ANNALES

PROCEEDINGS OF THE ACADEMY OF SCIENCES OF BOLOGNA

CLASS OF PHYSICAL SCIENCES

Bologna
University Press

ANNALES

PROCEEDINGS OF THE ACADEMY OF SCIENCES OF BOLOGNA

CLASS OF PHYSICAL SCIENCES

3

Annales. Proceedings of the Academy of Sciences of Bologna
Class of Physical Sciences

Fondazione Bologna University Press
Via Saragozza 10, 40123 Bologna
tel. (+39) 051 232 882
ISSN: 2975-2302
ISBN: 979-12-5477-702-2
ISBN online: 979-12-5477-703-9
DOI: 10.30682/annalesps2503

www.buponline.com
info@buponline.com

Cover: Pellegrino Tibaldi, *Odysseus and Ino-Leocothea*, 1550-1551,
detail (Bologna, Academy of Sciences)

Layout: Gianluca Bollina-DoppioClickArt (Bologna)

First edition: December 2025

Table of contents

The fear of AI: A simple story of complexity

Marco Roccetti

Department of Computer Science and Engineering, Alma Mater Studiorum - University of Bologna, Resident Corresponding Academician

Abstract

This brief essay reflects on the fears that humans express toward the latest generation of AI, arguing that these fears plausibly originate from the fact that current AI possesses many characteristics that assimilate it to complex systems, as defined in modern systems theory. The features of unpredictability, emergent behavior, and self-organization, that typically characterize complex systems, do not reassure humans who usually expect predictability and determinism from the systems they are entrusted with. The essay suggests that part of the complex nature of modern AI lies in the very algorithm used for learning from data, which, in the process of training these systems, highlights and amplifies the aforementioned characteristics. The essay concludes by offering a final reflection on the fear that stems from the uncertainty inherent in science as a whole, not just from the latest generation of AI.

Keywords

Symbolic and sub-symbolic AI, Concerns about AI, Complexity, Systems theory, Gradient descent, Uncertainty.

1. Introduction

As I begin to write this brief essay, I am struck by all the possible questions that my colleagues and I may face during our meeting at the Bologna Academy of Sciences, which is specifically dedicated to modern AI and its impact on society. In the end, it does not take much imagination because the questions, though many and phrased in different ways, all point to one major issue. This can be summed up by the question of whether modern AI will eventually replace humans, and if so, in which activities first, and, most importantly, with what implications and impact on humans and human society [1].

To begin answering from the end, I will state right away that I believe that a profound sense of fear is behind all these questions. I also argue that this fear is rational, as I will try to show in the rest of this essay. To get straight to the point, my conjecture is that this is a specific feeling of fear that humanity reserves for anything that we are used to calling complex [2]. I will return later to what I mean by this term, which, by the way, can be easily ascribed to the definition given in the context of systems theory, yet my conviction is that such a feeling is not irrational, as I have said before, because it is not so much directed at the system itself, as it is at its consequences and the presumed inability of humans to fully control and, above all, understand it.

I will proceed within this Introduction by immediately highlighting why any system, technically definable as complex, can generate fear. The reasons are essentially the following:

a) *Will we lose control?* Complex systems are characterized by multiple interconnected components that interact in non-linear and often unpredictable ways. This interconnectedness makes it nearly impossible to predict the outcome of an action with certainty. Consider the well-known issue of global climate: a small change in one region can have cascading effects almost worldwide. The feeling of losing control over what is happening is one of the deepest sources of fear.

b) *Are they predictable or not?* Complexity intrinsically leads to uncertainty. We cannot be sure of the consequences that will stem from events triggered within a context of complexity. This is true both at an individual level, for example when choosing a career in a rapidly evolving job market, and at a global level, such as in the management of pandemics or economic crises. Unpredictability makes planning difficult and can make us feel vulnerable.

c) *Can cascading consequences be triggered?* In complex systems, a small problem in one component can trigger a series of chain failures that lead to the collapse of the entire system. The crash of a financial market or a large-scale power blackout are striking examples. This kind of hidden fragility is particularly frightening because disaster can seem to come *out of nowhere*.

d) *How much do we understand about how they work?* Complexity can exceed our cognitive abilities. It is difficult for the human mind to process an enormous amount of data and interconnections simultaneously. We often rely on simplified mental models to understand reality, but these models fail to capture the true nature of complex systems, leading us to make wrong choices. The feeling of not understanding what is happening can generate frustration and anxiety.

