates learn. (Lapuschkin et al. 2019) Deep neural networks (dNNs) are merely “more sophisticated statistical techniques for fitting functions” and have nothing to do with real learning. (Darwiche 2018)

More precisely so-called deep neural networks (dNNs) are stochastic regression or classification models. Stochastic models are obtained by applying optimisation algorithms to the training tuples. The optimisation algorithms work under constraints with the goal of minimising the loss of the model, which means the deviation of the model from the reality of the observed outcomes. While the ability of highly sophisticated optimisation algorithms to autocompute dNN models across huge distributions is impressive, such stochastic models (and deterministic models as well) are always models of logic systems, because (a) they are executable on a Turing-machine, which is a logic system and, (b) Turing machines can only execute instructions that are logical in nature. Thus these models will not develop intentions—the equations are just functionals or operators relating an input vector to a certain output—in other words, they are nothing but a general form of regression models.

Furthermore, the nature of AI models as logic systems explains what Larson (2021, 155) calls “model saturation”, which is the phenomenon whereby stochastic models often reach a certain quality level but then cannot get any better despite the addition of new training data. The reason for this is the absolute limit, which is caused by the modelling of a complex system with a logical system. The logic system can never attain the performance of the complex system, which creates a quality hiatus that cannot be closed. (Landgrebe & Smith 2023, 147-9)

Exercising will and autonomy

Without will and the intentions and acts that flow therefrom, there is no possibility that a machine could become an autonomous agent. And if it is not autonomous, it cannot pursue any goals. It is the person who is the source of human will. (Scheler 1973) Persons are differentiated from animals, not only by their cognitive capabilities, but also by their ability to act based on their will. To create an artificial will, we need a complex of dispositions like the ones possessed by humans which can be realized in intentions, deliberations, and resolutions which all emanate from a complex system and none of which could be modelled mathematically. Hence there will be no AI will and no emulation of the will of any sort.

Moral judgment

It is impossible to teach machines moral judgement: “People need to understand that current AI—and the AI that we can foresee in the reasonable future—does not, and will not, have a moral sense or moral understanding of what is right and what is wrong” (Yoshua Bengio in Ford 2018, Chapter 2).

5. The nature of intelligence –

The difference between human intelligence and machine intelligence has scarcely gone unnoticed. An often cited example is chess. As Kasparov reminds us “(i)n what artificial intelligence and robotics experts call Moravec’s paradox, in chess, as in so many things, what machines are good at is where humans are weak, and vice versa. In 1988, the roboticist Hans Moravec wrote, “It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.” (8) “As Moravec’s paradox

dictates, computers are very good at chess calculation, which is the part humans have the most trouble with. Computers are poor at recognizing patterns and making analogical evaluations, a human strength.” (50) 2

It would seem that when projecting an “intelligence explosion” AGI theorists employ an erroneous definition of intelligence and profoundly misunderstand both the nature of intelligence and the behavior of recursively self-augmenting systems. Human intelligence depends on innate dispositions, on interaction with the environment (sensorimotor affordances), and on socialization; it can be exemplified only by a human being who is part of society. Complex real-world systems cannot be modelled using the Markov assumption. (Landgrebe & Smith 2023, 16, henceforth L&S 2023; Chollet 2017)

What, then, is human intelligence and what should machine intelligence look like if it is to emulate human intelligence?

On a general level and for clarity’s sake we can distinguish between two aspects of intelligence, which following L&S we can call “primal” and “objectifying” intelligence, respectively. Humans, of course, have only one type of intelligence, which is a fusion of both. The idea of what we are here calling “primal intelligence” was introduced by the philosopher Max Scheler as what he called “practical intelligence”. “Primal intelligence” is found in higher animals such as mammals and birds, and it may be present in other species also.

Primal intelligence is realised in non-human organisms always in an action through which the organism aims to fulfil a biological need such as drinking, eating, or life preservation through flight or fight. Animals (by which we mean here non-human animals) always live to fulfil immediate goals; they cannot create complex long-term plans. They live in the present situation and cannot abstract away from what holds only of their survival or, in higher species, the survival of their offspring. Animal perception is structurally restricted. Animals are blind to stimuli that are not related to the fulfilment of their immediate biological needs, which means that their worldview is highly restricted. Sensual clues that do not belong to the environment to which they have been adapted by evolution are ignored in something like the way that we humans, in normal circumstances, ignore ultraviolet light or radioactivity. (L&S 42) What sets humans apart is objectifying intelligence. If beavers or bower birds are removed from their habitat, they cannot survive unaided. Humans, by contrast,

have an intelligence that comprises, in addition to the spectrum of capabilities of primal intelligence, also the ability to conceive, and then deliberately plan and build artefacts that will enable them to survive even where there is no life at all—in polar barrens in the high arctic, for example, or in submarines, or in outer space.

Our genetic disposition for objectifying intelligence arose in tandem with the degeneration of our biological adaptation to the natural world. (Scheler 1961; Gehlen 1988) As homo sapiens lost the specialisation to natural environments which higher non-human mammals still enjoy, our species acquired—slowly, over millions of years of evolution—the general purpose adaptation which we are calling objectifying intelligence, and this capability has in modern times enabled humans to create their own environments summing up to the entire contemporary technosphere.

Where non-human vertebrates and all lower organisms relate to their environment in a pre-determined set of ways, objectifying intelligence allows homo sapiens to disengage himself from his environment in a way that allows him to see himself, other human beings, and the elements of this environment (both biological and non-biological) as objects, each with its own trajectory and its own array of properties and causal powers.

