AI for Controlling Our World Beyond Humanly Accessible Complexity

Gheorghe M. Ștefan

(draft)

Abstract

In the human world, two entities coexist with different dynamics of evolution and adaptation. The *living* and the *technological*. The living dominated by evolutions with a *logarithmic* dynamic and the technological dominated by an *exponential* dynamic. ***Humanity*** and ***Artificial Intelligence*** are representative of the two entities. Their relationship is increasingly tested by the complexity it generates in the human world, which comes with an additional dimension, that of the population that also grows exponentially. Can we hope, under these conditions, for a saving evolution that allows us to control an exponential complexity using our capacities that evolve logarithmically? Perhaps a well-tempered control of this complexity through human-AI coevolution can offer a solution. We believe that this prefigures a process of deeper understanding of what the human being is with all the limitations that we have not yet been able to accept.

The *noosphere* (Teilhard de Chardin [7, 8], Vladimir Ivanovich Vernadsky [49]) and the *technosphere* (Tudor Vianu) manifest themselves with maximum intensity in a product that they support convergently: Artificial Intelligence (AI). It is for the first time in the evolution of the human world in which the two spheres emanating from the *biosphere* collaborate in the emergence of an entity with characteristics borrowed from both spheres. Indeed, AI is not only a manifestation of the noetic nor only a technology. Through the uniqueness of the approach, will AI

be able to lead us beyond what humanity is today? Just as technological evolution has had a major impact on human behavior throughout history, we expect major transformations corresponding to the level of novelty and efficiency of AI [26, 38].

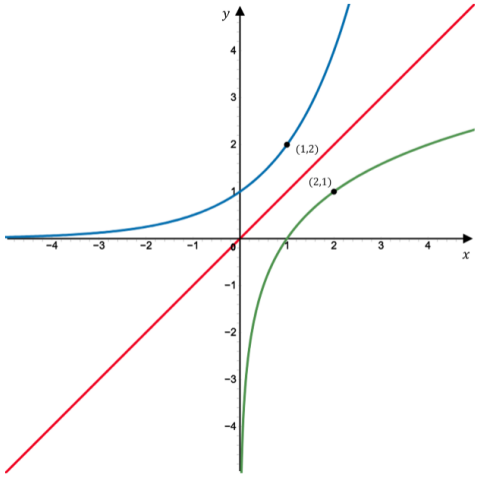


Figure 1: The shape of an *exponential* function (blue) and a *logarithmic* function (green) compared to the shape of a *linear* function (red).

The science that reveals and the technology that imposes both evolve under the impetus of the exponential growth of technological capabilities in the context of another exponential; the exponential growth of population. The human capacity for research and innovation, which evolves logarithmically (see Figure 1), adapts with difficulty to exponential developments due to a natural inertia, imposed biologically, culturally or institutionally. Likewise, the capacity of the population to assimilate and use new knowledge and technological products is limited by a similar inertia. Research funding policies and the appetite of investors in new technologies are fundamentally affected by the distorted way in which science and technology are perceived by decision-makers and the public. The incompetence of politicians and the greed of investors, in the context of the increasing complexity of the world dominated by more subtle self-organizing mechanisms than ever, require a radical solution, easy to define but very difficult to implement. It is

about ***education***, as part of the learning triad that involves *instruction - training - education*. Education, as the ability to decide where neither reason nor imagination succeeds, is the most ignored component of learning. Therefore, we will have to review our educational strategies starting with the fairy tales we tell preschool children.

Therefore, man must find how to cope with the complexity, brought by the two exponentials in the human world, with his logarithmically limited capacity to evolve.

1 Under the imperium of two exponentials

The human world, which emerged under the pressure of survival intelligently controlled by man, is characterized by processes of evolution and adaptation that lie between two extremes. At one extreme we have processes with a logarithmic dynamic (this concerns the behaviors of the living embodied in human life, but also in the structures of human communities), and at the other extreme processes with an exponential dynamic (this concerns technological developments). Linear developments are rare, episodic or masking exponential ones in the initial phases of growth.