In addition, the ethical dilemmas they pose should also be considered, due to their ambivalent nature (for example, who is responsible in a context where the interactions between intersubjective entities are so intricate?). The answers to the questions above explain why humans fear *complexity* and the systems it defines. It is a reaction to the fact that these systems challenge our perception of control, predictability, and understanding of the world. Nevertheless, and I will dedicate the last part of the essay to this, it is also true that humanity has always sought to tame complexity through science, technology, and new forms of organization, demonstrating that, alongside fear, there is also a strong desire for understanding and adaptation.

In closing this Introduction, I preview here what I will cover in the next sections. In the next one, I will give a concise definition of what complex systems are. In the third section I will argue why AI, especially this latest generation, also known as sub-symbolic AI, can rightly be classified as such, certainly much more so than the symbolic AI that preceded it. In the fourth section, I will try to illustrate the origin of this complexity. Unlike explanations provided in other articles by other colleagues, explanations that are certainly more learned than the present one, I will show how this origin is partly to be found in the algorithmic domain. Particularly, in the algorithm with which sub-symbolic AI systems that learn from data learn from the data itself. Not surprisingly, it has, relatively ancient roots in the method which was developed and published in 1847 by the French mathematician Augustin-Louis Cauchy, in an attempt to solve the problem of determining the orbit of a celestial body based on its equations of motion [3]. Finally, I will conclude by stating my doubt that everyone who fears this latest version of AI, even for valid reasons, has ever truly considered the real context (both scientific and otherwise) in which AI operates. I will argue that the implications of their fear could change based on the answer to that question.

2. Defining a complex system

The theory of complexity, or the theory of complex systems, is a branch of modern science that has established itself in recent decades. It was driven by the rise of computerization and a growing inclination to abandon the assumptions of linearity in dynamic systems to investigate their real behavior more deeply. The concept of complexity has its roots in the late 19th-century works of the physicist-mathematician Henri Poincaré and in the contributions of many mathematicians and physicists, such as Kolmogorov and Schrödinger, just to cite a few, from the first half of the 20th century [4, 5].

Decisively, the advent of computers came alongside these contributions, leading Edward Lorenz to confirm Poincaré's intuition. He did this by experimentally demonstrating how significant (finite) variations in a dynamic system can arise from infinitesimal variations in initial conditions [6,7]. Leaving aside for reasons of space the numerous, highly important contributions of mathematicians and physicists to the structuring of the concept of complexity, one cannot, however, forget to mention how, on the strength of these scientific results, the philosopher Edgar Morin proposed a monumental rationalization and generalization of the concept of *complexity*. Essentially, he moved it from systems to the very notion of thought, complex thought. While acknowledging its relevance, I find it hard not to notice, however, how due to narrative suggestions, the use of the term *complexity*, once it left the scientific context, has become *unstable*, *questionable* and very often employed in a way that is semantically far from the definition that generated it.

It is now worthwhile to approach the definition of a complex system, and to do so, one must start from the almost opposite concept of *linearity*, always in the sense attributed to these concepts by systems theory. In general, a problem is said to be linear if it can be broken down into a set of independent sub-problems. When, on the other hand, the various components of a problem interact with each other in such a way that it becomes impossible to separate them to solve the problem sub-problem by sub-problem, then we speak of non-linearity.

More specifically, a system is linear if it responds in a way that is directly proportional to the stimuli it receives. It is then said that for that system, the principle of superposition holds, in the sense that if to a first stimulus the system gives a certain response, and to a second stimulus it gives another response, then stimulating the system with both stimuli simultaneously will lead it to respond with a linear combination of the two previous single responses.

Although the problems that arise in nature are often non-linear, humanity has become accustomed to resorting to the hypothesis of linearity to achieve simplified solutions, even when the system is not linear. This leads to describing the problem mathematically with models that treat it as if it were linear, considering the effects of non-linearity to be negligible, in a first approximation.

Mathematically speaking, the coefficients of the corresponding polynomial function are treated as if they were independent of each other (or so weakly dependent that their interactions can be ignored).

Pretending linearity where there is none is useful. It provides simple, easily calculable solutions that are in turn easily extendable to other similar problems, and this fiction has undoubtedly led to enormous scientific and technological progress for humanity. The fact, however, is that our claim to linearize every phenomenon (natural or not) can often prove simplistic and therefore inadequate to describe that phenomenon and thus to adequately solve the problems it entails.

When we stop pretending that systems are linear and study them in their complexity, we often find ourselves dealing with systems composed of various interconnected components or sub-systems that can interact with each other (including through feedback loops). Therefore, it is no longer possible to analytically solve all the components and their interactions, and we need *holistic* approaches capable of computing the behaviors of the individual sub-systems together with their mutual interactions, *in toto*. The drawback of this is that it becomes necessary to rely on simulations or heuristics, which lead to partial or approximate solutions.

But this in itself would not be a serious problem. Even more crucial of complex systems is their intrinsic tendency towards so-called *emergent behaviors* and related *self-organization*. This means a complex system can produce a behavior that is unpredictable and cannot be derived from the simple sum of its component elements (for example, consider the unpredictably volatile trends of the financial markets). Just as unpredictable behaviors can emerge from the non-linear interactions between a system's components, these systems can also exhibit an unpredictable capacity for self-re-organization among their many non-linearly interacting parts. This, in turn, gives rise to global behaviors that cannot be explained by a single physical law (unlike in linear systems). Famous examples of these include communities of interacting people, vehicular traffic, and computer networks.

In closing, we are dealing with unstable aggregates of agents and connections that self-organize to ensure their adaptation and survival. Usually, these attitudes emerge coherently over time, adapting and organizing the system as a whole, without a specific entity deliberately designed to manage or control it, but simply through the constant redefinition of the relationship between the system itself and the environment in which it is immersed.

Ultimately, what bothers us humans about these systems is that from the simple non-linearity of the interaction between their components and their relative self-organization, a collective tendency arises to exhibit properties that are inexplicable based on the laws governing the individual components alone. If this has already worried the reader and made them understand in advance why current AI causes concern in some, I will not burden the discussion by dwelling on other possible implications of complex systems such as so-called chaotic behaviors, the theory of bifurcations and catastrophes, and so on [2].

3. The complexity of symbolic and sub-symbolic AI

To recap from the previous section, a complex system does not adhere to the principles of homogeneity and superposition. Its distinctive properties are: a) Emergent behavior: its behavior of the system as a whole is more than the sum of its parts. For example, the behavior of a flock of birds cannot be deduced by studying the flight of a single bird. The collective organization and formations that are created are properties that emerge only from the interaction of all the birds. b) Non-linearity: small variations in inputs or initial conditions can produce enormously different and unpredictable outcomes over time (often referred, after Lorenz, to as the *butterfly effect*). c) Self-organization: the system has the capacity to structure and organize itself autonomously in response to internal and external stimuli, without the need for an external, centralized control. In short, a complex system is characterized by non-linear interactions, emergence, and self-organization, which make it unpredictable and dynamic.

We now need to argue why and under what conditions AI can be considered a complex system, a trait that becomes more pronounced with new-generation AI. Let us start, then, with the so-called symbolic AI, by first trying to provide an intuitive definition of it, one that is easily comprehensible even in the operational procedures it implements. In summary, symbolic AI has been the first major approach to artificial intelligence [8]. It is based on the idea that intelligence can be replicated by manipulating symbols and applying logical rules, in a way similar to how the human mind was believed to work. The goal was to make the computer reason, not to have it learn from data. Hence, the functioning of symbolic AI is based on two main elements. *Symbols*: these are abstract representations of real-world concepts. Things like dog, cat, barks, has-four-legs or mammal are all symbols. *Logical Rules*: these are explicit, predefined instructions that link the symbols together. These rules are usually expressed in an IF... THEN... format.

Imagine we want to build a symbolic AI that identifies a dog. As the programmer, we should enter the following rules into the corresponding knowledge base: *Rule 1*: IF the animal has-four-legs AND barks, THEN the animal is a dog; *Rule 2*: IF the animal is a dog, THEN it is a mammal.

When the system is given the input: animal that has four legs and barks, the AI examines its knowledge base, applies Rule 1, and concludes that it is a dog. It can then apply Rule 2 and make a further inference: the animal is a mammal. This simple example illustrates the notable fact that, with symbolic AI, the decision-making process is completely transparent and explainable because one can always trace back to the rule that led to that conclusion. Its power lies in its ability to make logical deductions from a set of explicitly defined symbols and rules.

Unlike the simplicity of the example, things can also get complicated, for instance, because the knowledge base can become very large, or because, instead of the classical *modus ponens*, more complex logical rules could be adopted that complicate the inferential process. But above all, this happens because symbolic AI systems are typically used through backward chaining, which means interrogating the system to judge the veracity of a fact, thus triggering decision-making processes that start from what one wants to prove and work their way back up the decisional tree until the fact is proved (or negated). Finally, the fact that probabilistic weights can be linked to rules does not fundamentally change this approach; on the contrary, it complicates the semantic interpretation. This is because probability theory has its own formal structure, which already provides for the possibility of updating beliefs based on Bayes' theorem. As a result, update mechanisms based on the formula for conditional probability are often not easily compatible with the admittedly simpler ones that stem from inference rules like *modus ponens* or *tollens* [9].

Now, I will not spend too much time on the definition of non-symbolic AI, both for reasons of space and due to the wide range of important differences between its various instances, which would lead us too far afield. For explanatory purposes, it is enough to state that sub-symbolic AI is a field of artificial intelligence that focuses on systems that learn from data and process information without relying on explicit symbolic representations or rules, unlike traditional AI. Instead, it uses numerical values and complex mathematical models to represent knowledge, an approach inspired by how the human brain processes information through interconnected networks of neurons. A key example of this is deep learning, which uses neural networks made up of many layers. These networks learn by being trained on vast amounts of data, such as images, text, or audio. During training, the network adjusts the numerical weights and biases of its connections. When an output is incorrect, the network uses a process called *backpropagation* to adjust these weights to reduce the error. Over many iterations, the network learns to recognize patterns and make accurate predictions [10].

For instance, to teach a sub-symbolic AI to identify cats, one should feed it thousands of labeled images. The network would initially make random guesses, but through backpropagation, it would gradually learn which combinations of weights and biases lead to correct identifications. Consequently, the knowledge it gains is not a set of rules like *a cat has pointy ears*, but rather a complex configuration of numerical weights that allows it to successfully identify a cat. This approach has been proven incredibly effective for tasks like image recognition, natural language processing and speech understanding, for example.

But having briefly described both, let us now get to the point we wanted to make: namely, when we talk about symbolic or sub-symbolic AI, are we talking about entities that meet the requirements of complex systems? In my opinion, the answer is yes, even though, and this is perhaps the turning point of this whole situation, it is in a much more pronounced way for sub-symbolic AI.

In particular, regarding symbolic AI, one could argue that symbolic AI is complex because it is not linear, where its complexity emerges from:

- *Vastness*: its complexity is primarily *combinatorial*. A symbolic system can have an enormous number of rules and a logical search tree with billions of branches. Its complexity lies in managing this immense network of possibilities.
- *Complex internal structure:* with symbolic AI, there is no true learning or adaptation in the biological sense of the term. The rules and symbols are explicitly entered by a human. In this sense, the system's behavior would be, therefore, in principle, predictable because all the rules are known, however in the practical use of these systems it is not always easy to trace their lines of reasoning, a fact Gary Kasparov learned firsthand way back in 1997 in New York City.

When it comes to sub-symbolic AI, things are decidedly worse (or better, depending on how one looks at it), because sub-symbolic AI is to be considered *more complex*. Its complexity is not just combinatorial, but also emergent. Its systems are not governed by explicit rules, but instead learn to create their own internal representations of the data, thus revealing the following characteristics.

- *Emergence:* the intelligence and ability of the system to recognize a cat or translate a text are not the result of a coded rule, but of emergent behaviors that arise from billions of interactions between the network's neurons during training. This is a key trait of complex systems.
- *Adaptation:* the system autonomously adapts and self-organizes. The learning process, in which the weights of the neural connections are modified to minimize an error, is a form of dynamic adaptation to new inputs, a key characteristic of complex adaptive systems.
- *Intrinsic non-linearity:* unlike linear systems, neural networks are intrinsically non-linear thanks to their activation functions and layered structure. This non-linearity is not a simple sum of rules, but a complex transformation that generates unpredictable behavior.

In summary, we can conclude that symbolic AI is complex because it is non-linear and requires managing a vast amount of logical information. However, its complexity is stable and determined by the rules that govern it. Sub-symbolic AI, on the other hand, should be considered vastly more complex because its intelligence emerges, learns, and adapts somewhat autonomously, creating a behavior that cannot be predicted *a priori* from simply knowing the code or the data it learns from.

4. To the sources of the complexity of sub-symbolic AI

In 1986, a seminal paper in *Nature*, co-authored by Geoffrey Hinton, David Rumelhart, and Ronald J. Williams, popularized the *backpropagation* algorithm [11]. This was a critical step in training multi-layer neural networks, even though the group was not the first to develop the method. Backpropagation (or error backpropagation) is really a fundamental method used to train neural networks: in simple terms, it is an algorithm that efficiently calculates the gradient, which is the direction and magnitude by which the network's parameters (the weights) must be changed to minimize the error committed by the system. Its operation is based on a clever application of the chain rule. Essentially, it unfolds in two main phases: a) A *forward* step where the system receives an input and processes it, layer by layer, until it produces an output. This output is then compared with the expected result to calculate the error. b) A *backward* step where backpropagation comes into play. Instead of performing redundant calculations, the algorithm starts from the last layer and moves backward. At each layer, it calculates how much each weight contributed to the total error and determines how to adjust it. In fact, the term backpropagation strictly refers only to this gradient calculation phase, but it is often more generally used to denote the entire learning process. Once the gradient is calculated, the network/system uses a method like *gradient descent* to update the weights in the direction that reduces the error.

Underpinning all of this is indeed this mentioned concept of *gradient descent*, a fundamental method for unconstrained mathematical optimization, generally attributed to the French mathematician Augustin-Louis Cauchy, who first suggested the approach in 1847. This is an iterative procedure designed to minimize a differentiable multivariate function. The core idea is to repeatedly take small steps in the opposite direction of the function's gradient at the current point. This works because the gradient points in the direction of the steepest increase, so moving in the opposite direction is the fastest way to descend to the function's minimum. Conversely, stepping in the direction of the gradient will lead to a maximum, a process known, in turn, as gradient ascent. This method is particularly useful in machine/deep learning for minimizing the cost or loss function.

Therefore, this method is used to train neural networks on data because, under certain conditions (*i.e.*, differentiability and convexity), it guarantees the minimization of the loss function. It is also true that in most real-world training scenarios, the loss functions turn out to be non-convex. Therefore, the iterative process of calculating the gradient and moving in the opposite direction leads to finding only local minima. This is also why much of modern deep-learning research is dedicated to developing more sophisticated optimizers that can navigate these non-convex landscapes and, hopefully, find a global minimum (or at least a very good local minimum). Ultimately, this process progressively reduces the error and enables the network to recognize what it is observing, or, more precisely, what humans are showing it with near-perfect accuracy.

In closing this section, far from considering the gradient descent algorithm itself intrinsically complex, as it is the very mechanism that allows much of sub-symbolic AI to function, it is nevertheless a key component that contributes to the overall complexity of a neural network's learning process, being the main tool that allows this system to evolve.

Here is how I argue that gradient descent contributes to this complexity, alongside a variety of other factors, like the number of layers and the trainable parameters of a given neural network:

– *Feedback*: the algorithm uses a continuous feedback loop. The network's output is compared to the expected value (the error), and this error is used to adjust the weights. This non-linear feedback loop is a hallmark of complex systems because the action (weight update) directly depends on the outcome (error), influencing future actions in unpredictable ways.

- *Dynamic behavior*: a neural network learning process is itself a dynamic system. The system's trajectory, that is the evolution of weights over time, is guided by gradient descent. This path is not predetermined; it is an iterative movement in a high-dimensional space (the weight space) that seeks to minimize the cost/loss function.
- *Emergent properties*: when applied to a neural network, gradient descent allows for the emergence of unexpected properties. The network does not learn explicit rules (like if-then-else statements) but instead develops an internal representation of reality, such as the ability to recognize faces or understand languages. This ability is not programmed; it emerges from the dynamic interactions of the nodes, guided by the gradient descent.

In summary, gradient descent is the engine that allows a neural network, a complex system, to self-organize. Its main function is to guide the system's evolution toward a state of minimum error. The complexity of the learning process lies in the fact that, even though the algorithm is simple in its logic, its cumulative effect on thousands or millions of weights leads to emergent and sometimes unpredictable behaviors.

5. Concluding remarks

The idea developed in this concluding section is not to provide the usual summary of what has already been discussed in the previous ones, but to use this space and time to propose a further point for reflection. This reflection aims to offer at least one soothing, or consoling, argument to those who fear AI, and with good reason based on what was discussed earlier. This argument is loosely based on a recently published essay that discusses the concept of *truth* and how *science* has contributed to its determination [12].

Simplifying as much as possible, it is assumed that there are (at least) two ways to handle the concept of *truth* in science, and the manner with which it is confirmed through proper verification.

The first context is that of mathematical truth, which since the time of Euclid has proposed a point of view entirely governed by the duality of axioms and theorems. The axioms are the premises, accepted as true without any discussion. Theorems, instead, are knowledge logically derivable from the axioms. From this perspective, the solutions that humanity has found for real-world problems with mathematical thinking are simply constructive procedures that advance infallibly, but only by starting from premises we have held to be true, along with the rules that have helped us in the construction. But what if such premises were not true? And what if the very rules of inference we have relied on were to prove, at some point, to be inadequate in capturing the reality that surrounds us? For example, who says that if it rains, our heads will always be wet, even when we wear a waterproof hat or open an umbrella?

Let us, now, move on to the second context, the one, all in all, more recent, in which humanity has become accustomed to deriving that kind of knowledge on the basis of which modern society developed. This is the typical context of natural and physical sciences, those that require the sophisticated step of experimentation to validate a theory (*i.e.*, a true fact). In this second context, all the truths we arrive at are, in themselves, contingent and empirical. Each one is demonstrated, yet always with a margin of doubt, and thus of risk, which, however, we consider bearable. Their coherence with the real world is always dependent on the amount of evidence (or proofs) we have managed to gather during the experimentation phase. Unlike the mathematical context, here there is not even a need to rhetorically advance the *what if* doubt, because these truths are precisely built with *doubts inside*. This is so indisputable that it often happens that in extreme crisis situations we tend to forget the limits, and at the same time the true strength, of truths acquired through the technique of validation by experimentation. If the reader is not convinced by what has just been said, a worthwhile effort would be to think back to the recent COVID-19 pandemic and to consider both hydroxychloroquine and mRNA vaccines [12]. One should try to ask themselves if they remember which of the two methods emerged as scientifically valid in the sense mentioned before. But above all, one might try to ask if, in those moments or even afterward, personal beliefs aligned with the messages that science was providing, and whether those messages were found to be comforting or, instead, misleading.

In closing, this entire final digression would like to present the reader with the following, hopefully comforting, consideration. What we call science (and we have actually only been calling it that since a relatively short time) has generally brought great advantages to humanity, but it has always had to confront uncertainty, unpredictability, the emergence of unexpected phenomena, and problems that have emerged as side effects of solutions to other problems. This has been true whether it occurred under the apparently infallible guise of logical-mathematical reasoning or in the form of an incessant and continuous experiential effort to find evidence for or against a given scientific theory.

So my final question to readers is why worry, beyond the limits of what has been explained to be rational, about the artificial (intelligence)? After all, the artificial is human, just as the human has so much of the artificial within it. Uncertainty and unpredictability are both outside of us, and also inside of us. To fear artificial intelligence is perhaps just one of the many ways we feel fear for human intelligence, and also a great deal for human unintelligence.

References

1. Eloundou, T.; Manning, S.; Mishkin, P.; Rock, D. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 2024, 384(6702):1306-1308. doi: 10.1126/science.adj0998.
2. Mitchell, M. *Complexity: A Guided Tour*. Oxford University Press: Oxford (UK), 2009. doi: 10.1093/oso/9780195124415.001.0001.
3. Cauchy, A.L. Méthode générale pour la résolution des systèmes d'équations simultanées. *Comptes Rendus de l'Académie des Sciences*, 1847, 25:536-538.
4. Ladyman, J.; Lambert, J.; Wiesner, K. What is a complex system?. *European Journal for Philosophy of Science*, 2013, 3:33-67. doi: 10.1007/s13194-012-0056-8.
5. Skrimizea, E.; Haniotou, H.; Parra, C. On the complexity turn in planning: An adaptive rationale to navigate spaces and times of uncertainty. *Planning Theory*, 2019, 18:122-142. doi: 10.1177/1473095218780515.
6. Lorenz, E.N. Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*, 1963, 20(2):130-141. doi: 10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.

7. Poincaré, H. Sur le problème des trois corps et les équations de la dynamique. *Acta Matematica*, 1890, 13(1-2):VII-XII. doi: 10.1007/BF02392505.
8. Russel, S.J.; Norvig, P. *Artificial Intelligence: A Modern Approach*. Pearson Education: Hoboken NJ (USA), 2020. ISBN: 9780134671932.
9. Roccetti, M.; Teolis, A.G.B. Sulla correttezza probabilistica dei sistemi esperti. *Statistica*, 1991, 51:267-299. doi: 10.6092/issn.1973-2201/871.
10. Haykin, S. *Neural Networks: A Comprehensive Foundation*. Prentice Hall: Hoboken, NJ (USA), 1999. ISBN: 978-0132733502.
11. Rumelhart D.E.; Hinton G.E.; Williams R.J. Learning representations by back-propagating errors. *Nature*, 1986, 323:533-536. doi: 10.1038/323533a0.
12. Li Calzi M. Chi è Custode della Verità?. *Il Foglio*, 2025. 189, 3, anno XXX.