We can characterise the capability of objectifying intelligence as involving (L&S 46-7):

- the ability to objectify both the person's environment and her own self; each person can serve as target not only of her own but also of the others' conscious acts; and each person is aware that they can themselves become the target of the conscious acts of others;
- the ability to focus on and to track objects through time in a way that enables both short- and longterm planning (potentially extending across multiple generations), including the setting aside of resources for the future; investment in the creation of enduring physical artefacts (churches, factories, roads, theatres) and institutions (governments, legal and financial systems, religions);
- the ability to make sense of the world in terms of causality and teleology; to understand object persistence for different categories of object; to associate specific categories of processes, dispositions, capabilities, and functions with specific categories of objects;

- and to differentially and consciously value objects (including other persons) in light of their different contributions to the realisation of one's goals;
- the power of language, including the ability to think of and to categorise objects under universals and to exploit such linguistically mediated categorisations to enable more complex activities, including activities involving shared agency;
 - a heightened degree of independence (relative to what is the case for lower animals) from immediate organic necessities, which manifests itself in having and realising intentions of new sorts, including intentions belonging to cultural worlds;
 - self-distancing, which means the ability to stand outside natural life also in the sense that we are able to reflect upon ourselves as taking the point of view of an observer in relation to other objects in the world;
 - distance from the world: this means that humans have a wide range of choices as to which parts of reality they will direct their attention and interests, where animals are restricted to modes of interaction with the world that are optimised to the environmental niche into which they have evolved;
 - the ability to modify our directedness towards targets by cancelling the belief-moment. It is this which allows all forms of imaginative directedness towards objects, in the literary and visual arts as well as in planning for the future and in all forms of speculation and hypothetical reasoning. The ability to direct one's thinking to entirely new kinds of objects is a characteristic feature of human creativity.

In view of the foregoing how, then, could we obtain a definition of AI that is useful and applicable in real user settings?

If we are talking of AGI, then we would certainly want a machine with -- not merely primal but also and foremost -- objectifying intelligence. (L&S 60 ff.) For example, a robot with the ability to engage in conversations with humans in which it would be perceived as a useful interlocutor because it has, for example, the ability to understand an ambiguous order (such as: 'Give me the bottle', where there are multiple bottles standing on the shelf), disambiguate the order by asking clarificatory questions, and execute the order by moving over to the shelf and reaching out with its robot arm. Objectifying intelligence is required for this purpose because execution of the order presupposes an objectification of reality

analogous to that performed by humans. Thus the requirement for useful AI is: (not merely primal but) objectifying intelligence—including self-objectification—which would in any case be required for all purposes in which the artificial agent is required to move freely among and interact with humans. For the agent would need to move and behave in a way that is compatible with the ways humans move and behave in relation to each other in real environments and thus in a way that would make the agent, too, a part of what we can think of as the human world.

An extensive review and discussion of representative definitions of the term “intelligence” provided by the leading proponents of AI, and specifically of AGI, starting with what is in the AGI community the most influential and still the most widely accepted definition, which was put forward by Legg and Hutter in a paper entitled “Universal Intelligence: A Definition of Machine Intelligence” published in 2007, falls outside the scope of this paper.

It would appear, however, that without exception these definitions, when measured against the previously specified requirement, throw no light at all on human intelligence in either of its two aspects of primal and objectifying, and therefore do not yield machines that can fulfill this requirement; neither will they yield machines that will have the capacity to go significantly beyond traditional “narrow” AI. (L&S Chapter 3) 3

Still the account of human intelligence can be used to throw light in the reverse direction on what AI research itself has really achieved and will be able to continue to achieve in the future, using primal and human intelligence as a benchmark.

6. The general argument: the missing mathematics of complex systems and the impossibility of AGI

Whether it was John Searle’s Chinese Room argument (Searle, 1980) or Roger Penrose’s argument of the non-computable nature of a mathematician’s insight – an argument that was based on Gödel’s Incompleteness theorem (Penrose, 1989), we have always had skeptics that questioned the possibility of realizing strong Artificial Intelligence, or what has become known as Artificial General Intelligence (AGI). (Van Den Hauwe 2020) Many of the possible

objections to AI were foreseen by Alan Turing, the first person to define AI, before they were subsequently raised by others. (Turing 1950)

But the strongest and most convincing argument elaborated up till present that AGI is simply impossible has been put forward recently by Jobst Landgrebe and Barry Smith in their 2023 book *Why Machines Will Never Rule the World – Artificial Intelligence without Fear* (L&S 2023). The central question of this book is the possibility of the emulation of the most complex single-organism complex system on earth, namely the human mind-body continuum.

In the authors' view the human mind is an integral part of the human body or rather of what they call the human mind-body continuum. There is no separation of mind and body; there is only one whole. Their position can be called a "no layers" approach in that it embraces a materialistic monist view according to which mental processes are physical processes. Contrary to a computer, that is a machine that creates a numerical output based on some numerical input using a mathematical model (Turing 1937), the human brain and the human mind-body continuum are not machines of any kind. (L& S Chapters 7 & 8)

They convincingly defend the thesis that it is impossible to obtain synoptic and adequate mathematical models of complex systems, which means: models that would allow us to engineer AI systems that can fulfill the requirements such systems must satisfy if they are to emulate human-level intelligence. 4

The overall argument is quite simple and consists of two steps:

(a) Anything we engineer (a computer or any other machine) must ultimately be a system that can be modelled mathematically. That is, any engine we engineer is in the end a logical system that can be formally modelled and described by the mathematics available to us. Artificial intelligence, no matter what problems it is applied to, would have to reach its solutions by executing a set of mathematical functions that are each computable in the Church-Turing sense. Any AI algorithm must be Church-Turing computable and only algorithms that can be formulated as a sequence of elementary recursive functions are computable. (Enderton 2010) This requirement places a restriction on the sorts of programs that can be executed by a computer: they must be based on some mathematical model

whose outputs are Turing-computable from their inputs.

Any mathematical model that runs on a Turing machine can only model comprehensively and adequately what we call logic systems. This is because to be computable it must be isomorphic to an algorithm which can be expressed using the basic recursive Church functions. Each model consisting of a combination of these functions is always a model of a logic system, even if the latter is used to approximate a complex system. Only logic systems, that is, systems that can be successfully modelled using propositions of mathematics linked together by logical relations, allow models that can predict their behavior almost exactly. Computable models are models of logic systems; they all belong to the (extended) Newtonian paradigm of mathematical modelling of reality.

Stochastic models of complex systems are obtained using derivatives of loss functions, which are used to find local minima of multivariate functionals. The result is a very long, differentiable equation. Due to the mathematical properties of every dNN, this equation obeys relaxed Newtonian requirements. This means that it does not require the interactions between its variables to be always the same, and it also does not require that these interactions have to be homogeneous over the entire neural net. However, the importance of any given interaction must decrease over space or time in a regular fashion; in other words, every neural network must still have a weak Markov property over space or time. And neural networks still require most of the properties of Newtonian models in order to be computable.

Simplifying logic systems satisfy the following four conditions (L&S 122-3):

- (1) The system behaviour can be explained by reference only to one of the four fundamental interactions of gravity, electromagnetic force, and the weak and strong nuclear force.
- (2) The system behaviour of interest is dominated by a single homogeneous and isotropic force in such a way that the effects of the other interactions are so small, in the context of the modelled aspect, that they can be neglected. If there is more than one relevant force in a system, for example gravity and electromagnetic force, their effects can be modelled

separately, given that each force dominates relative to its effects on corresponding separate aspects of the system's behaviour. The interaction with other forces can be neglected.

(Thurner et al. 2018)

(3) In each system there are groups consisting of elements of the same type. The elements of each such group interact with each other in an identical manner, and they also interact with the elements of other such groups again in an identical manner (which may be different for different groups). All interaction patterns are in this sense homogeneous. For example, in the solar system, the sun and the planets can be seen as a group of elements (of type: lump of matter) which interact via gravitation. But the sun is a star and the earth, Mars, as well as the other satellites of the sun are planets, and the sun (seen as a star) also interacts with these satellites through its electromagnetic radiation.

(4) The boundary conditions of the system can be assumed to be fixed without invalidating the model, so that the system can be considered context-free, and thus the context in which the system is embedded can be abstracted away without detriment to the predictive power of the model.

(b) The mind, however, is not a logical system but a dynamic complex system that no known mathematics can model or describe. The nature of complex systems prevents their synoptic and adequate modelling.

Excursus: history of the concept of complex systems

One of the first to argue that for all animate systems we are unable to create predictive models was Henri Bergson in 1907. In part under Bergson's influence, the mathematics of complex systems was pioneered by Ilya Prigogine in his work on what he called "dissipative structures", specifically in his Introduction to Thermodynamics of Irreversible Processes (Prigogine 1955). Prigogine identified many mathematical properties of complex systems, for example relating to the ways in which such systems exhibit processes which involve a constant passage away from equilibrium.

Let's try to explain. The complexity of modeling mental processes is not simply a function of

their complex temporal or stochastic behavior; rather, it is because these processes are dynamic, adaptive, continuously evolving, and constitute systems whose behavior affects and is affected by the environment they function in. This is the source of limitations of modern-day machine learning techniques: While one can “train” a deep network on a set of input-output pairs, beyond any narrow domain no set of training data can adequately predict the future environment since the state of that environment itself is a function of the very system that we are training. Such cyclical cause-and-effect behavior of complex systems cannot be modelled by any known mathematics.

More precisely complex systems are marked by the following seven properties (L&S Chapter 7; also Thurner et al. 2018):

Property 1: Change and evolutionary character—sudden continuous and potentially non-differentiable or non-continuous changes of element types and element (type) combinations, which include changing behaviors on the part of all instances of a type. Contrary to the types of relations among the elements of logic systems that do not change over time, so that the types of behaviours manifested by these elements are given and fixed, a complex system has a creative character, which means that it can at any time create new elements and new patterns of interaction.

Each mathematical model requires a vector space -- often a coordinate space over an algebraic space F -- but with the changing variables and interactions that we find in complex systems, there is no coordinate space over which models can be defined. Since each and every model is defined for a specific vector space, it becomes invalid if the reality targeted by the model differs from the vector space for which the model was originally defined. The more it differs, the stronger the deviation and the less accurate the model becomes. This is one of the main reasons why we cannot model complex systems mathematically.

All this is related to the evolutionary character of complex systems. Evolutionary systems are adaptive and robust at the same time, a phenomenon that is very hard to model because robustness requires lack of divergence from a fixed set of states while adaptation requires the exploration of new phase spaces. Evolutionary systems are also such as to manifest path-dependence in their development and thus show a strong and long-lasting memory (in the sense that the relation of their present to their past cannot be captured using Markov

models). Such systems are therefore both non-ergodic (they cannot be modelled by averaging over space and time without losing information) and non-Markovian (their behaviour depends not just on one or two immediately preceding steps). The lack of ergodicity is one of the chief obstacles to using stochastic AI for complex systems and another main reason why we cannot model complex systems mathematically.

In probability theory, multivariate distributions can be thought of as resulting from stochastic processes, such as the Gaussian process, which is ergodic and creates a continuum of multivariate normal distributions. Each ergodic process creates a series of data which can be modelled as samples from a stable multivariate distribution which can be represented explicitly in mathematical form.

Suppose that we have a complex system and we wish to use observations of its behaviour to obtain a representative sample of the sort that we can use to train an AI application. For this to be possible, the sample data would have to correspond to a multivariate distribution that is representative of the system's behaviour, which can often be assumed for logic systems as well as for certain artificial systems such as Go and chess, where the observable behaviour is constrained by strict rules. However, there are many, many cases for which no such distribution exists. This may be, for example, because the evolutionary nature of the system will imply that the coordinates of the vector system which models its phase space are continually changing. Second, it may be because, even in the absence of such change, the observations modelled by the distribution emanate from a non-ergodic system, so that the distribution of data points in the vector space cannot be modelled adequately with either a parametric or a non-parametric distribution. This is because it is impossible to draw adequate samples from a distribution of this sort, because there is no representative subspace from which the needed training samples could be drawn. Under these conditions, there is no process that can yield a representative sample.

Ergodic distributions are rare, and the distributions we encounter in real-world data are in most cases non-parametric. This means that we cannot use parameters to build an equation to represent them mathematically, as contrasted with what is the case for distributions resulting, for example, from a Gaussian process. In cases where the data do not come from a distribution of this sort, but rather from a non-ergodic process or from a distribution generated by a complex system the stochastic model obtained by using such data will fail

when faced with new observations. This is because the latter emanate from a distribution that will diverge from the training distribution in a proportion of cases in a way that will at best ensure a poor performance and at worst make the model useless. Due to the nature of complex systems, this divergence may be unnoticeable immediately after training, but it will typically increase over time.

Property 2: Element-dependent interactions—which lead to irregularity and non-repeatability. Irregularity means that the system does not behave in a way that can be formalized using equations. Non-repeatability signifies a behavior that cannot be reproduced experimentally. When bodies are related to each other in the sorts of logic systems described in classical physics, for example through the force of gravitation, their interaction is homogeneous and not specifically related to the bodies involved—it depends only on the mass of the bodies and on the distance between them. In contrast to this, the elements of complex systems have relations specific to their nature, the interaction types are dependent on the types of the elements they relate.

Importantly, in a logic system, whether natural or artificial, an element can change its state but not its type. For example, the gravitational force a planet exerts on other bodies depends solely on its mass, no matter which state of matter it is in. However, in the sorts of complex systems we find in biology elements can dynamically change their function, and when such changes occur this interacts with their state. What this means is that when the function of an element, for example a membrane protein of a myocyte, changes due to phosphorylation, then this brings about changes in the set of its measurable non-invariant property values. It can acquire new states due to the functional change. The former are dynamically dependent on the latter. There is no way to model this sort of change mathematically for many elements and states at the same time, which is why models of complex systems can model, at best, only certain narrow aspects of a system's behaviour.

Property 3: Force overlay—several forces acting at the same time and thereby potentially interacting. This property is often correlated with anisotropy (which means that the effect resulting from force overlay does not propagate with the same magnitude in all directions). All system behaviour, including the behaviour of complex systems, is the result of the four basic physical interactions (electromagnetic, gravitational, strong, and weak). But these forces interact with each other and are overlaid upon each other in such a complicated way

in complex systems that it is impossible to model how the observed behaviour of such systems is generated.

Property 4: Non-ergodic phase spaces—which cannot be predicted from the system elements and lead to time-irreversibility. A time-irreversible process is a process which cannot be described by equations which are invariant or symmetrical under a change in the sign of time.

Complex systems have a rich phase space, which is to say that the set of all elements and their states that would be needed to describe the entire workings of the system is very large. Some directly observable macrostates such as temperature, pressure, or density are explainable exhaustively from microstates at lower granular levels (for example, from states of molecules in Brownian motion). The former, in other words, can be predicted from the latter. In complex systems, however, we observe macrostates that emerge in a fashion that cannot be predicted or derived from knowledge about the microstates which compose them. For example, we cannot adequately model regional or global average temperatures (a macrostate) from the microstates of the earth's climate system in the case where adequacy would mean that the model could predict the temperature time series with good accuracy over decades.

Yet more obstacles to modelling are created where we are dealing with non-ergodic processes, which produce events in which we cannot identify any law-like pattern that can be modeled mathematically. The reason for this is that non-ergodic processes do not yield distributions from which representative samples can be drawn.

An additional obstacle turns on the fact that the traces of non-ergodic processes—in other words the data series which such processes generate—provide no adequate target spaces for stochastic sampling. The samples drawn from such complex traces are never representative of the process behaviour due to the non-ergodic character of the process. There is here no distribution to sample from. This systematically prevents stochastic modelling of such processes.

Property 5: Drivenness—either involving some external energy force or resulting from some sort of inner drive; drivenness implies the lack of an equilibrium state to which the system

would constantly be converging. This lack of equilibrium is caused by an energy gradient and results in energy dissipation. Complex systems are often driven in the technical sense that is defined in physics (more precisely in statistical mechanics). Driven systems undergo a flow of energy, which prevents them from converging or moving to an equilibrium; the energy flow pushes them ever onward from one state to the next. The mathematical difficulties in dealing with out-of-equilibrium or non-equilibrium systems are tremendous and beyond analytical reach.

Property 6: Context-dependence—non-fixable boundary conditions and embeddedness in one or more wider environments. In complex systems, the boundary conditions at the interface between system and environment are constantly changing. This is why a complex system cannot be modelled by assuming that its boundary conditions (formed by the elements at the boundary) are fixed: doing this would create an invalid model. In other words, one cannot abstract from this environment without fundamentally mismodelling the behaviour of the systems it contains. When dealing with logic systems, in contrast, one can abstract from the context; the boundary conditions of the system can be assumed to be fixed, and the system itself is in this sense context-free. Because complex systems are context-dependent; their boundary conditions massively determine how they work.

The context-dependence property of complex systems has the consequence that the system will use a different phase space following different principles depending on the context in which it is situated. Yet neural networks always rely on the assumption that all the input-output-relationships they model via their training samples are context free. The distribution from which they are drawn has no further context. Crucially, this means that they cannot cope with the non-ergodic system events which are characteristic of complex systems as the networks are trained using large sets of events over which they merely average. No matter how large the model parameterisation becomes, this training process cannot yield models of complex systems which are both synoptical and adequate. In other words, when data are sampled from a complex system, they are never representative of the system, for the system's behaviour never has a multivariate distribution from which one could draw representative samples. Context-dependence is another main reason why we cannot model complex systems mathematically.

Property 7: Chaos—inability to predict system behavior due to inability to obtain exact measurements of starting conditions. Chaotic behaviour results from the dependence of a system on its starting conditions and is referred to as deterministic chaos in physics. It arises not only in complex systems, but also in simple systems, for which it was first described. In such systems, we know exactly which laws govern a physical process and can model it with a number of variables that is sufficiently small to allow us, in principle, to obtain a predictive model. However we fail to do so because we are unable to measure the starting conditions with sufficient exactness. No matter which type of system we are dealing with, chaos cannot be predictively modelled—the divergence from the real outcome may sometimes be low over very short observation intervals, but it increases exponentially over time. While there are non-chaotic simple (Newtonian) systems, complex systems are in every case chaotic.

Clearly, very many of the systems we encounter in nature, including the global climate and plate tectonic systems, and almost all the systems we encounter in the realm of living organisms, are complex. This means that they cannot be modelled in a way that would yield the sorts of mathematical predictions that can be reliably used in technological applications.

Most processes in nature, even many seemingly simple inanimate processes, cannot be modelled mathematically. We cannot write down or automatically generate equations which describe, explain, or predict such processes accurately.

The class of problems in relation to which mathematical modelling has been singularly successful in generating exact or almost exact predictions belongs to the domain of physics where we can usefully employ “extended Newtonian mathematics”, comprising the entirety of those mathematical resources that have the sort of predictive power first unleashed by the invention by Newton and Leibniz of the differential calculus. But the structure of extended Newtonian mathematics and the limitations of its models that have been brought to light through the development of chaos theory and the theory of complex systems have far-reaching implications as concerns the possibility of our creating models with the ability to predict the behaviors of complex chaotic systems such as the human brain. The latter would require a major revolution in mathematics of a type which has been ruled out as impossible by leaders in the field, and no traces of which are even on the horizon. If we are restricted to using extended Newtonian mathematics, and so long as we are constrained to use those algorithms of extended Newtonian mathematics which can be executed on universal Turing

machines, it is not conceivable that we will be able to mathematically model, and thereby to engineer, a system with the complexity required to emulate human intelligence. In other words, there is no way to model the behavior of a complex system with the accuracy necessary to support sound technical applications and attempts to apply extended Newtonian mathematics to complex systems lead to failures in most settings, and this applies not least to the human central nervous system.

Summarizing, both the argument that the mind or some faculties of the mind are complex systems that are dynamic, adaptive, continuously evolving, and are systems whose behavior affects and is affected by the environment they function in, and the argument that the behavior of such systems is beyond any known mathematics are very compelling and certainly also refute any claim that an AGI is conceivable that could mathematically or algorithmically emulate (or go beyond) human entrepreneurial creativity.

Schematically the argument can be summarized as follows: (1) In order to emulate entrepreneurial creativity with the help of AI we would have to simulate these creative processes computationally; (2) Entrepreneurial creativity is a capability of the complex dynamical system which is the human mind-body-environment continuum; (3) Therefore an emulation of entrepreneurial creativity with the help of machines would require to simulate computationally the workings of complex dynamical systems; (4) Simulating a complex dynamical system computationally requires adequate mathematical models of such systems. (5) Adequate mathematical models of complex dynamical systems are impossible. (6) Therefore, it is impossible to emulate entrepreneurial creativity with the help of machines.

This critique was clearly anticipated by Jesús Huerta de Soto when he wrote:

“ (...) mathematicians have yet to (and may never) take up the challenge of conceiving and developing a whole new “mathematics” which permits the analysis of human creative capacity with all of its implications.” (Huerta de Soto 2008, 108)

Some of today’s AI proponents believe that the currently fashionable AI paradigm of “deep neural networks”—connectionist as opposed to symbolic AI—can mimic the way the brain functions; L&S show that, again for mathematical reasons, this is not so, not only for deep neural networks but for any other type of AI software that might be invented in the future. 5

The argument against the possibility of AGI is in more than one respect analogous to and can elucidate the argument of Mises and Hayek against the possibility socialism as L&S also recognize. (L&S 157-8) Both the human brain and the economic system are complex systems that are not amenable to effective and satisfactory mathematical modelling. 6

As L & S recognize economics yields mostly descriptive and interpretative models, involving no mathematical causality and yielding no exact predictions. Macroeconomics for instance provides no causal explanations, but rather (at best) very helpful causal interpretations. No economic model can predict exactly any single economic quantity for any selected time or time interval in the future, whether this be the price of a good or the excess capacity of a production method. Nor can the causation of economic phenomena be modelled causally in such a way as to yield a scientific explanation—again, because of the complexity of the system.

Let's summarize. There are hard boundaries to the modelling of complex systems, so that causal explanations and exact predictions—even of single traits of these systems—are in almost all cases mathematically impossible. This is so because for such systems we are unable to formulate equations that yield the needed predictions. For an AGI designed to substitute for humans in the performance of complex tasks in natural environments, inexact predictions are insufficient: the AGI will not pass even minimal safety checks. The problem here is that, if we measure the behaviour of complex systems by assigning numbers to the observable events which these systems (co-)generate, we obtain data to which no predictive model can be made to fit, no matter which procedure we use. An example is the system formed by two human beings when they engage in a dialogue.

However, many partial aspects and properties of complex systems can be modelled descriptively or approximatively. Economics -- in its "mainstream" variant -- is only one of a number of disciplines in the life sciences (biology, biochemistry, medicine, pharmacology, and so forth) and also in the humanities and certain other social science disciplines (psychology, anthropology, ethnology...) all dealing with complex systems that widely use mathematical models for descriptive, interpretative, and approximatively and partially

predictive modelling. But the nature of complex systems sets tight boundaries on what such descriptive modelling can achieve. It is important to understand that synoptic and adequate models of complex systems are not possible.

Mathematicians who have become aware of the inadequacy of Newtonian mathematics for the modelling of complex systems have tried to develop more sophisticated (non-naïve) approaches, using mathematical frameworks which can cope with the properties of complex systems and yet remain computable. The study of these approaches falls outside the scope of this paper. Non-naïve approaches to complex system modelling are often mathematically interesting and contribute to our descriptive and interpretative understanding of aspects of the phenomena under study. However, they do not give a procedure to obtain exact causal or predictive mathematical models of complex systems, in most cases not even for single traits of such systems. Such a procedure can be found only for simple (logic) systems that are man-made and artificially driven. Predictive mathematical models for the behaviour of any complex system have thus far not been provided on any approach.

Excursus: the uniqueness of the methodology of the Austrian School of economics

As I have pointed out elsewhere (Van Den Hauwe 2009, 213-4) and want to repeat here, the economists of the Austrian School of economics, in particular Ludwig von Mises and his followers, have developed a unique theoretical method, the method of praxeology, that can be interpreted as a method and device to cope with the complexity of economic phenomena. This method is both exact and non-mathematical, both predictive and non-quantitative. An elaboration of this theme falls outside the scope of this paper, however, which is devoted to the relationship between entrepreneurship and artificial intelligence.

7. Implications

The “general impossibility” is exemplified by some more specific impossibilities that equally render AGI impossible. Prominent among these are:

(1) Machines will not master human language. (L & S Chapters 4, 5 & 10)

Language is a prerequisite to any AGI but since linguistic communication – comprising open interactive dialogues -- is itself a complex system that no mathematics can model, again no AGI is possible. In a real dialogue the interpretation of some utterance must be a function of previous utterances and the overall context that has been built so far. But since responses cannot be predicted in any meaningful way, the overall context is not well defined, and so the entire interaction cannot be mathematically modelled.

The most striking capability which distinguishes human beings from other animals is our ability to speak, and more specifically to conduct conversations. Language is the most important observable expression of our objectifying intelligence. Animals have no language, and they have no non-verbal abstract symbols such as badges or insignia, no ability to manipulate numbers, and no objectifying intelligence.

L & S lay out the role that language plays for humans and describes language complexity to let us appreciate the challenge that lies in the attempt to mathematically model language in a way that would be required to create an AI. (L & S Chapters 4 & 5) Humans produce meaningful language and assign meaning to the language produced by others in a dynamic process. L & S summarize the current view of language production and interpretation on the part of philosophers of language and of linguists. (L & S Chapter 5) The result is then used as basis for understanding their argument in later chapters to the effect that it is impossible to model mathematically either of these capabilities of the human mind in a way that is adequate in the sense that it is able to generate the sorts of predictions needed to support machine emulation of human language use.⁷

As L & S conclude:

“When a conversation occurs between human beings, multiple complex systems, each with its own evolving sets of intentions and realizing its own sets of capabilities, are interacting with each other. Interactions of this sort are analogous to those which occur when other sorts of complex systems interact—for instance when the earth’s tidal system interacts with the ecological systems of coastal wetlands. We can describe and explain some of what occurs in the course of such interactions; but we cannot build mathematical models that will enable us to predict what will occur. The two sorts of systems simply interact. That is what they do.

And so, too, in the case of many sorts of interactions, both linguistic and non-linguistic, involving humans: humans do not consciously or unconsciously compute these interactions (because the human mind-body continuum is not any sort of computer). Rather, they simply interact in a way that involves, at the level of ultimate physics, a constantly self-adjusting sequence of interactions between the different sets of fundamental forces deriving from the different human beings involved.” (89)

(2) Machines will not master social interaction. (L & S Chapters 6 & 11)

We will never be able to engineer machines with the social and ethical capabilities of human beings. In preparation for drawing this conclusion we need to understand what these capabilities are. To this end L & S engage in an accelerated grand tour through sociology and social ontology, focusing on three sets of issues, relating to (a) social behaviour in communities, societies, and institutions, (b) perspective-taking and intersubjectivity, and (3) social norms, including legal and moral norms. In chapter 11 L & S then address the implications of this for the possibility of emulating ethics in the machine.

As Adam Smith was perhaps the first to recognize, in all social interactions—from shaking hands in order to seal a deal, to assisting in someone’s suicide, to the public dialogue between magistrate and thief that precedes the thief’s being condemned to the stocks—a successful outcome requires that all parties have been able to use their social capabilities to understand the situation they are in and the norms thereby entailed. It requires also that they each use these same capabilities to understand the intentions of the other parties, and the power gradients that obtain between them (Smith 1790, I.i.1.3). Value consciousness and the ability to integrate social norms, intersubjectivity, and power relationships consciously into a coherent, deliberate form of behaviour is a capability exclusive to humans. Animals can recognize very simple value differentials (for example between pleasure and pain) and perform elementary integrations of social norms and social rank; but they do not have the capability to apprehend values of higher order or to perform the conscious integration of values, feelings, and intentions that humans are capable of. (L & S 106)

Since we can emulate neither human intelligence nor human language in the machine because we lack the mathematical models that would be needed to do so, it follows that we

cannot emulate human social capabilities either, since these require both intelligence and mastery of language. There can be no machine intersubjectivity, no machine social norms, no law-abiding behavior or emulation of morality by machines. (L & S Chapter 11)

8. Conclusion: machines will not replace entrepreneurs

Human and machine intelligence are radically different. The myth of AI insists that the differences are only temporary, in the sense that, step-by-step, more powerful AI systems will erase them. Yet the success achieved by focusing on narrow AI applications gets us not one step closer to general intelligence. No algorithm exists for general intelligence. And we have good reason to be skeptical that such an algorithm will emerge through further efforts on deep learning systems or any other approach popular today.

At the intuitive level the contrast between the materialistic worldview underlying most of AI research on the one hand and the immaterial aspects of entrepreneurship on the other, already casts serious doubts upon any claim to the effect that entrepreneurial creativity could be emulated algorithmically by a computer.

Summarizing some tenets of Austrian entrepreneurship theory, in particular highlighting the immaterial and spiritual nature of the phenomenon and confronting these with the assumptions underlying AGI research has allowed us to perceive the incongruence of any attempt to explain entrepreneurship in materialistic (deterministic, reductionistic...) terms. However, even without assuming any mind-body discontinuity, that is, even if mental processes are themselves physical processes, the impossibility of AGI can be demonstrated relying on scientific contributions from a range of disciplines, and any claims regarding the prospects of emulating entrepreneurship algorithmically and someday replacing entrepreneurs by machines or robots are clearly unfounded. The core of the argument relates to the fact the emulation of entrepreneurial creativity with the help of machines would require the synoptic and adequate mathematical modelling of the complex dynamical system which is the human-mind-environment continuum which is impossible. If AGI is

defined as a form of machine intelligence that allows the construction of a synoptic and adequate model of human-level intelligence and creativity, it is for the same reason impossible.

Whatever the useful implications of the development of AI for the economy are and will be – see e.g. HBR 2019 -- and despite enormous advances in (narrow) AI, machines will not replace entrepreneurs and genuine human entrepreneurs will remain the driving force of the market economy. This conclusion warrants optimism regarding the prospects of future research into the nature of entrepreneurship along lines initiated by Austrian economists.

Notes

1 It is not quite correct that machines engage in inductive reasoning; they rather compute local minima for loss functions, which can be seen as a very primitive emulation of induction from data because a functional is indeed obtained from observations (individual data). However, machines do not perform the induction themselves; they merely compute human-designed optimization algorithms which emulate a narrow form of human induction.

2 IBM's famous Deep Blue prevailed in chess over Gary Kasparov, and more recently, AI systems have prevailed in other games, e.g. Jeopardy! and Go, which is an illustration of the fact that in certain focused areas machines can out-perform human minds. There are two fundamental types of computable system models: deterministic and stochastic. The former comprise, for example, models expressed using propositional, predicate or modal logic, and including what are called expert systems or rule systems. The chess-playing algorithm Deep Blue that beat Kasparov in 1996 was deterministic; it used an α - β -search algorithm (Heineman et al. 2008, chapter 7).

3 The definitions of intelligence based on utility functions proposed by the AGI community identify the intelligence of a machine on the basis of the fact that the machine is endowed with an optimisation framework for obtaining some extremum for a high-dimensional functional for which derivatives can be calculated. This formulation is just an alternative way of stating that, as on all connectivist approaches to AI, they obtain a model which is defined via a loss function, or in other words that they execute a recipe found using optimisation.

This brings one advantage over AI based on symbolic logic (GOF-AI), namely that the connectionist AI algorithms can be generated automatically, where GOF-AI requires algorithms that are designed explicitly. In this way, the new utility-based AI yields an approach that can scale to apply in areas where we have to deal with very large bodies of data with a certain degree of variance. But it is an approach which works only where we can assemble training samples with a variance which is representative of the variance in the target data. This is possible only along certain very narrow lanes. Alternative definitions of intelligence are unlikely to yield anything that can fulfill the requirements described earlier. For no matter how we generate an alternative AI, it will have to emulate what we call a 'logic system', which is a system such as a simple device engineered in such a way that its behaviour can be predicted using the equations of physics and the rules of logic.

4 To enable a classification of such models according to their utility, L & S introduce the notions of synoptic and adequate models. A synoptic model is a model that can be used either 1. to engineer a system or system component of a specified sort (for example, a combustion engine or an artificial heart), or 2. to emulate the behaviour of a system or system component (for example, the behaviour of a tiger as emulated in a computer game, or the behaviour of a clerk in a travel agency using a chatbot). A model is adequate relative to some set of specified requirements if it can be used to engineer an artefact, or to create an emulation, that satisfies all the requirements of that set. (112)

5 Even a nervous system made of only a few hundred neurons is much more complex than an artificial dNN with billions of parameters, which is merely a (big) logic-system-modelling equation. This is because each neuron contains millions of signal-integrating molecules and is connected to other neurons via synapses using a plenitude of neurotransmitters which elicit many different reactions based on the state of the post-synaptic neuron. Furthermore, the neurons of higher organisms also depend on humoral factors (hormones and other signalling molecules in the blood). They are living cells, which are driven and thus never in equilibrium, but they produce and consume energy all the time. In short, unlike stochastic models (such as dNNs), which are logic systems and can thus be executed on computers (to approximate complex systems), nervous systems are complex systems in their own right. (L&S 168-9)

6 In this respect the theory of complex systems comprises some lessons not only for AI enthusiasts, but also for economic methodologists. There are three types of models: descriptive, explanatory, and predictive. (L & S 111-2) There are two types of explanations: 1.

Interpretative explanation of effects of certain types, in which important causes of the effect types can be listed and the relationship between cause and effect types can be qualitatively described. 2. Full causal explanation, in which the physically relevant types of causes and their effects can be enumerated, and their relationships can be modelled quantitatively and exactly using an equation or a set of equations. Prediction refers to those cases where we can model the behaviour of a system in such a way that we have an assurance that, given an input of the sort for which the model is designed, the model will yield an output (a prediction) that is in accordance with the behaviour of the modelled system. Predictive models can be exact or approximative. In the latter case they are stochastic, where a simple example is a model of the outcome of throwing a dice. All stochastic AI models, such as classical statistical learning models or deep neural networks, are of this approximatively predictive type. It is exact models that enable strict scientific knowledge, including both exact causal explanatory and exact predictive models. This is the sort of knowledge that we can obtain in physics, in chemistry, and in certain areas of biology.

7 For mathematical models *predict* is not restricted in its meaning to the prediction of future events (as in weather forecasting). Rather, it is used more generally to denote the calculation or computation of model output from some model input. In artificial dialogue systems, the computation of a machine utterance based on the utterance of a human being is also a prediction from the perspective of mathematics; from a user perspective, however, it is rather simply a succeeding utterance.

References

Bergson, Henri. 1911. *Creative evolution* [1907]. New York: Henry Holt & Co.

Boettke Peter J. and Rosolino A. Candela, 2023. On the feasibility of technosocialism, *Journal of Economic Behavior & Organization*, Volume 205, Pages 44-54, <https://doi.org/10.1016/j.jebo.2022.10.046>.

Bylund, Per L. 2022a. *A Modern Guide to Austrian Economics*, Cheltenham: Edward Elgar.

Bylund, Per L. 2022b. Entrepreneurship and the market process, Chapter 5 in Bylund 2022a, 84-102.

Chollet, François. 2017. The implausibility of intelligence explosion, Retrieved from: <https://medium.com/@francois.chollet/the-impossibility-of-intelligence-explosion-5be4a9eda6ec>

- Chollet, François. 2021. Deep learning with Python. 2nd Edition, Shelter Island, NY: Manning Publications Company.
- Darwiche, Adnan. 2018. Human-level intelligence or animal-like abilities? Communications of the ACM 61 (10): 56–67.
- Enderston, Herbert B. 2010. Computability theory: an introduction to recursion theory. Cambridge, MA: Academic Press.
- Fernández-Villaverde, Jesús. 2020. Simple rules for a complex world with artificial intelligence. <https://economics.sas.upenn.edu/pier/working-paper/2020/simple-rules-complex-world-artificial-intelligence>.
- Ford, Martin. 2018. Architects of intelligence: the truth about AI from the people building it. Birmingham: Packt Publishing Ltd.
- Foss N. J. & Klein P. G. 2012. Organizing Entrepreneurial Judgment : A New Approach to the Firm, Cambridge: Cambridge University Press
- Gehlen, Arnold. 1988. Man: his nature and place in the world [1940]. New York: Columbia University Press.
- Harper, D. 1996. Entrepreneurship and the Market Process -- An enquiry into the growth of knowledge, London: Routledge.
- Harvard Business Review (HBR) 2019. Artificial Intelligence, Boston: Harvard Business Review Press.
- Hayek, Friedrich August von. 1945. The use of knowledge in society. The American Economic Review 35 (4): 519–530.
- Heineman, George T., Gary Pollice, and Stanley Selkow. 2008. Algorithms in a nutshell. Sebastopol, CA: O'Reilly.
- Huerta de Soto, J. 2008. The Austrian School Market Order and Entrepreneurial Creativity, Cheltenham: Edward Elgar.
- Huerta de Soto, J. 2010. Socialism, Economic Calculation and Entrepreneurship, Cheltenham: Edward Elgar.
- Kasparov, G. 2017. Deep Thinking Where Machine Intelligence Ends and Human Creativity Begins, Great Britain: John Murray
- King B. and R. Petty 2021 The Rise of Technosocialism: How Inequality, AI and Climate will Usher in a New World, Marshall Cavendish International.
- Kirzner, I. M. 1973. Competition and Entrepreneurship, Chicago IL: University of Chicago Press.
- Klenke, Achim. 2013. Probability theory: a comprehensive course. 2nd ed. New York and Berlin: Springer.

- Knight, F. H. 1921. *Risk, Uncertainty, and Profit*. Boston, MA: Houghton Mifflin.
- Landgrebe J. and B. Smith 2023. *Why Machines Will Never Rule the World -- Artificial Intelligence without Fear*, New York: Routledge.
- Lapuschkin, Sebastian, Stephan Wäldchen, et al. 2019. Unmasking Clever Hans predictors and assessing what machines really learn. *Nature Communications* 10 (1): 1–8.
- Larson, Erik J. 2021. *The myth of artificial intelligence: why computers can't think the way we do*. Cambridge, MA: Harvard University Press.
- Legg, Shane, and Marcus Hutter. 2007. Universal intelligence: a definition of machine intelligence. *Minds and Machines* 17: 391–444.
- Pearl, Judea. 2020. The limitations of opaque learning machines. In *Possible minds: twenty-five ways of looking at AI*, edited by John Brockman, 13–19. New York: Penguin Books.
- Penrose, R. (1989) *The emperor's new mind: Concerning computers, minds, and the laws of physics*. Oxford: Oxford University Press.
- Prigogine, Ilya. 1955. *Introduction to thermodynamics of irreversible processes*. New York: Interscience Publishers.
- Prigogine, Ilya, and René Lefever. 1973. Theory of dissipative structures. In *Synergetics*, edited by H. Haken, 124–135. Wiesbaden: Vieweg+Teubner Verlag.
- Saba, Walid S. December 2022. Review of Landgrebe J. and B. Smith 2023, *Journal of Knowledge Structures & Systems*, Vol.: 3 Issue: 4 Pages: 38-41.
- Sautet, F. 2022. Alertness: an Aristotelian approach, in: Bylund op. cit., 39-63.
- Scheler, Max. 1961. *Man's place in nature*. New York: The Noonday Press.
- Scheler, Max. 1973. *Formalism in ethics and non-formal ethics of values*. Evanston: Northwestern University Press.
- Searle, J. R. 1980. Minds, brains, and programs. *BBS*, 3, 417–457.
- Smith, Adam. 1790. *Theory of moral sentiments, or an essay towards an analysis of the principles by which men naturally judge concerning the conduct and character*. 6th ed. The Strand, London: A. Strahan/T. Cadell.
- Thurner, Stefan, Peter Klimek, and Rudolf Hanel. 2018. *Introduction to the theory of complex systems*. Oxford: Oxford University Press.
- Turing, Alan. 1937. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* 42 (1): 230–265.
- Turing, Alan. 1950. Computing machinery and intelligence. *Mind* LIX: 433–460.
- Van Den Hauwe, L. (2009). *Foundations of Business Cycle Research, Volume II*, Saarbrücken: VDM Verlag Dr. Müller.

Van Den Hauwe, L. 2020. Entrepreneurship and Artificial Intelligence. unpublished manuscript.

Wolpert, David H., and William G. Macready. 1997. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1 (1): 67–82.

-