In Figure 1 the three types of evolution are represented. In the human world, extreme types co-exist in a tension that, until now, has stimulated a historical evolution along the path of progress that almost all of us recognize.

No one guarantees that the current evolution has or will have an upward course from the point of view of quality (to the extent that a qualitatively upward evolution is easy to characterize). For this reason we limit ourselves to quantitative assessments. In this sense, the evolutions of the last millennia, even if they were exponential, occurred apparently linearly until around 1800, because exponentials had not yet shown their effects[1]. But in the last 70-80 years we have entered the area of evolution in which exponentials begin to show their "explosive" effects.

In this context we must accept that the human world is marked by

[1] As we know, an exponential function ($f(x) = a^x$) has an evolution that can be sufficiently well approximated as linear for very small x.

two developments that have proven exponential: population and technology.

1.1 The two exponentials

1.1.1 Exponential population growth

The exponential nature of population growth was not evident until the early 19th century. But it has become very evident in the last century. In Figure 2 we see a deceptively linear increase until around 1800, then a markedly explosive one after 1950.

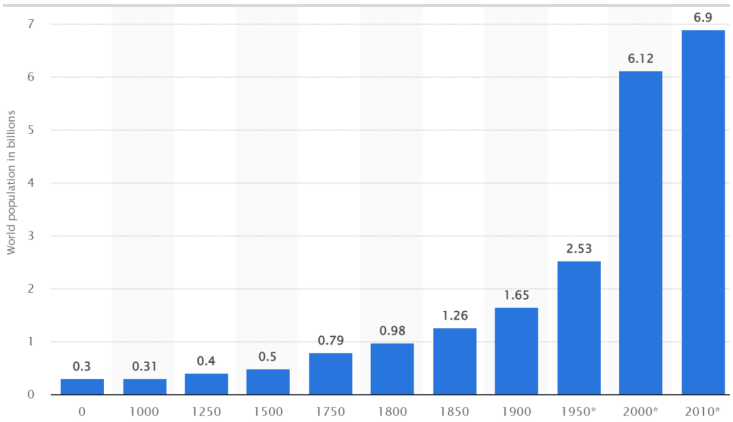


Figure 2: In about 2000 years, between 1950 and 1960, the population multiplied about 10 times, while in the last 70 years the population has multiplied > 3 times.

In conditions of exponential quantitative growth of the population, the quality of this population becomes increasingly important. Thus, the most important problems are related to the conscious behavior of the average person. The human world depends on many factors, but we are used to evaluating its quality mainly from the perspective of hierarchical organizational control structures, forgetting that everything could start from the quality of the individuals who populate it. Thus, the quality of the exponentially growing population becomes particularly important, so that in addition to the total number, the composition of the population is important. Interesting from this point of view is

Robert Kegan's approach which highlights 5 *orders of consciousness* [22, 23] offering a powerful framework for understanding how humans make meaning across the lifespan. At its core, his approach argues that development is not just about gaining knowledge but about transforming the structure of our knowing — shifting what we are subject to into what we can hold as object and reflect upon.

Kegan's model is one of the most influential constructive-developmental theories because it explains why adults differ so much in how they handle complexity, relationships, and inner life. Its strength lies in showing that psychological growth continues well into adulthood, but its limitation is that few people reach the later stages, and the model can be hard to empirically measure. Each order represents a qualitative shift in how a person constructs reality:

First Order: Impulsive mind

> Typical for early childhood (2-6 years). Experience is driven by perceptions and impulses; typical of early childhood. The child's experience is dominated by raw perception and moment-to-moment impulses. There is no stable sense of self across time. The world is experienced as a series of immediate events.

Second Order: Instrumental mind

> Typical for older children, adolescents. The self is guided by concrete needs, preferences, and rules; others matter only insofar as they affect one's own goals. The person now has stable preferences and can follow rules, but sees relationships in instrumental terms: "What do I get out of this?" Others matter only insofar as they affect one's own needs. She/he can plan and delay gratification, can follow rules and negotiate, or can understand that others have different perspectives but not coordinate them.

Third Order: Socialized mind

> Typical for most adults (unfortunately $\sim 58\%$). Subject to expectations of others, group norms, identity roles. The self is authored by the environment. Identity is shaped by relationships, expectations, social roles, belonging, loyalty, and alignment with groups, institutions, or belief systems. The person internalizes the values

of family, culture, religion, or profession. The self coheres by its alignments with, or loyalty to, the people or ideas with which it has *uncritically* become identified.

Fourth Order: Self-authored mind

Typical for a minority of adults (unfortunately only $\sim 35\%$). Subject to one's own ideology, self-constructed system. The individual becomes capable of authoring their own belief system, rather than inheriting it. They can evaluate external expectations rather than be defined by them. They can hold a coherent ideology or personal philosophy that guides decisions. The individual creates their own internal value system and can evaluate external expectations rather than be defined by them. Self directed and self authoring individuals having identity and ideology given by interpersonal relationships and reciprocity. Shortly, a self-authored mind take responsibility for its own development.

Fifth Order: Self-transforming mind

Characteristic for extremely rare adults ($\sim 1\%$). Subject to one's own self-authored system. Identity becomes fluid, dynamic, and multi-perspectival. The person can hold multiple systems simultaneously, see their limitations, and integrate contradictions. They are not defined by any single ideology, identity, or worldview.

Kegan's framework is one of the most insightful ways to understand:

- why people differ in their ability to handle complexity
- how adults can grow through transformative learning
- what psychological capacities modern life demands.

His work is especially relevant today, when individuals face unprecedented complexity and must navigate multiple, often conflicting, systems of meaning.

When more than half of the exponentially growing population is unable to critically respond to the challenges of the world it floods, ***a radical inversion is urgently needed***, in the sense that we must reach, by an appropriate ***educational reform***, a population with $\sim 35\%$ in 3-O and $\sim 58\%$ in 4-O.

1.1.2 Exponential technological growth

In the field of technologies, exponential growth is more spectacular. We witness growths qualified as *super-exponential* in technologies that evolve under the influence of their own evolution. This is the case of some technologies related to the *externalization*[2] of functions through which the noosphere and the technosphere was shaped in the human world. This is the succinct but wrongly named field of computation. In fact, it is about the externalization of the ***decision*** function in a technological space where ***information*** is the dominant actor.

The history of computation has its origins at the confluence of several subdomains of the noosphere that have interacted, over a history of more than two millennia, with advances in the technosphere space. It was not just one history, but we are dealing with several histories.

A first history, in the space of the noosphere, is a conceptual one that begins with the impossibility of deciding on a proposition uttered sometime in the 6th century BC, when the Cretan Epimenides of Knossos said: "All Cretans are liars". More than two and a half millennia followed until David Hilbert [20, 21] rigorously reformulated the problem, and Kurt Gödel [18] demonstrated, in 1931, that it has no solution, formulating the most important negative result in the history of logic and mathematics. This result triggered the appearance in 1935 of 4 works that mathematically substantiated the theory of computation [9, 24, 36, 47]. It is important to mention here Alan Turing's key paper on intelligence is "Computing Machinery and Intelligence" (1950), published in the journal Mind [48]. This is the paper in which he introduces what later became known as the Turing Test, proposing the *Imitation Game* as a practical way to discuss whether machines can "think". AI was therefore a concern from the beginning in the founders' attention.

A second history, let's call it mythical history, has its origins also in antiquity and also in the noetic space, that of the imaginary. We will list only a few moments of this history: Pygmalion, Golem, Robot. They are widely interpreted as mythic or symbolic "proto-AI milestones" be-

[2]The term *externalization* was coined by André Leroi-Gourhan. It refers to functions that, being performed by the human being with limited performance, are externalized in the technosphere [27, 28].

cause they mark humanity's recurring fascination with creating artificial beings. They are not technical steps in AI development, but they form a cultural and conceptual, i.e. noetic, lineage that predates modern computing and helps explain why AI feels both inspiring and unsettling. These stories prefigure the hopes and anxieties that later shaped real AI research, which formally begins only in the 1950's with Alan Turing, John McCarthy, and the Dartmouth workshop. Let us remember that John McCarthy is one of the "founding fathers" of AI – together with Alan Turing, Marvin Minsky, Allen Newell, and Herbert A. Simon. McCarthy, Minsky, Nathaniel Rochester and Claude E. Shannon – who coined the term "artificial intelligence" in a proposal wrote for the famous Dartmouth conference in Summer 1956. This conference started AI as a research field. Let's just remember that in the late 1950s, McCarthy discovered that primitive recursive functions could be extended to compute with symbolic expressions, resulting in the ***Lisp*** programming language [30]. We mention this event here as part of a mythical endeavor because that's how it seemed at the time it was formulated. And that's how it turned out, because it took a few ups and downs until, after passing through AI winters, we reached what is today the explosion of AI technology.

The third history is about early artifacts. *Antikythera mechanism* is the earliest physical example of a complex gear-driven machine ever discovered. The well-matched interaction of gears allows for the performance of simple arithmetic operations used to model real physical processes; for example, the motion of planets. A series of achievements, spread over the history of the next two millennia, mark the slow but steady evolution of technologies. Almost certainly many of the achievements that mark this upward trajectory are forgotten, but a few notable moments are known to us: Gerbert d'Aurillac (999), Wilhelm Schickard (1623), Blaise Pascal (1642), Gottfried Wilhelm Leibniz (1763), Charles Babbage (1822) (see Appendix A). Claude Shannon's master's thesis allows the replacement of gears with logical devices based on Boolean algebra [53], and thus we enter the era of binary-based implementations. It is the moment when the evolution of computing technology begins to manifest its exponential character. And with the advent of semiconductor technologies, Moore's law begins to act mercilessly.

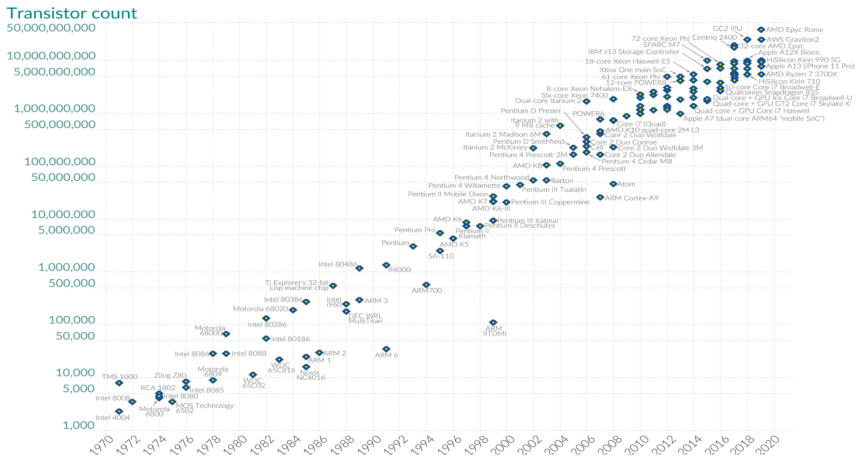


Figure 3: The evolution of the number of components per chip starting with the appearance of the first silicon memories and the first processors

When the performance of electronic technologies doubles within a 12-18 month period, during the active professional life of an engineer, increases in the tens of millions are obtained. In 50 years, from 1970 to 2020, we can see in Figure 3 how the maximum number of components on silicon has increased 50 million times! To avoid only numbers, we can use the expressiveness of images (see Figure 4).



Figure 4: How hard it was to move a 5MB IBM drive in 1956, while today's iPhones store hundreds of GB.

Why do some artificial systems grow exponentially? In semiconductor technology, software, networks, or artificial intelligence, growth is not constrained by biological limits. Artificial systems, compared to living systems, can be scaled by:

- **external addition of resources (servers, energy, memory)**

  Artificial systems can grow simply by allocating more hardware or computational capacity. Unlike organisms, which must internally generate energy and build new cells, engineered systems can instantly expand their capabilities by plugging in new servers, increasing cloud capacity, or supplying more electrical power. This allows scaling to happen at industrial speeds rather than biological ones, enabling rapid jumps in processing power, storage, or throughput.

- **automation (growth does not depend on metabolism)**

  Biological growth requires slow, resource-intensive metabolic processes. Artificial systems, by contrast, can automate nearly every step of their own expansion: deploying new software instances, replicating models, or managing infrastructure without human intervention. Because these processes are algorithmic rather than metabolic, they can operate continuously, at machine speed, and without the bottlenecks that limit living organisms.

- **network effects (each new node increases the value of the system)**

  Many artificial systems – social networks, communication platforms, distributed computing clusters – become more powerful as they grow. Every additional user, device, or node increases the total number of possible connections and interactions. This creates positive feedback loops: growth increases value, increased value attracts more growth, and the cycle accelerates. This dynamic is a major driver of exponential expansion in digital ecosystems.

- **reproduction without biological cost (instant copying)**

  Software and digital models can be duplicated perfectly and instantly, with no equivalent to the biological costs of gestation,

development, or mutation. A program can be copied millions of times in seconds, and each copy is fully functional from the moment it exists. This ability to replicate without time delays or resource-intensive processes allows artificial systems to scale at rates that biology cannot match.

Exponential quantitative developments are spectacularly transforming into qualitative developments, so that a high-tech field like artificial intelligence has entered an area where its creators can design its architecture and training, but can no longer fully understand its internal workings.

The mechanisms of technological development reach the level at which autonomous evolutionary processes begin to manifest themselves, whereby evolution tends to come under its own control. We thus externalize, at least partially, the strictly human attribute of control of the development of a highly complex high-tech field such as AI. Much greater complexities in the technological field can thus be achieved with minimized human effort accompanied by reduced control of the result.

1.2 Exponential complexity growth

We find that both exponentials that dominate quantitative developments in the human world come with qualitative problems. To face them, specific problems that were not previously encountered must be solved. The population must be educated and the autonomy of technologies tempered. Moreover, because both population and technology evolve exponentially, and their interaction produces multiplicative effects because when two exponential processes feed each other (more people $\rightarrow$ more technology $\rightarrow$ more connected people $\rightarrow$ more technology), the result is a *super-exponential* increase in complexity.

The main effect of exponential quantitative growth is the increase in complexity in the human world in which man must integrate responsibly towards himself, towards his world, but also towards nature. Isn't that too much?

1.2.1 Population exponential complexity

The functioning of our world is disrupted by the immense complexity that characterizes it, which can no longer be mastered at any of the levels at which it manifests itself:

- **State-level**: characterized by incompetence & corruption

  As states grow in scale and responsibility, they become labyrinths of competing interests, opaque procedures, and inherited structures. The modern state aspires to omniscience-regulating economies, safeguarding rights, managing crises-yet it remains bound by the limits of human judgment and institutional inertia. In such a vast machinery, incompetence is not merely a personal failing but a structural inevitability, and corruption becomes the shadow cast by power when transparency dissolves. The state, once imagined as the rational guardian of the common good, now often appears overwhelmed by the very complexity it helped create.

- **Corporate-level**: driven by greediness

  Corporations have become the great engines of modern civilization, yet their logic is fundamentally instrumental: to grow, to accumulate, to outcompete. In a world where value is measured in quarterly increments, the pursuit of profit becomes detached from any broader conception of human flourishing. Greed here is not a vice of individuals but the governing principle of an economic order that rewards extraction over stewardship. Corporations, in their relentless expansion, mirror the boundless appetite of the systems they inhabit, systems that know how to produce wealth but not how to define its purpose.

- **NGO-level**: dominated by destructive competition

  Even institutions founded on ideals are not immune to the pressures of scarcity and visibility. NGOs, meant to embody moral intention, often find themselves locked in subtle battles for funding, legitimacy, and narrative control. Their missions converge, their resources diverge, and cooperation becomes a luxury they cannot afford. In this paradox, organizations dedicated to healing

the world can inadvertently reproduce its fractures, as competition erodes solidarity and the pursuit of impact becomes entangled with the pursuit of survival.

- **Community-level**: limited to tribalism

  Faced with global uncertainty and overwhelming information, communities retreat into smaller circles of identity and belonging. Tribalism emerges not only as a social instinct but as a defense mechanism against the disorienting vastness of modern life. The local becomes a refuge, the familiar a shield. Yet this narrowing of perspective fragments the social fabric, replacing shared horizons with competing narratives. What once bound communities together – common purpose, mutual recognition – gives way to suspicion and the comfort of echo chambers.

- **Human-level**: lack of education as the ability to decide between good and evil

  At the deepest level lies the individual, confronted with a world too complex to grasp and too fast to interpret. Education, in its truest sense, is the cultivation of discernment: the capacity to distinguish truth from illusion, justice from convenience, good from harm. But when education becomes mere training, or information becomes noise, this moral compass falters. The human being, overwhelmed by choices and deprived of guidance, becomes vulnerable to manipulation, confusion, and moral drift. In such a condition, the ancient task of choosing the good becomes a struggle against the entropy of the modern mind.

At all these levels, the living, represented mainly by the human mind, must cope with the complexity that is growing super-exponentially. With what capacity for adaptation and evolution? In no case with one that is also growing exponentially.

1.2.2 AI technology exponential complexity

Empirical evidence: AI capacity doubles every ~ 7 months. AI technology grows exponentially because several reinforcing forces accelerate one another, creating a self-amplifying cycle of improvement. At its

core, AI advances quickly because better hardware, better algorithms, more data, and more investment all compound together, each enabling the next leap forward.

The main factors behind this exponential grow are:

- **Explosive improvements in computing hardware**

  Modern AI runs on highly specialized chips that improve far faster than traditional CPUs. For example, Nvidia's H100 GPU delivers roughly six times the compute performance of the previous A100 generation, and generates text 2× faster in practical use. These leaps allow researchers to train models in weeks instead of months, enabling rapid iteration and scaling.

- **Breakthrough algorithms that scale with compute**

  The 2017 introduction of the Transformer architecture allowed AI models to process information in parallel rather than step-by-step. This removed a major bottleneck and made it possible to train enormous models simply by adding more hardware. This architectural shift underlies nearly all modern AI systems.

- **Predictable scaling laws**

  Researchers discovered that AI performance improves in a mathematically predictable way when you increase model size and training data together. This turned AI progress into an engineering discipline: if you double compute, data, and parameters, you can reliably predict how much better the model will become.

- **Massive growth in available data**

  The global *datasphere* is exploding – expected to reach 175 zettabytes by 2025 – providing AI systems with unprecedented training material. More data means richer patterns to learn from, which directly boosts model capability.

- **Unprecedented investment and industrial focus**

  Governments and corporations are pouring billions into AI research, infrastructure, and national strategies. Over 60 countries now have formal AI strategies, accelerating development and deployment.

- **Positive feedback loops**

  As AI becomes more capable, it becomes more useful. As it becomes more useful, more people adopt it. As adoption grows, more data, money, and talent flow into the field—accelerating the next wave of breakthroughs. This compounding dynamic is the essence of exponential growth.

Each factor amplifies the others, creating a technological ecosystem that accelerates itself.

Moreover, AI technology is developing as a product whose complexity is increased by the form in which it materializes: a suite of weight matrices whose direct interpretation is impossible. If a conventional software product, even that of AI products in the early stages of development of this technology, is presented as a program written in a programming language that could be read and understood by a human programmer, the matrices associated with current AI applications have no meaning to those who inspect them. We cannot answer the question ”how does an AI program work?” because the scalar matrices that define an application are established through an iterative approximation process used to ”train” artificial neural networks.

In other words, the complexity of this tool is also accentuated by the obscurity produced by our inability to verify the result of the ”programming” except by evaluating the running results which are marked by approximations that cannot be reduced to zero.

2 The logarithmic behavior of living

How can we explain the logarithmic growth (see the green form in Figure 1) observed in the living world? In short, logarithmic growth stems from biological constraints that limit growth in organisms, while artificial systems can be scaled almost indefinitely by external addition of resources, leading to exponential growth. More precisely, in biology, growth is limited by physical laws, finite resources, and internal regulatory mechanisms. Even though some biological processes may start exponentially (e.g., bacterial multiplication), almost all quickly reach a saturation zone, where the rate slows down. This slowdown produces

logistic[3] or logarithmic curves. Key factors that determine this type of growth are:

- **Limited resources: nutrients, space, metabolic energy**

  Every living system is born into a world of finitude. Nutrients are not infinite, space is not boundless, and energy must be constantly renewed through fragile metabolic pathways. As a population grows, each individual competes for the same shrinking pool of necessities. What begins as abundance gradually becomes scarcity, and scarcity imposes its own mathematics. The exponential curve bends not because life lacks ambition, but because the world imposes limits that no organism can transcend.

- **Homeostasis: organisms cannot grow indefinitely; there are strict internal balances**

  Life is not merely expansion, it is equilibrium. Every organism carries within it a delicate architecture of feedback loops, hormonal signals, and regulatory mechanisms designed to preserve internal stability. Growth is permitted only insofar as it does not threaten this balance. Beyond a certain point, further expansion becomes destabilizing, and the organism must slow down or stop. Homeostasis is the quiet philosopher of biology, reminding every living being that harmony, not magnitude, is the true condition of survival.

- **Energy cost: each additional unit of growth costs more (law of diminishing returns)**

  In the early stages of growth, energy is used efficiently: the simplest tasks yield the greatest returns. But as complexity increases, so does the cost of maintaining it. Larger bodies require more energy to sustain; more cells mean more repair, more coordination, more vulnerability. The law of diminishing returns is not merely an economic principle—it is a biological truth. Nature whispers that growth is easy at the beginning, but each step upward demands a heavier toll.

[3] We postpone the investigation of logistics processes for a later specialized text.

- **Natural selection: evolution favors stability, not unlimited expansion**

  Evolution does not reward the organism that grows the fastest, but the one that endures. Unlimited expansion is inherently unstable: it exhausts resources, destabilizes ecosystems, and invites collapse. Over millions of years, natural selection has sculpted life toward strategies that balance growth with restraint. The species that survive are those that learn to live within limits, not those that attempt to transcend them. In this sense, evolution is a teacher of moderation, not excess.

The result: rapid growth at first, then progressive slowdown—a pattern that, when translated into mathematics, becomes the familiar shape of logarithmic or logistic growth. It is the curve of life itself: a brief surge of possibility, followed by the wisdom of limits. As long as the evolution in the technosphere and the noosphere was well approximated linearly, the interaction between the two spheres was moderate. But, once the exponentials started to show their specificity, the mutual actions become very complex. From the technosphere we will take into consideration a high-tech, that is, AI, and from the living sphere the human mind to see how they evolve and co-evolve.

3 AI vs. human evolution

The differences between the exponential way AI evolves and the logarithmic way humans evolve are mainly due to the following facts:

- **The living is optimized for survival, not for unlimited growth**

  Life, in its deepest essence, is conservative. Every organism is shaped by millions of years of evolution whose primary mandate is not expansion but persistence. Survival requires balance, restraint, and adaptation to finite environments. Growth is permitted only insofar as it does not endanger the organism or the ecosystem that sustains it. