



Intelligence at any price? A criterion for defining AI

Mihai Nadin¹

Received: 3 April 2023 / Accepted: 9 May 2023

© The Author(s), under exclusive licence to Springer-Verlag London Ltd., part of Springer Nature 2023

Abstract

According to how AI has defined itself from its beginning, thinking in non-living matter, i.e., without life, is possible. The premise of symbolic AI is that operating on representations of reality machines can understand it. When this assumption did not work as expected, the mathematical model of the neuron became the engine of artificial “brains.” Connectionism followed. Currently, in the context of Machine Learning success, attempts are made at integrating the symbolic and connectionist paths. There is hope that Artificial General Intelligence (AGI) performance can be achieved. As encouraging as neuro-symbolic AI seems to be, it remains unclear whether AGI is actually a moving target as long as AI itself remains ambiguously defined. This paper makes the argument that the intelligence of machines, expressed in their performance, reflects how adequate the means used for achieving it are. Therefore, energy use and the amount of data necessary qualify as a good metric for comparing natural and artificial performance.

Keywords Data · Energy · Matter · Living matter · Anticipation · Adaptivity · Neuro-symbolic AI

1 Introduction

The timeline of developments associated with AI is also a record of high hopes and failures. In retrospect, the decoding of Enigma (1942)—the cipher device used by Nazi Germany in WWII—and Turing’s test for machine intelligence (1950), stole from McCarthy et al. (2006), the glory of giving birth to a discipline, but not of being the first in naming it. The name is more a curse than a blessing given the consequences of the labeling adopted since the Dartmouth Conference of 1956. What follows on the timeline is the failure of machine translation (1966)—currently no longer a problem—and even, as surprising as it might sound, of connectionism (1970).

That the symbolic beginnings of AI where disappointing does not need to be rehashed here. The same holds for the opaque nature of artificial neural networks (ANN) performance—impressive, but still not easily explainable. It is still unclear, more than 50 years later, whether the two can be related, never mind unified in a coherent manner. They have

in common the qualifier intelligence, more a target than the outcome of a clearly defined goal.

The current state and the future of AI—symbolic, neural, or combined—cannot be meaningfully assessed in the absence of adequate evaluation means and methods. Adequacy itself is reflective of the understanding of the subject of evaluation. It is expressed in the specific metric, the benchmark, to be applied when evaluation of performance is undertaken. The evaluation ought to be agnostic of how the goal of conceiving, designing, and building machines labeled intelligent is reached. In other words, whether computer-based, or of any other nature (there is future even after the Turing machine, Turing 1948), the machine’s purpose cannot be confounded with how it might be achieved, or even with its output. The moving of goalposts (the “AI effect,” i.e., what was once labeled AI and became routine data processing (McCorduck 2004, and later the philosopher Nick Bostrom, in many lectures) will not help in defining how AI is different, or not, from the broad understanding of algorithmic computation. This is made even more critical as neuro-symbolic directions are pursued (and the optimism reaches almost euphoria, Hitzler and Sarker 2022).

Awareness of how difficult the task is (see Chollet 2019; Hernandez-Orallo 2017; Legg and Hutter 2007; Howe and Cohen 1988, among others) is indicative of the acknowledged need to go beyond the “inaugural banner” of the

✉ Mihai Nadin
nadin@utdallas.edu

¹ Institute for Research in Anticipatory Systems, University of Texas, Dallas, USA

Dartmouth Conference of 1956. All known attempts at a metric of progress in AI have so far failed. The Turing Test,¹ which Chollet (2019, *op. cit.*, p.3) described as a way to “out-source the task” of distinguishing machine intelligence from human intelligence to “unreliable human judges,” who themselves do not know better, together with its “relatives”—e.g., the Total Turing Test, the Loebner Prize—is but the better known example of such failures. The reason for this unavoidable outcome—i.e., its inadequacy—is obvious: the curse of circularity. Minsky’s definition—“AI is the science of making machines capable of performing tasks that would require intelligence if performed by humans” (Minsky 1968)—points to the “center” of the circle in which AI, as a particular form of computer science, keeps moving. Extending the radius (from expert systems, i.e. symbolic AI to deep learning, i.e., neural AI, to whatever else) does not break the circularity of the enterprise: intelligence is whatever is considered as the outcome of intelligent action. The novelty of ML is that input is now understood as training, and output as operations (inferences, in the first place) performed on request (interactively) on training data (the more, the merrier the outcome). But there is no knowledge to account for. Moreover, there is no way to quantify in advance, the cost of the performance.

This paper (part of a series of elaborations on the subject, Nadin 2013, 2015, 2017, 2019, 2020, 2022a, b, c) starts by placing the subject in its proper context: How does the living, regardless of scale, achieve its intelligent performance? Which means: How does the living understand reality in order to deal with it for survival or for some other purpose? Based on this question, a first attempt at defining artificial intelligence, and what it takes to achieve it, is expressed in the form of an *evaluation principle*. It pertains to data and energy—parameters which can be measured. Therefore, applications of the principle are testable.

2 Minimum energy principle (MEP)

Entities embodied in lifeless matter (such as robots, but even cellular phones, or adaptive materials, for example) succeed more and more in emulating activities associated with organisms embodied in living matter. This is the domain of the artificial. AI, in particular symbolic AI, and more recently Machine Learning (ML), as a particular form of AI are part of this domain. The discussion of the Singularity (Vinge

1993; Kurzweil 2005), the presumed state in which the artificial, in its ever-increasing variety, outperforms the living, was reignited by recent accomplishments. Large Language Model (LLM) developments, using the Transformer Architecture, integrate algorithms for mimicking understanding of language, and for generating similes of human language. To quote:

GPT-4 is a large multimodal model (accepting image and text inputs, emitting text outputs) that, while less capable than humans in many real-world scenarios, exhibits human-level performance on various professional and academic benchmarks (GPT-4, openai.com).

But before GPT-4 there was the OpenAI and high target: Our mission is to ensure that artificial general intelligence—AI systems that are generally smarter than humans—benefits all of humanity. For this purpose, large internet datasets were used to train various kinds of ANN. GPT-3, as a precursor, was aimed at natural language answering of questions, but also to translate between languages and coherently generate improvised text. For images, DALL-E, a deep learning model that can generate digital images from natural language descriptions was developed. ChatGPT, a newly conceived chatbot based on GPT-3.5, stole the show, but it is by far not unique (Google, Microsoft, MetaAI, etc. have their own developments), neither in performance nor in the basic principle applied and represented by Transformers of all kind.

The context is clear: computation in non-living matter and language performance characteristic of the human (embodied in living matter) are compared to each other. After all, AGI is the goal, i.e. the general intelligence of humans, and not the specialized performance (playing chess, Go, etc., Hassabis et al. 2017, Silver et al. 2018) displayed not long ago by IBM’s Watson. It is therefore justified to see how the distinction between matter and living matter can help in understanding what this all means.

In living matter across scales—from cells to organisms, to species—activities for the preservation of life cannot consume more energy than what metabolism affords. This is the Minimum Energy Principle (MEP). At all scales, from mono-cells to complex organisms, intelligence guides choices that translate into survival. It is an evolving intelligence, reflecting adaptivity to new conditions of life. Purpose is what drives choices.

Evolution explains why the MEP does not hold for the human being. From the entire realm of the living, only the species *homo sapiens*—*thinking* being the ultimate identifier—consumes more energy in its self-preservation than what metabolism alone contributes. What became known as *culture*—i.e., the tamed nature within which human activity takes place—is the outcome of progressively increasing

¹ Diderot: s’il se trouvait un perroquet qui répondit à tout, je prononcerais sans balancer que c’est un être pensant (cf. *Pensées Philosophique*), translated as..if there was a parrot which could answer every question, I should say at once that it was a thinking being (cf. *Philosophic thoughts*, page 37).

energy use, and thus the continuous remaking of oneself as energy dependent. No other form of life on Earth has this behavioral pattern. The human being redefined itself in respect to physical abilities—augmented by tools and machines—and to thinking—cognitive abilities, associated with a larger brain. The augmented capabilities are energy and data dependent. Since intelligence-driven evolution of humanity reached the level at which resource consumption became an issue of its survival, to ignore the consequences of this situation means to ignore the perspective of sustainability.

Just for the sake of grounding this observation in measured data: ChatGPT got one billion inquiries in its first month after release (February 2023). It generated ca. one trillion tokens (i.e., answers to inquiries ranging from the frivolous to serious tests). Experts translate this into seven metric tons of CO₂ per day—very large carbon footprint. Inference processes triggered by inquiries reached a carbon footprint level higher than those of training, using over one trillion words. A glut of spectacular tests corralled the public into either unrestrained optimism—the dream of Artificial General Intelligence (God-like intelligence)—or despair—we will be replaced, as doctors were warned, when Watson (by now gone) was making the rounds.

Endowing non-living matter with capabilities that can be associated with natural intelligence—the cutting edge of science today—is an energy-and data-hungry endeavor. There is the Landauer Principle, a theory concerning the lower limit of energy consumption for computation (Landauer 1961; and there is the estimation of energy consumption in machine learning (Garcia-Martin et al. 2019). When it is practiced unintelligently, i.e. as brute force, it does not justify the outcome. Yann LeCun, who contributed to convolutional neural networks (CNN), observed (in a Tweet, April 2023, @ylecun) that “Human don’t need to learn from 1 trillion words to reach intelligence.”

For the record: outsourcing natural functions to artifacts starts with the use of tools. The lever increased the force applied to an object (e.g., moving a heavy stone). The tool itself is not intelligent, but it allowed an individual to perform something (move the stone) for which in the absence of the lever, several persons were necessary. It could be used as needed. *Homo habilis* is the human being making itself in the process of conceiving and using tools. They are in anticipation of their use, i.e., a way to multiply future possibilities. Tools date back to the first identifiable human forms of activity—foraging and hunting. From their hard condition—matter (e.g., stones) made into artifacts—to their soft condition—programs to activate various machines for more work—they represent knowledge put into action. The immediate result of this pattern—from hardware to software—is the disconnect between means of existence—ecological sources of energy—and progressively reduced

natural expression. From the natural cycle of day and night, of seasons, to artificial cycles of machines, the change is such that artificial time replaces the naturalness of the human being. The declining anticipatory function—appropriate to natural interactions—corresponds to the insertion between goals and achievements of intermediaries. The lever and the hammer are examples, as is the shovel; so are gears and pulleys. Such tools do not understand the WHY of the actions in which they are used. The Generative Pre-trained transformer (GPT) and the ChatGPT—to name the newest impressive “hammer,” based on an autoregressive language model, that everyone wants to play with—is the new tool to be followed by more and more. This tool was trained on 570 GB of data acquired from books, Webtexts, Wikipedia, articles and other pieces of writing on the internet. And it generates texts or images that imitate human written texts or drawings. But it does not have any understanding of what it generates.

3 Intelligence means understanding

First, an answer from the ChatGPT system:

As an AI language4 model, I don’t possess any comprehension of understanding of language at all. My responses are generated purely based on statistical patterns and associations learned from training data without any subjective experiences or understanding of the world. Therefore, I cannot make any claim about appropriateness of accuracy of my responses.

There is no sense of future, and accordingly, no anticipatory process to be expected.

With the focus on understanding how change takes place—including their own changes over time—humans effectively substituted their innate anticipatory abilities with artificially constructed models of the future inspired by the past. The minimum energy threshold characteristic of survival was effectively overwritten by the optimistic principle of Everything Is Possible (EIS)—at the expense of the ecological system. The human species lives at the expense of the rest of the environment of its own existence. And it is the only species *devolving* into overpopulation. No other living being could afford activities in which the outcome is less than the effort. The flipping of the Upside/Downside Ratio, i.e., the negative yield of human activities dependent more and more on energy use, is characteristic of a new stage in the life of societies. It documents the assertion, as impressive as the conversation on AI and Machine learning, that human beings live more and more at the expense of the future. Not understanding the purpose of action, i.e., lacking intelligence, has existential consequences.

Although the MEP does not hold for humans, it is justified to define intelligence in the perspective of the MEP, in conjunction with its particular expression as the Data Minimum Principle (DMP).² Taking in reality through the senses consumes energy. Given the behavior conditioned by the MEP, it follows that the living, through perception, measures reality to the minimum possible. This means that data pertinent to life interactions cannot be less than the minimum it takes for the preservation of life.

Given the obsession with higher performance in imitating life, it is justified to evaluate the performance of artificial means in order to assess their viability. Based on this understanding, an evaluation principle can be formulated:

Artificial entities could justifiably claim intelligence if, in executing a task, they would use as much energy or less, and as much data or less, than a living entity performing the same task.

Energy use and the necessary data are a good metric for comparing natural and artificial performance: how much energy and how much data is used in a well-defined activity. Some machine learning applications (chess, or *Go* playing, Dickson 2020, Labbe 2021) can use as much electricity as a small town over the duration of the performance. GPT-related energy consumption (the cost of training on large scale data), and the cost of use, were scrutinized already (and caused justified warnings).

The energy consumed by the miniscule bar-tailed godwit during migration is acquired through metabolism. The data processed is acquired through “measuring,” i.e., sensing the environment. (Data acquisition also involves energy expenditure.) The take-off, ascent, gliding, bonding, soaring, and continuous forward flight through flapping wings are energetically different. Some actions are more “expensive” than others. Altitude is yet another factor: less oxygen, for example, addressed by a different motoric for consuming less of it (Alexander 1998). For the sake of the discussion, it suffices to mention that a power of 4.3 watts is actually used for the flight (Wikelski et al. 2003). They are by some orders of magnitude less than what would be needed to guide an artificial bird of similar size and weight. Winning or losing a game of chess or of *Go* would not require a power plant if performed by a human being. The data processed by humans in playing the games is in the order of kilobytes. This is way smaller than the huge data amounts (order of 10^{120}) guiding the artificial playing machine. The human brain operates on 20–30 watts—less than an LED source. Even the plankton inhabiting the oceans is much more intelligent than what the

most sophisticated machinery, based on the deterministic science dominating civilization, can achieve. This pertains to the energy used and the data collected. It is expressed in its adaptive performance. Indeed, the dynamics of the plankton is non-deterministic. The plankton navigates the oceans under terrible conditions, finding survival niches for which we do not have names, never mind knowledge about. In anticipation of adverse conditions, swarms of migrating birds or of fish change, respectively, flight altitude or swimming depth. Even under the most generous assumptions of scientific and technological progress, performance comparable to that of living entities in a continuous state of anticipation is not even on the agenda of current science and technology. Such a performance is as impossible as doubling a cube using a compass and ruler or squaring the circle.

What makes the difference is the anticipatory component of the activity. Migratory behavior exhibits adaptive characteristics associated with anticipatory processes driven by the possible future. The timeline (migration start) and the trajectory are fine-tuned to possible storms, as though the migrating birds, or migrating fish or animals, are prescient of what might affect—possible future—their respective journeys (on the predictive performance of veeries, see McGlashan 2019, one reference from among many). Artificial entities embodied in non-living matter are “fired up” with energy from the outside and with data from measurements of similar activities. To be precise: the living senses the environment. Sensing involves energy use: to measure is an activity engaging the entire living being. When the available energy acquired through metabolism or stored is too low, the living ceases to take in “reality.”

The living is in a state of anticipation from the start of life until its end. It is a continuous state, with various forms of expression and variable intensity. It engages the entirety of the organism, at all its levels. It depends on metabolism and on perceptual activity. Avoiding danger, as opposed to reacting to it, is, from the perspective of the data involved and the energy consumed, quite different. Outperforming others in the context of the competitive nature of life and in securing evolutionary advantage takes place also on account of energy use and data processing appropriate to the circumstances.

Inventions qualify as examples of activities driven by anticipation. Anticipatory processes effectively extend awareness of cause-and-effect into the richer sense of causality that integrates past, present, and possible future. “Sensing” the future, i.e., virtually living it before it becomes real means awareness of consequences. From an energy and data perspective, this is different from the practice of predicting it on account of measuring reality and inferring from a current state to a future state. It succeeds (or fails) if the energy expense undermines life. That is, if the data goes beyond what a specific living entity can afford to acquire,

² Occam (thought that it is vain to do with more than can be done with less (Frustra fit per plura quod potest fieri per pauciora). *Summa Totius Logicae*, Loux 1974).

there is no future state to account for. The minimum energy for the human being is not predicated by the threshold of life, i.e., what is needed to maintain life, but rather by gaining independence from environmental limitations. The human being can extract from the environment more energy than is needed to survive and multiply. Humans also acquire more data than what would be needed to maintain life. This explains their evolutionary success, but also raises questions. Is living at the expense of the future (in terms of exhausting resources, or in affecting the viability of the species) a viable option?

A metric for artificial intelligence is not only a means for comparing natural and artificial performance. Open AI benefits from the public success of ChatGPT or GPT-4, but the real profit goes to the providers of brute-force computation tools (in particular, Nvidia's GPU Technology data guzzlers). It could help in the broader evaluation of what it takes to achieve sustainability to understand the cost of converting brute force into the appearance of intelligence. On the current course of AI and ML spectacular achievements, i.e. obtained in the absence of intelligence substituted by brute force computation, sustainability cannot be hoped for.

Data availability statement Not applicable.

References

- Alexander McNR (1998) When is migration worthwhile for animals that walk, swim or fly? *J Avian Biol* 29(4):387–394
- Chollet F (2019) On the measure of intelligence, arXiv, 5 November, [arXiv:1911.01547](https://arxiv.org/abs/1911.01547) [cs.AI]
- Dickson B (2020) Energy consumption of AI poses environmental problems. <https://www.pcmag.com/news/ai-could-save-the-world-if-it-doesnt-ruin-the-environment-first>
- Garcia-Martin E et al (2019) Estimation of energy consumption in machine learning. *J Parallel Distributed Comput* 134:75–88
- Hassabis D, Kumaran D, Summerfield C, Botvinik M (2017) Neuroscience-inspired artificial intelligence. *Neuron* 95(2):245–258
- Hernandez-Orallo J (2017) Evaluation in artificial intelligence: from task-oriented to ability-oriented measurement. *Artif Intell Rev* 48(3):397–447
- Hitzler P, Sarker MdK (eds) *Neuro-Symbolic Artificial Intelligence: The State of the Art*, Vol. 342, *Frontiers in Artificial Intelligence and Applications*
- Howe AE, Cohen PR (1988) How evaluation guides ai research: the message still counts more than the medium. *AI Mag* 9:35
- Kurzweil R (2005) *The Singularity Is Near: When Humans Transcend Biology*, Viking
- Landauer (1961) Irreversibility and heat generation in the computing process. *IBM J Res Dev* 5(3):183–191
- Legg S, Hutter M (2007) Universal intelligence: a definition of machine intelligence. *Mind Mach* 17(4):391–444
- Labbe M (2021) Energy consumption of AI poses environmental problems, *Tech Target Enterprise AI*
- Loux (Trans) MJ (1974) Ockham's theory of terms (Part 1 of the *Summa Logicae*). Notre Dame University Press, Notre Dame IN
- McCarthy J et al (2006) Proposal for the dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Mag* 27:4
- McCorduck P (2004) *Machines who think*. A.K. Peters, New York
- McGlashan A (2019) Are these birds better than computers at predicting hurricane seasons? *Audubon Magazine*, <https://www.audubon.org/news/are-these-birds-better-computers-predicting-hurricane-seasons>. Accessed 28 Oct 2021
- Minsky ML (ed) (1968) *Semantic information processing*. MIT Press, Cambridge MA
- Nadin M (2013) The intractable and the undecidable—computation and anticipatory processes. *Int J Appl Res Inf Technol Comput* 4(3):99–121
- Nadin M (2015) Anticipation and computation. Is anticipatory computing possible? In: Nadin M (ed) *Anticipation across disciplines, cognitive science monographs*, 29. Springer, Cham, pp 283–339
- Nadin M (2017) Predictive and anticipatory computation. In: Laplante P (ed) *Encyclopedia of computer science and technology*, 2nd edn. Taylor and Francis/CRC Press, London, pp 645–659
- Nadin M (2019) Machine intelligence: a chimera. *AI Soc* 34:215–242
- Nadin M (2020) Aiming AI at a moving target: health (or disease). *AI Soc* 35:841–849
- Nadin M (2022a) *Cumpăna Științei*. Spandugino, Bucharest
- Nadin M (2022b) *Anticipation informed design*. eBook. DAE Publishers, Eindhoven
- Nadin M (2022c) *Anticipation and epigenetics*. Springer, Cham
- Silver D et al (2018) A general reinforcement learning algorithm that masters chess, shogi and Go through self-play. *Science* 362(6419):1140–1144
- Turing A (1948) *Intelligent machinery*. National Physical Laboratory Report, London
- Vinge V (1993) *The coming technological singularity*, *Whole Earth Review*, Point. p 88–95
- Wikelski M et al (2003) Costs of migration in free-flying songbirds. *Nature* 423:704

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.