

Complexity and Particularity: An Argument for the Impossibility of Artificial Intelligence

EMANUELE MARTINELLI
University of Zurich

Abstract: Landgrebe and Smith (2022) have recently offered an important *mathematical* argument against the possibility of Artificial General Intelligence (AGI): human intelligence is a complex system; complex systems have some properties that cannot be modelled mathematically; hence we have no viable way to build an AI that would be able to emulate human intelligence. The issue of complexity is thus at the heart of the Landgrebe and Smith approach, and they tackle this issue by postulating a set of conditions, derived from mathematics, for a system to be complex. I argue that Hayek’s “Theory of Complex Phenomena” offers an interesting alternative path to understanding what makes complex systems complex. This offers an argument complementary to that of Landgrebe and Smith, which approaches the issue of complexity from an *ontological* angle, and which is designed to show that AI systems are not the right kind of entity that may interact meaningfully with complex systems.

Keywords: Artificial intelligence, AGI, Complexity, Agency, Particularity

1. INTRODUCTION

Artificial General Intelligence (AGI) is one popular name for the idea that AI systems may one day reach (Goertzel and Pennachin 2007)—or even exceed (Bostrom 2014)—human-level intelligence. Delimiting what this entails with precision is easier said than done. It essentially depends on one’s account of intelligence and of what marks out human intelligence in contrast with other possible forms of intelligent behavior. Turing (2009) Cite the original paper suggested replacing the question of whether a machine can possess human-level intelligence with the more tractable question of whether a machine can have the ability to convince some human interrogator that he or she is talking to a fellow human. Others focus on more substantial accounts of human intelligence, involving for instance the requirement of a genuine understanding of meaning as opposed to mere syntactic manipulation (Searle 1980).¹

Landgrebe and Smith (2022) (hereafter: ‘L&S’) stand close to the latter approach. They define what they call ‘objectifying intelligence’ as the characteristic human ability for a human:

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 (L&S, p. 46).

In essence, L&S argue that there are fundamental mathematical limits to the possibility of machine modeling real-world systems to the degree needed to mimic human-level AI. Thus they defend the thesis that:

it is impossible to obtain what we shall call 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 intelligence (L&S).²

This, they go on, ‘will prove that the lack of success in creating a general AI is not, as some claim, something that can and will be overcome by increasing the processing power and memory size of computers’ (L&S, p. 8).

Complexity

As mentioned already, what marks human-environment interaction is our capacity to interact with *complex* environments. In fact, unrestrained environments of the sort that we encounter in all kinds of scenarios—from taking the subway to walking on the street to having lunch in a restaurant to tracking a deer in the forest—all display the seven characteristic features of ‘complexity’ listed by L&S in their chapter 7, which is the heart of the volume. On the other hand, all of today’s AI applications, which realize so-called ‘narrow AI’, work in artificially restrained environments, such as abstract games (chess, GO) or controlled lanes for self-driving cars. The aim of this artificial restriction of the AI agent’s context is precisely to banish the complex manifestations of real-world environments. This statement includes the most recent developments of Large Language Models (LLMs) such as ChatGPT, which may compose applications with broader goals than, say, a chess-playing machine, but must nonetheless function in the artificial prompt-followed-by-response contexts determined by the user interface. Complexity thus plays a central role in the argument of the L&S book. We might advance the thesis that a core condition for intelligent behavior at the human level is the ability to interact adequately with novel complex systems. L&S thus dedicate a large amount of space to their treatment of what complex systems are, a treatment which draws on the seminal work of Thurner, Klimek, and Hanel (2018). The L&S definition of complexity lists a series of typical features of complex systems, features which preclude the creation of mathematical models of the sort that would be required if we would want to manufacture a computer with AGI. The L&S argument is mathematical in nature: there are principled limits to what we can do with the mathematical tools available to a Turing machine, and the emulation of systems with these features lies beyond what we can achieve with these tools.

Their work is of absolute importance and will hopefully transform the direction of research in the AI field at large in the future. However, I think that a deeper perspective on the notion of complexity may offer an argument against AGI which is complementary to what L&S provide. In this paper, I will work with a definition of complexity that is fundamentally compatible with that of Thurner, *et al.*, and yet has a more *ontological* flavor.

I will take as my starting point the account of complexity sketched by Hayek as a tool for his economic research. I believe that this is an interesting starting point to explore what makes a system complex in the first place, in order to ground the mentioned seven features of complexity that inhibit AGI engineering.

As a result, we will have the basis for an ontological argument against AGI which has the advantage that it will still hold even if we are one day confronted with a mathematical Copernican revolution that will revolutionize our current mathematics to the point where we can model complex systems mathematically after all.³

2. WHAT IS IT FOR A COMPLEX SYSTEM TO BE COMPLEX?

2.1 The Many Meanings of ‘Complexity’

To find a univocal definition of complexity is a very complicated issue in itself. It is not even clear whether elaborating a theory that would encompass all the various ways in which a system or a phenomenon may be classified as complex is a reasonable research goal (Mitchell 2009). Lloyd (2001) lists 42 different accounts of complexity, which capture both different aspects of complexity and different ways in which something may be complex. He groups them into three main categories:

1. Difficulty of description: How hard is it to describe?
2. Difficulty of creation: How hard is it to create?
3. Degree of organization: What is its degree of organization?

We share with L&S a concern with complexity as predicated of systems. We share also the definition of ‘system’ as a collection of ‘dynamically interrelated elements’ (L&S, p. 117). An agent’s environment, in this perspective, is indeed the system of all the elements that compose its surroundings and with which it may interact (Barandiaran et al. 2009; Beer 1995). What we need is thus something found in the third category of accounts. Unfortunately, the most popular notions of complexity, especially those found in the AI literature, seem to belong to the other two. ‘Computational complexity’, taken as ‘the amount of a computational resource (usually time or memory) that it takes to solve a class of problems’ (Edmonds 1999, p. 7), would fit into the ‘difficulty of creation’ category. Here, the complexity of a system is defined in terms of what it would take to *manufacture* the system, for example in such a way that would be able to undertake some specified tasks. There is however no evident relation between the complexity of the process needed to manufacture a machine and its level of intelligent behavior.⁴

Another popular notion of complexity is ‘Kolmogorov complexity’, which Lloyd (2001) classifies under the class 2—‘minimum description length’—heading. Kolmogorov complexity measures the complexity of a phenomenon according to the ‘minimum length of a Turing machine program needed to generate a pattern’ that describes the phenomenon in question (Edmonds 1999, p. 7).

More formally:

$$K(x) = \min l(p) \mid L(p) = x$$

where x is the described object, p is a computer program expressing x , $l(p)$ is the length of p , L is the language to perform the description, and x is the resulting description of the phenomenon (see Kolmogorov 1968). This account belongs to class 3, the ‘difficulty of description’ category, focusing on the complexity attached to a single element of a system—the program—rather than on a systematic assessment of the interactions between the AI agent and its environment.⁵

In this essay, however, we are dealing not with approaches based on computational aspects but rather with attempts to define complexity as a property of complex systems themselves. An account of this sort must deal with the degree of organization of a system. In particular, for the purposes of understanding the interaction between an AI agent and its environment, we need a notion that refers to the set of relevant circumstances in the agent’s surroundings.

2.2 Systemic Complexity

For Thurner, *et al.*, ‘complex systems are co-evolving multilayer networks’ (2018, p. 22). This definition simply repeats that ‘systems’ are ‘networks’ of elements that are interrelated, without delving into the question what makes these systems complex. Their approach is rather to provide an account of what makes a system

simple, thus enumerating complex systems as those systems that are *not simple*. I will attempt to address this issue in section 4. Thurner, *et al.* go on to define a handful of characteristic features that allow us to identify them as complex. Let us briefly review them:⁶

- 1) *Evolutionary processes.* The relations between the elements of the system and the overall states of the system emerge from the interaction between the system in question and surrounding systems, or the environment as a whole (Mitchell 2009). Evolution may bring about other characteristic features, like adaptivity, co-evolution, or the ‘edge of chaos’ effect (Edmonds 1999; Thurner et al. 2018).
- 2) *Path-dependency.* Present states of the system depend on past states; similarly, elements or sets of elements depend on other elements or sets of elements. Path-dependency, among other things, is an element of learning, understood as the ability of an agent to alter its actions based on past decisions or events (Russell and Norvig 2021). Lipman and Srivastava (1990) treat path-dependency as the central condition for complexity (working under a game theory approach).⁷

Non-Markovian character. As L&S highlight, this effect makes complex systems non-Markovian: their present states depend on an arbitrary number of past states, as opposed to the immediate previous state (so-called ‘Markov property’).

- 3) *Multi-level dynamic character.* The relations between the elements of the system are non-linear and may not be described by deterministic methods, e.g. systems of differential equations. We lack the ability to predict the next state of the system with certainty, and thus our modeling efforts need to resort to probabilistic means, such as stochastic methods (Goertzel and Pennachin, 2007; L&S). Meteorology or thermodynamics, for instance, exemplify this feature:

Non-linearity. Non-linearity is a consequence of the multilayer character of complex systems, in the sense that the various levels of descriptions pertaining to the system may influence each other in its development (Havel 1993; Thurner et al. 2018). For instance, the neurophysiological makeup of a person’s brain indeed impacts his psychological makeup to some extent, and possibly vice versa. This is to assert that not only the elements (or, in modeling terms, the parameters’ values) of the system vary, but even the relations between the elements (or the parameters themselves) may vary: ‘states change as a function (deterministic or stochastic) of the interaction network and, simultaneously, interactions change as a function of the states’ (Thurner et al. 2018, p. 23).

- 4) *Openness.* The boundaries of the system are not set a priori, and thus the elements and relations that compose the system are not pre-determined (Carissimo and Korecki 2023; L&S). It may be impossible to determine a priori whether this cloud in particular composes the meteorological system of London’s weather or not at this time.

Chaotic nature. As a consequence, complex systems are chaotic: due to the large number of variables and the indeterminacy of what even counts as a variable, it may be impossible to compute all the initial conditions of the system’s state at any given moment.

- 5) *Drivenness.* The system does not tend to a state of equilibrium as a rule, but its processes are fueled by an internal source of energy. The state changes because the system is pushed by this source of energy until it ultimately dries up (L&S; Thurner et al. 2018). The sources of energy can be either inanimate, as in a steam engine, or animate, as in organisms.

- 6) *Context-dependency*. The same elements may entertain different relations, or the same parameters describing the system may have some values under one description and some other values under another description. This structural ambiguity, for instance, allows the human sciences to study humans from the psychological, cultural, biological, and sociological points of view simultaneously, by considering one level of description at a time.
- 7) *Non-ergodicity*. The state of certain elements of the system (or values of parameters) cannot be found by averaging over an arbitrary set of past states of the system: this makes them ‘non-ergodic’ (L&S; North 1990). This bestows complex systems with potentially erratic behavior that ‘makes it impossible to model how and where in the system its elements distribute’ (L&S, p. 139). By contrast, in simple systems, usually, the distribution of states follows the typical Bell’s curve structure, so that we may calculate with acceptable certainty the probability that the value of a given parameter approached the average value at some moment in the past.

Structural unpredictability. As a consequence, in complex system the possibility that an unpredictable event, which could not have been possibly inferred from the present or past states of the system, ‘just happens’ cannot be ruled out in principle (Ayers 2017).

Let us consider an example that displays the seven properties to verify that it is a complex system: the state of the economy of the central bazaar in Marrakesh, Morocco. The system comprises all the individuals (workers, owners, entrepreneurs, customers), shops, relations, agreements, contracts, activities, taking place within the geographical boundaries of the central bazaar. The behavior of all the economic actors in the bazaar has influenced by past transactions, and depends on what has worked and what has not in the past (2). This has brought about the shape of all the activities, customs, and prices within the bazaar’s economy over time, based on the interaction between different actors and their interests (1). The decisions that give shape to the central bazaar derive from considerations on multiple levels that freely influence each other: economic considerations may be influenced by personal preferences or psychological states, for instance, e.g. personal acquaintances among customers and sellers (3). A single indecisive tourist, who is in doubt on whether he should spend the morning roaming across the streets of the bazaar, blurs the boundaries of the bazaar as a whole (4). The activities of the bazaar, in general, depend on the motivation and interests of the economic actors involved, so that, for instance, activities cease at night and resume at dawn (5). Economic evaluations may vary depending on the context, e.g. the price of sweets on a regular weekday versus on the eve of the final celebrations of the Ramadan (6). Finally, the actions and decisions of the individuals involved in the bazaar, or the composition of the bazaar itself, are intrinsically dependent on the structural possibility of unpredictable events—such as a sudden terrorist attack that forces the shops to close on rush hour, and investors to lose their money in ways they could not have possibly predicted (7).

2.3 Complexity of the Agent Versus Complexity of the Environment

In essence, L&S adopt this notion of complexity. Complex system, in this sense, yield features that radically resist mathematical modeling—e.g. non-Markovian character and non-ergodicity. Stochastic methods of the most advanced form must proceed from a dataset with clear probability distributions, the so-called ‘phase space’ of the model. The phase space represents the set of all the possible states of the system and is composed of the probability that every parameter of the model assumes any possible value. As long as all the manifestations of new events fall within the model’s phase space, then the model is in principle able to predict them, the accuracy of this probability being a mere technical problem. Unfortunately, in complex system, there is no guarantee that the next state of the system will fall within the phase space that has been valid up until this point in time (Felin et al. 2014). As in the previous central bazaar example, the dataset comprising all that happened and could happen within the market area need not contain a terrorist attack from a new group that formed overnight for unforeseen reasons. And this is a radical limit to the possibility to predict circumstances that, in fact, may prove to perturbate the system substantially.

AI systems of the most recent sort, like machine learning applications, function by being trained on huge datasets that represent a myriad of events that happened in a certain context. For instance, ChatGPT has been trained on millions and millions of conversations, statements, requests, or linguistic exchanges in several languages, in general, and now the system is capable of generating plausible answers to new inputs. However, any novel request that picks out things or events that could have not been included in that dataset is bound to produce very poor predictions or replies (Floridi 2023).

This kind of reasons substantiate the L&S argument against AGI:

- I. complex systems resist mathematical modeling,
- II. AI systems essentially are mathematical models,
- III. AGI would entail modeling complex systems (as opposed to Narrow AI that, as discussed, works with controlled, simple systems),
- IV. hence AGI is impossible.

Now, premise (III) needs further specification, of course. AGI shall not be possible because ‘the only way to engineer such technology is to create a software emulation of the human neurocognitive system’ (L&S). This means that the mathematical argument relies on impossibility to build a model of the source of human intelligence: the complex system behind general intelligence, being it the brain, the mind, or a human being as a whole. In the following, I want to develop a related claim: (III*) a general intelligent agent must be able to interact meaningfully with other complex systems. By ‘meaningfully’ I mean that the kind of response by the agent must be appropriate to the nature of the input, aiming at maintaining the agent’s goal in spite of environmental perturbations (Di Paolo 2005).

Both (III) and (III*) focus on complexity as applied to systems. (III), however, focuses in particular, we may say, on the complexity of the agent *qua* active system. (III*) focuses on the complexity of the environment *qua* interactive system, and demands that the agent be equipped with the ability to grapple with the environment’s complexity. That is, the agent must be the right kind of entity that can interact meaningfully with *any* manifestation of the environment, given the special cases warranted by complex systems. I will now formulate a new account of complexity that may accommodate claim (III*). In section 4, finally, I will claim that AI systems are not the right kind of entity that can meet this requirement.

3. WHAT MAKES COMPLEX SYSTEMS COMPLEX

To recap: Thurner, Klimek, and Hanel’s notion of complexity enumerates the features of complex systems. In doing this, they reveal the precondition required to have a synoptic and adequate model of complex systems. Thus, such notion nicely complements L&S’s mathematical argument (III) for conclusion (IV). Based on the works by Hayek on complexity (1967, 2014a), I will extract an alternative definition of complexity, which should be understood as a compatible, yet deeper take on Thurner, Klimek, and Hanel’s take. The focal point of this definition of complexity is what makes a complex system *complex*—and, as a consequence, what grounds the above features of complex systems (1-7): a complex system is a system whose synoptic description necessarily entails reference to a particular entity.

3.1 The Austrian School and the Theory of Complexity

A very interesting contribution to the notion of complexity exactly from the angle that I want to take has been suggested by the Austrian economist and philosopher Friedrich von Hayek. The intersection between Hayek’s works and the concept of complexity is mostly evident in his economic theory. Economic phenomena, he argues, are shaped by countless microeconomic transactions dependent on even more variables; as a consequence, social configurations like the market order cannot be designed or manipulated by any human mind or steering committee of human minds (Hayek 2014b). The market, Hayek argues, resists controllable

modeling precisely because it features the typical manifestations of a complex system. Hence, economic phenomena should be left to the spontaneous order emerging from uncoordinated interactions rather than to the design of planned orders like political institutions (Hayek 1982).

Less attention is typically offered to the preliminary work Hayek has dedicated to the notion of complexity per se. In *The Theory of Complex Phenomena* (Hayek 1967), Hayek introduces a class of ‘essentially complex phenomena’ that are typically investigated by the social sciences (among which economics and psychology), linguistics, or evolutionary biology for instance—in contrast to ‘simple phenomena’ that are the subject of the hard sciences like mathematics and physics. Complex phenomena intrinsically yield a lower ‘degree of falsifiability’ which makes prediction possible only at the level of ‘certain general features of a situation which may be compatible with a great many particular circumstances’ (Hayek 1967, pp. 28-29). Statistics is the method we typically use to treat complex phenomena in response to the lack of clear data on ‘the relations between the individual elements with different attributes’ (Hayek 1967, p. 30)—although, Hayek warns us, even these devices yield limited results in terms of predictability and falsifiability. In *The Pretense of Knowledge* (Hayek 2014a), before deploying the notion to illustrate the limits of political economy as a discipline, Hayek depicts the features of complex phenomena in a way not too far from Thurner, Klimek, and Havel (2018):

- ‘The aspects of [essentially complex phenomena] to be accounted for about which we can get quantitative data are necessarily limited and may not include the important ones’ (p. 363).
- They yield ‘structures whose characteristic properties can be exhibited only by models made up of relatively large numbers of variables’ (p. 365).
- ‘Complexity here means that the character of the structures showing it depends not only on the properties of the individual elements of which they are composed, and the relative frequency with which they occur, but also on the manner in which the individual elements are connected with each other’ (Ibid.).

Finally, Hayek remarks that ‘ascertaining all the data determining a particular manifestation of the phenomenon in question [is] a difficulty which is often insurmountable in practice and sometimes even an *absolute* one’ (Hayek 1967, p. 27; emphasis added). By ‘absolute’ Hayek refers to an argument, delivered in *The Sensory Order* (Hayek 1976, pp. 184-190), that it is logically impossible for any agent *x* to reproduce, manipulate, or fully comprehend any entity *y* that is more complex than *x*. Very clearly, this claim resonates with L&S statement (3) against AGI. We can find traces of the same line of reasoning, very briefly, in other Austrian authors—to remark the importance that Austrian economic theory played in the formulation of this view of systemic complexity.⁵ Mises, in *The Human Action* (Mises 1996, pp. 129-142), clearly distinguishes between two uses of probability in grappling with the radical uncertainty of the future. In one sense, ‘class probability’, the likelihood of events of a certain set of possibilities (today we would say ‘phase space’ of a the problem) can be mathematically estimated; in the sense that has to do with the way that individuals actually deal with the unpredictability on reality in their actions, however, ‘case probability’ escapes mathematical tractability as it refers to ‘individual, unique, and nonrepeatable’ cases (Mises 1996, p. 135)—in a very similar fashion to what I have called ‘unpredictable events’.

Israel Kirzner, throughout his theory of entrepreneurship, insisted on a similar note that the characteristic ability of the individual human agent in society is the ‘alertness’ to unforeseen errors in the allocations of resources, whose correction brings profit to the entrepreneur (Kirzner 1963, 1997). Felin et al. (Felin et al. 2014), as we shall see in a minute, move precisely from the Austrian view of entrepreneurship as interaction with a radically unpredictable world to generalize the complexity of economic systems unto the complexity of other domains, e.g. biological evolution.

3.2 Particularity as the Key to Complexity

Now that we have established that Hayek and others in the Austrian tradition have more or less implicitly defended employed a notion of complexity similar to one that suits our needs, let me elaborate on this definition. I find that the Austrian sources that I have reported offer a basis to develop an account of complexity to spell out what actually makes a complex system complex. In short, a system is complex iff any synoptic description of it necessarily entails reference to at least one particular entity. The reader might be surprised to note that the promised ‘ontological’ definition of systemic complexity focuses on how we *describe* the system, rather than saying something substantial of the system in and of itself. However, this comes from the fact that a ‘system’, properly speaking, is not a substance, with an independent nature and existence. A system is a set of interrelated substances that the observer may stipulate to grasp their emergent properties (Smith 2000). There is no such thing as the ‘Milan Malpensa International Airport’, but rather what we are dealing with is a collection of substances (e.g. airplanes, machines, infrastructural components, workers, passengers, police forces, companies, etc.), relations (e.g. distances, hiring contracts, roles, assignments of seats), and processes (e.g. departures, arrivals, delays, exchanges, etc.) that, taken as a whole, may be said to have certain properties (e.g. be tidy, efficient, soon to be renovated, etc.). New planes or workers may come and go, facilities may be open and closed, and we would regard the Milan Malpensa International Airport as the same system over time, while undergoing substantial transformations.⁸ Indeed, this reading of systems as fiat entities, or quasi-abstract objects, is shared by L&S (pp. 117-119) as well. Hence, a proper definition of what qualifies a *system* as complex must reflect the way we qualify fiat entities, i.e. the way we describe them.

Similarly to L&S, I take a synoptic description of *x* to be a complete description of *x* that represents at least all the substances, relations, and processes necessary to represent all of *x*’s aspects. The definition means that, whenever it is necessary to make reference to one particular entity, be it an object or an event, to obtain a synoptic description of a system *x*, then *x* is a complex system. Let us make an example. A ball rolling off a slide, as represented in physical models for instance, is a *simple* system. To have a description of all that is going on and defines the system, it is only necessary to say that any ball (anything belonging to the kind of ‘balls’) interacts with any slide (*ibidem*) by entertaining any process that could qualify as ‘rolling off’—however we may define those terms. On the contrary, the economy of Marrakech’s central bazaar is a complex system. It is impossible to convey a complete description of the economy of Marrakech’s central bazaar without, say, making reference to the Mr. Al-Amin’s historic shoe stand located in the North-Western corner of the market’s area, or its decision to buy an extra stock of Adidas on March 31, 2023. That is, without information about *particular* objects or events, as opposed to general or universal kinds, we would not be describing the economy of Marrakech’s central bazaar.

This definition that views particularity as the ground for complexity is never actually spelt out by Hayek himself, but I believe it to be the logical conclusion of his intuitions. In his own words, for instance, ‘a theory of essentially complex phenomena must refer to a large number of particular facts; and to derive a prediction from it, or to test it, we have to ascertain all these particular facts’ (Hayek 2014a, p. 370). As of why particularity is the distinguishing feature that makes systems complex, this is because it is the need to reference particular entities that bounds the intrinsic capacity of that system to be captured by our models with the same predictive power that the models of hard sciences, dealing with simple systems, have.

3.3 The Particular Character of Structural Unpredictability

As one can easily see, many complex systems may be approximated by simple systems.⁹ For instance, we can say that the ball-rolling-off-a-slide physical system represents a general version of my old blue soccer ball rolling off my hometown’s slide in the central park on a summer day of 2005. Conversely, economic models usually describe a generalized version of real-world markets, abstracting away particular shops or actors and just postulating that certain assets lie in the hands of rational and self-interested human beings.

This possibility to artificially constrain systems so that they become simple is at the heart of the fact that we can actually have scientific knowledge of the world (Hayek 1952). For example, we may derive from a generic economic model general tendencies like a prediction that the price of shoes will go up if a stock gets destroyed. Yet, the prediction of particular facts about unrestrained, real-world complex systems is beyond our capacity.

Hayek, in fact, reports that of complex phenomena we can only have knowledge of ‘the sort of pattern that will appear and not its particular manifestation’ (Hayek 1967, p. 28). Again, ‘without specific information about the individual elements we shall be confined to [...] predictions of some of the general attributes of the structures that will form themselves’ (Hayek 2014a). The point is that there are some real-world systems, or ‘phenomena’, that just require reference to particular entities. Generalized, simple versions of them are not synoptic descriptions of them—just like generic economic models do not represent to the fullest the economy of Marrakech’s central bazaar in particular.

Thus, for those complex systems where particularity is eliminable only at the price of describing a different system altogether, we have a structural form of unpredictability. This points to a somewhat paradoxical result: of simple systems, where reference to the level of the particular is not necessary, we can have predictions of particular facts (e.g. that the ball hit the bottom, at such and such conditions, at x velocity); of complex systems, where reference to particular entities is necessary, we can only have predictions of general facts (e.g. that the price of shoes will be pushed up by the disruption of a stock decreasing offer). In Hayek’s own example, biological evolution may teach us that horses will not grow wings in the next hundred years, and yet we cannot say anything certain about the mutations that this particular horse will transmit to her offspring next year (Hayek 1967).

This reduced degree of predictability is not only a matter of practical feasibility of predictions: there is something structural in complex systems that make them unpredictable. This Mises had in mind with his distinction between ‘class probability’ and ‘case probability’. The uniqueness of circumstances then grounds, among other things, condition (7) of the Thurner-Klimek-Havel notion of complexity—and, from here, all the other manifestations of complex systems that make them intractable with our mathematical tools. The presence of particular objects or events in the picture, and the radical possibility that new particular objects or events enter the picture at some point, leaves us with the evergreen possibility that something that could not have been possibly inferred will happen. As Felin et al. put it: ‘the problem is not only one of comparison among the best uses and functions of objects and spaces, but even the very generation of the full list is not algorithmically feasible’ (Felin et al. 2014, p. 274). The fact that complex systems are resistant to generalization and abstraction, in other terms, make ‘pertinent variables shift constantly’ (p. 277) and thus prevent any model’s phase space from fully depicting the synoptic set of all possible states that the system could ever assume.

4. THE RELATION WITH THE PRE-INTENTIONAL WORLD

The ontological argument, thus, is the following. We indeed may predict some general facts, or sets of possible states if you will, of a complex system by artificially restraining our description of it. For example, meteorologists do this when they infer that the high-pressure front flying over Italy next week will likely bring rain, even without bothering to measure the exact position and velocity of all the water molecules and wind currents involved in the particular occasion. However, there are some cases in which general predictions of simple systems are not enough. When an agent is faced with a real-world environment, the agent needs to be capable of fully interacting with it. In other words, it must have the ability to tackle the unpredictable, particular manifestations of the system—as opposed to manipulating general rules over sets of possible states of the system.

Now, I will argue in this section that AI systems are not entities of a sort that can do this. The capacity to meaningfully interact with complex system pertains to organisms that are the offspring of evolutionary

processes themselves. AGI is based on mathematical models which are different in kind from evolutionary processes, thus AGI intrinsically cannot interact meaningfully with complex systems.

4.1 The Intentional and the Pre-Intentional World

Objectifying intelligence, as the term suggests, entails the capacity by the intelligent agent to treat the elements of its environment (including itself) as objects. In other words, it presupposes that we develop and maintain intentional states—intentionality being classically defined as ‘that property of many mental states and events by which they are directed at or about or of objects and states of affairs in the world’ (Searle 1983, p. 1). Intentional states include beliefs, desires, fears, thoughts, and possibly (even though there is no point in defending this extension here) functionally analogous states that are non-mental but rather reside, somehow, inside a computer.

Now, it is worth mentioning that the definition of any intentional state relies on an indefinite number of related intentional states, forming a ‘network’ of intentional states. Some of these are explicit, i.e. consciously present to the subject’s mind, while others remain implicit, unconscious, until they become relevant enough to be explicit (Searle 1983, pp. 140-142). In just the same way, my belief that China is bigger than Japan, or my desire to eat ramen on a regular basis, usually stay submerged outside the realm of my consciousness, just until empirical experience or rational thought make considering those intentional states useful (Popper 1995). Some assumptions in the network of intentional states are so fundamental that they rarely are called to consciousness: Mark will open the door likely without thinking to himself ‘Beware: I have to pull with some force because this door will resist my touch’.

Eventually, Searle suggests, this network of intentional states is grounded on a ‘bedrock of mental capabilities that do not themselves consist in Intentional states (representations), but nonetheless form the preconditions for the functioning of Intentional states’ (Searle 1983, p. 143). This ‘background’ of intentionality is the setting for all the nonrepresentational contents of the mind that are pre-intentional and ground intentional states. What is the nature of the contents that compose the background of intentionality? Searle himself suggests that these nonrepresentational mental contents are forms of non-propositional, practical knowledge:

In order that I can now have the Intentional states that I do I must have certain kinds of know-how: I must know how things are and I must know how to do things, but the kinds of “know-how” in question are not, in these cases, forms of ‘knowing that’. [...] It is important to emphasize that there is no sharp dividing line between ‘how things are for me’ and ‘how I do things’. It is, for example, part of my preintentional stance towards the world that I recognize degrees of the hardness of things as parts of ‘how things are’ and that I have numerous physical skills as part of ‘how to do things’. But I cannot activate my preintentional skill of, say, peeling oranges independently of my preintentional stance toward the hardness of things (Searle 1983, p. 143).

This background of pre-intentional material represents the direct relation between the subject and the world. The condition of possibility of the interpretation and elaboration of experience is, indeed, the experience of the world. For example, think of what the case must be in order for Mark to form the intention to go and pet his cat: he must know how to stand, how to move through space, how to walk, how to recognize his cat, how to pet it, how to pet it in such a way that it is pleased and not bothered, and so on. All these capabilities are necessary to form and pursue the intentions of agents and help to compose their actions. As a consequence, the network of intentional states is grounded on the agent’s relation with the particular entities that make up the environment.

The fact that agent *x* is able to develop intentional state *y* toward element *z* is ultimately grounded in the ability of agent *x* to entertain a relation with *z*. For instance, I can desire to own a football because I have perceived, kicked, or heard of actual footballs in the past.¹⁰ And it is fundamental to note that intentional

states, which compose the basis of objectifying intelligence, depend on actual footfalls and other *particular* entities. To take stock, then, we may advance the thesis that the precondition for the meaningful interaction with real-world environments that is of objectifying intelligence, is in turn grounded on an immediate relation between the agent and the pre-intentional world of particular entities in the agent's environment.

4.2 The Impressive Power of Organisms

So, what does it take for something to develop the ability to entertain an immediate relation to the pre-intentional world? I am going to defend what has been called the 'organic view' (Torrance 2008): there is something special about organisms that endows them with a unique degree of flexibility in their interaction with complex systems. Thus, organisms are the right kind of entity to develop an immediate relation with the environment to the level of particular entities, thereby tackling the unpredictable manifestations of complex systems meaningfully.

Organisms are self-organizing and self-maintaining entities that are composed by a plurality of functional units, i.e. organs, whose different functions coordinate to the conservation of the whole system. Any living organism, as simple as it may be, possesses a number of specialized functional units that develop specific functions, with the aim to serve the conservation of the whole. The heart has the function to deliver oxygen and nutrients to all parts of the organism by pumping blood; the lungs have the function to draw in air that will keep the blood of the organism oxygenated; the liver has the function to extract toxins from the blood to keep the organism healthy, etc. This distribution of functional units is organized by the organism itself, meaning that any organism has feedback-response mechanisms allowing it to adjust its internal composition responsively with respect to the perturbations of the environment. This view on organisms is based on the definition by Maturana and Varela (1980) of organisms as 'autopoietic machines':

An autopoietic machine is a machine organized (defined as a unity) as a network of processes of production (transformation and destruction) of components that produces the components which: (i) through their interactions and transformations continuously regenerate and realize the network of processes (relations) that produced them; and (ii) constitute it (the machine) as a concrete unity in the space in which they (the components) exist by specifying the topological domain of its realization as such a network. It follows that an autopoietic machine continuously generates and specifies its own organization through its operation as a system of production of its own components, and does this in an endless turnover of components under conditions of continuous perturbations and compensation of perturbations (Maturana and Varela 1980, pp. 78-79).

To be an organism is to instantiate an autopoietic system. The core difference between a *living* organizational structure and any machine or merely physical mechanism is that the former's processes are essentially directed to the production of the very components of the structure. Whereas my heart's and my lungs' functions are finalized to the preservation of my bodily parts, a car, or even a driverless car, produces effects that are independent from the goal of preserving the machine (Maturana and Varela 1980).

One of the most important consequences of the organism's ability to redirect the organs' functions is the capacity to distance itself from the environment. This is the establishment of a sense of *individuality*, e.g. a living cell generating a membrane to separate itself from the surroundings (Burge 2009; Maturana and Varela 1980). This, in turn, comes together with an essentially *autoreferential* functioning (Maturana & Varela 1980): the organism is capable of redirecting the functions of its organs to the preservation of the whole, as distinct from the environment, by creating new goals and assigning a new meaning and importance to any *particular* element of the environment depending on the particular circumstances. Hence, it is the sophisticated emergence of a centralized control system of a cluster of biological functions that allows the organism to really entertain a relation with particular entities in the environment. In essence, the organic structure is the precondition for entertaining an immediate relation with the pre-intentional world.

4.3 Artificial Agents and Mediated Intentionality

I have insisted that the flexibility required to grapple with complexity and unpredictability comes from an *immediate* access to the pre-intentional world. This qualification comes from the fact that artificial agents, among which, most notably, AI applications, have a *mediated* access to the pre-intentional world. And this, in turn, is motivated by their structure that is different from the makeup of organisms. List (2021) and Laukyte (2017) have shown that AI agents and collective agents (e.g. clubs, nations, firms, armies) have in common that they enjoy a form of ‘bounded autonomy’ restricting their ability to interact with their environment.

I would further claim that, due to their structure, their intentional states do not emerge spontaneously from a direct relation with the particular components of their environment. Properly speaking, AI agents (as well as collective agents) interact with kinds and patterns, rather than with the actual objects and events in the world. For instance, an automatic translator like Deepl.com delivers its output based on the patterns generated by the model and its training data set, rather than considering the particular instances of the utterances it is confronted with in the particular context at hand. More in depth, List and Pettit (2011) suggest that, in the case of collective agents, attitudes are generated by some ‘attitude aggregation mechanism’; the appropriate output intentional state is generated by the group’s members elaborating on a series of other pre-determined intentional states, so that the collective agent never creates genuine intentional states grounded on pre-intentional states. Collective agents form collective intentions or beliefs, for instance, based on the intentions and beliefs of their individual members, as they are aggregated following a certain algorithm (e.g. majority voting). This reasoning may easily be extended to all sorts of artificial agents. AI agents, in the same way, form automatic intentions or beliefs by generalizing on the intentions or beliefs crystallized in the training data set, or otherwise encapsulated in the instructions of their source code.

As we see, artificial agents simply act in a different way. This limits the quality of their interaction with the environment, and ultimately prevents a meaningful interaction with truly complex systems. Why so? The short answer is that AI agents (and collective agents too)¹¹ are not organisms. They do not spontaneously centralize a cluster of functional units: they merely execute functions, however convoluted their algorithms or foundational goals might be. The longer answer is that they are not brought about by the same kind of process.

Organisms, as far as we can observe today, are formed by the workings of biological evolution. As Hayek himself points out, evolution is the spontaneous result of a myriad of interactions in a myriad of different contexts, whose result is so complex in nature that it cannot be designed by any one of the agents taking part (Hayek 1967, 2014b; Vasconcelos Vilaça 2010). AI agents, on the other hand, are not formed by any means through evolution. They are essentially optimizing algorithms, which execute mathematical functions that try to maximize some given variable or set of variables (Carissimo and Korecki 2023). Evolution does not optimize, and, in fact, it is quite often the case that biological organisms endure trial and error processes where mutations fail to bring about the most optimal adaptation to given problems (Felin et al. 2014).

Optimization, then, is the process that characterizes mathematical models, which suffer from all the problems set forth by L&S, and by myself in the foregoing. And just to mention them, Carissimo and Korecki (2023) identify three kinds of limitations of optimizing algorithms with respect to evolutionary processes:

- *Object limits*, concerning the model and the subject of optimization. For instance, a sociological model may not capture all the aspects of the subject’s behavior, namely the dynamics of the social group in question.
- *Objective limits*, concerning the maximizing function. For instance, the optimization of subjective well-being within the social group might be flawed by the difficulty we face in quantifying subjective well-being where different individuals are involved.
- *Process limits*, concerning the impact that the optimizing process itself may have on the subject of optimization. For instance, the optimizing process may involve practices or policies (e.g. re-

distribution of resources within the group) that incentivize people to conceal their assets or their preferences, thereby invalidating the model.¹²

5. CONCLUSION

To sum up, I have delivered the following argument against the possibility of AGI. According to my definition of a complex system, a system is complex iff any synoptic description of it necessarily entails reference to at least one particular entity. Human-level intelligence presupposes the ability to interact meaningfully with complex systems. Hence, human-level intelligence presupposes the ability to interact meaningfully with particular entities. Now, to interact meaningfully with particular entities requires an immediate relation with the pre-intentional world. Only organisms entertain an immediate relation with the pre-intentional world. AI systems are not organisms, both because a) their intentional states are the result of a relation with the pre-intentional world mediated by other (human) agents' attitudes, and because b) they are optimizing algorithms rather than the product of evolutionary processes.

As for L&S so also here, the impossibility of AGI is not advanced as an empirical fact: there are a priori reasons why machines will never acquire human-level intelligence, creativity, and flexibility regardless of our technological progress in the future. I say that an ontological argument is more fundamental than a mathematical one, however, because the latter can in principle fall prey to the objection that, were humanity to experience a major breakthrough in mathematical modeling, the limitations on the plausibility of AGI might eventually be lifted. Conversely, the ontological argument concerns not our capacity to tackle engineering or technological problems but how things are in the world, independently of human interventions.

To be sure, AI systems are nevertheless an incredible technological achievement and will be absolutely useful and influential in a number of future endeavors. After all, for the majority of applications, it is sufficient that an AI system can fare better than the average human being at the same activity. The argument simply entails that there will be no general artificial intelligence and that research and development should focus on narrow artificial intelligence, i.e. applications tailored to treat particular tasks in well-defined, constrained environments. As a consequence, humans will not and should not be pushed completely out of the loop for those applications of AI where systems will be required to operate in complex, real-world environments, e.g. self-driving cars or medical diagnostics. In sum, this paper more or less confirms the empirical results of those studies claiming that the automatization of work will only concern those fields where human creativity (or the human 'touch') is not fundamental (Eloundou et al. 2023).

One final word should be spent on replying to one still unanswered question. What if we managed to build AI machines that are not built up like optimizing algorithms but actually reproduce the structure of organisms? This approach has been brought forth by two separate streams of research: the 'evolutionary/organic programming' (Bonabeau et al. 1999; Nakashima 1999) project and the 'artificial life' or 'A-life' project (Ray 1991). Both projects attempt to either emulate evolutionary processes or to build up algorithms or to train artificial agents by recreating evolutionary settings. My argument, as I framed it, is dependent on the success of these projects, and on the possibility to actually create artificial organisms. To tackle the chances of success of these projects would be out of the scope of the paper at this point (but see Carissimo and Korecki (2023) and L&S themselves for reasons to think these approaches are fundamentally a dead end). If anything, my argument would still prove that those would be the only paths forward toward building AGI in the future.¹³

NOTES

- 1 For L&S, AIs cannot understand anything at all, not even to the level required for the sort of ‘primal intelligence’ exhibited already by primates (pp. 55ff.; see also Landgrebe and Smith 2019).
- 2 A model of a system is synoptic if it can be used to engineer that system or to emulate its behaviour. A model is adequate relative to some set of requirements if it can be used to engineer an entity or an emulation that satisfies all requirements in that set (309f.).
- 3 As L&S point out (p. xii), it is extremely unlikely that this might happen. For it would require an overturning of our current mathematical models that would be as revolutionary as the achievements of Leibniz and Newton in inventing the differential calculus. For instance, it would be necessary to devise ways to treat mathematically systems whose structure, expressed in terms of the relations between their parameters, is allowed to vary freely across time.
- 4 Logically speaking, that is. Whether it is empirically necessary that adaptation to a complex system is itself complex in some sense is left open.
- 5 Let alone the fact that this account only works for the digital artifacts with which programs are associated, thus lacking generalization.
- 6 L&S (pp. 138-139) too identify a list of features of complex systems, although different from the one I use here.
- 7 Strategy s is more complex than s' iff the perturbations from the equilibrium path inflicted by past strategies is greater for s than for s' . More formally:

$$K_{P(i)}(s, h^*) = \sum_{(t, h)} P_i(t, h, h^*)$$

where P_i is the weight attributed by player i , s is the relevant strategy, h is the history of token activities t_n performed in past occurrences of the game, (t, h) stands for the present act t coupled with history h , and h^* is the equilibrium path or the history ideally leading to equilibrium (Lipman and Srivastava 1990, p. 279).

- 8 For more on this, cf. Moreno-Casas (2022).
- 9 One may say that there are systems of interactions between particulars that are predictable to some extent, and hence would qualify as simple systems, e.g. the orbit of the Earth around the Sun. As I will imply in what follows, as L&S also suggest, it is not that any complex system comprising particulars are utterly unpredictable under all of their respects. We may indeed predict some aspects of complex systems, especially those that can be traced to an approximation of the complex system as an abstract simple system. For instance, this is what the physician does when recognizing the symptoms of the seasonal flu on the inside of your own very particular throat. Or, to say with the mentioned example, the rate of solar eclipses on Earth due to the Earth-Sun-Moon system is totally predictable because it can be captured by the physical, simple system of three orbiting bodies with given masses and distances between them. The real-world Earth-Sun system’s predictability is limited though: there is no reliable way to predict the effect of the next solar storm’s twentieth X-ray pinnacle on this particular TV broadcasting antenna in New York. Many thanks to Janna Hastings for pointing this out.
- 10 To qualify the relation between the agent and the pre-intentional world of particular entities within the agent’s environment is beyond the scope of the paper. But a possibility, which Searle himself seems to implicitly take into consideration (while leaving the question open as well), is that the sort of practical knowledge about particular entities is the agent’s (qualitative) experience of particular entities (Lewis 1999) – to reinforce the Aristotelian proviso that passions of the soul proceed from sensibility.
- 11 Of course, this argument is disproven if one assumes that social groups may acquire the degree of sophistication of social organisms under the right conditions. Those who support this claim say that social groups, with enough advanced institutionalization can become ‘corporate persons’ with a will and attitudes of their own, over and above their individual members. And, coherently, the proponents of this view readily identify the need to protect the interests and capacities of social persons in a degree comparable to how we protect the interests and capacities of individual human persons (a view sometimes called ‘collectivism’). See for instance Luhmann (1995) or Scruton and Finnis (1989).
- 12 Cf. L&S’s notion of ‘model-induced escape’.
- 13 Many thanks to the University of Zurich and the Digital Society Initiative for funding this project.

REFERENCES

- Ayers, M. R. 2017. *The Refutation of Determinism*. New York and Abingdon: Routledge. <https://doi.org/10.4324/9781315188706>
- Barandiaran, X. E., Di Paolo, E., and Rohde, M. 2009. Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17(5): 367-386. <https://doi.org/10.1177/1059712309343819>
- Beer, R. D. 1995. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1-2):173-215. [https://doi.org/10.1016/0004-3702\(94\)00005-L](https://doi.org/10.1016/0004-3702(94)00005-L)
- Bonabeau, E., Dorigo, M., and Theraulaz, G. 1999. *Swarm Intelligence - From Natural to Artificial Systems*. Santa Fe Institute Studies in the Sciences of Complexity.
- Bostrom, N. 2014. *Superintelligence. Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Burge, T. 2009. Primitive Agency and Natural Norms. *Philosophy and Phenomenological Research*, 79(2):252-278.
- Carissimo, C. and Korecki, M. 2023. Limits of Optimization. *Minds and Machines*. <https://link.springer.com/article/10.1007/s11023-023-09633-1>
- Di Paolo, E. A. 2005. Autopoiesis, Adaptivity, Teleology, Agency. *Phenomenology and the Cognitive Sciences*, 4(4):429-452. <https://doi.org/10.1007/s11097-005-9002-y>
- Edmonds, B. 1999. What Is Complexity? The Philosophy of Complexity Per Se with Application to Some Examples in Evolution. In F. Heylighen and D. Aerts (eds.), *The Evolution of Complexity* (pp. 1-13). Boston: Kluwer.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. 2023. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *ArXiv:2303.10130v4*. <http://arxiv.org/abs/2303.10130>
- Felin, T., Kauffman, S., Koppl, R., and Longo, G. 2014. Economic opportunity and evolution: Beyond landscapes and bounded rationality. *Strategic Entrepreneurship Journal*, 8(4), 269–282. <https://doi.org/10.1002/sej.1184>
- Floridi, L. 2023. AI as Agency Without Intelligence: On ChatGPT, Large Language Models, and Other Generative Models. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4358789>
- Goertzel, Ben and Pennachin, Cassio. 2007. *Artificial General Intelligence*. Rockville, MD: Springer.
- Havel, I. 1993. Artificial Thought and Emergent Mind. *Proceedings of the 13th International Joint Conference on Artificial Intelligence*.
- Hayek, F. 1952. *The Counter-Revolution of Science. Studies on the Abuse of Reason*. Glencoe: Free Press. <https://doi.org/10.2307/2549519>
- _____. 1967. The Theory of Complex Phenomena. In *Studies in Philosophy, Politics, and Economics* (Vol. 1, pp. 22-42). Chicago: University of Chicago Press.
- _____. 1976. *The Sensory Order. An Inquiry into the Foundations of Theoretical Psychology* (2nd ed.). Chicago: University of Chicago Press.
- _____. 1982. Cosmos and Taxis. In: *Law, Legislation and Liberty. A New Statement of the Liberal Principles of Justice and Political Economy* (Vol. 2, pp. 35–54). New York and Abingdon: Routledge.
- _____. 2014a. The Pretense of Knowledge. In: B. Caldwell (ed.), *The Market and Other Orders* (pp. 362-372). Chicago: University of Chicago Press.
- _____. 2014b. The Results of Human Action but Not of Human Design. In: B. Caldwell (ed.), *The Market and Other Orders* (pp. 293-303). Chicago: University of Chicago Press.
- Kirzner, I. 1963. *Market Theory and the Price System* (1st ed.). New York: D. Van Nostrand Co., Inc.
- _____. 1997. Entrepreneurial Discovery and the Competitive Market Process: An Austrian Approach. *Journal of Economic Literature*, 35:60-85.
- Kolmogorov, A. 1968. Logical Basis for Information Theory and Probability Theory. *IEEE Transactions on Information Theory*, 14(5):662-664.
- Landgrebe, J., and Smith, B. 2019. *Making AI Meaningful Again*. <https://Arxiv.Org/Abs/1901.02918>.
- _____. 2022. Why Machines Will Never Rule the World. Artificial Intelligence without Fear. In: *Why Machines Will Never Rule the World: Artificial Intelligence without Fear*. New York and Abingdon: Routledge. <https://doi.org/10.4324/9781003310105>
- Laukyte, M. 2017. Artificial Agents Among Us: Should We Recognize Them as Agents Proper? *Ethics and Information Technology*, 19(1):1-17. <https://doi.org/10.1007/s10676-016-9411-3>
- Lewis, D. 1999. What Experience Teaches. In: *Papers in Metaphysics and Epistemology* (pp. 262–290). Cambridge: Cambridge University Press.
- Lipman, B. L., and Srivastava, S. 1990. Informational Requirements and Strategic Complexity in Repeated Games. *Games and Economic Behavior*, 2(3):273-290. [https://doi.org/10.1016/0899-8256\(90\)90022-M](https://doi.org/10.1016/0899-8256(90)90022-M)
- List, C. 2021. Group Agency and Artificial Intelligence. *Philosophy and Technology*, 34(4):1213-1242. <https://doi.org/10.1007/s13347-021-00454-7>
- List, C. and Pettit, P. 2011. *Group Agency. The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Lloyd, S. 2001. Measures of Complexity: A Non-Exhaustive List. *IEEE Control Systems Magazine*, 21(4):7-8. <https://doi.org/10.1109/MCS.2001.939938>

- Luhmann, N. 1995. *Social Systems*. Redwood City: Stanford University Press.
- Maturana, H. R. and Varela, F. J. 1980. *Autopoiesis and Cognition* (Vol. 42). Dordrecht: Springer. <https://doi.org/10.1007/978-94-009-8947-4>
- Mises, L. Von. 1996. *Human Action. A Treatise on Economics* (4th ed.). New York: Fox & Wilkes.
- Mitchell, M. 2009. *Complexity: A Guided Tour*. Oxford: Oxford University Press. <https://doi.org/10.1063/1.3326990>
- Moreno-Casas, V. 2022. Ludwig von Mises and Complexity Economics. *SSRN Electronic Journal*, <https://ssrn.com/abstract=4211886>. <https://doi.org/10.2139/ssrn.4211886>
- Nakashima, H. 1999. AI as Complex Information Processing. *Minds and Machines*, 9: 57-80.
- North, D. C. 1990. *Institutions, Institutional Change and Economic Performance*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511808678>
- Popper, K. 1995. *Knowledge and the Body-Mind Problem: In Defence of Interaction* (1st ed.). New York and Abingdon: Routledge.
- Ray, T. 1991. An Approach to the Synthesis of Life. *Physics*. <https://www.semanticscholar.org/paper/An-Approach-to-the-Synthesis-of-Life-Ray/2e140a4435f22a04570e35b1b442a22ae90053bd>
- Russell, S. and Norvig, P. 2021. *Artificial Intelligence: A Modern Approach* (4th ed.). London: Pearson.
- Scruton, R. and Finnis, J. 1989. Corporate Persons. *Aristotelian Society Supplementary Volume*, 63(1):239-274.
- Searle, J. R. 1980. Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3):417-424. <https://doi.org/10.1017/S0140525X00005756>
- _____. 1983. *Intentionality. An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Smith, B. 2000. Objects and Their Environments: From Aristotle to Ecological Ontology. In: *Life and Motion of Socio-Economic Units* (pp. 84-102). London: CRC Press. <https://doi.org/10.1201/9781482268096-14>
- Turner, S., Klimek, P., and Hanel, R. 2018. *Introduction to the Theory of Complex Systems* (Vol. 1). Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198821939.001.0001>
- Torrance, S. 2008. Ethics and Consciousness in Artificial Agents. *AI and Society*, 22(4):495-521.
- Turing, A. M. 2009. Computing Machinery and Intelligence. In: *Parsing the Turing Test*, pp. 23-65. New York: Springer.
- Vasconcelos Vilaça, G. 2010. From Hayek's Spontaneous Orders to Luhmann's Autopoietic Systems. *Studies in Emergent Order*, 3:50-81.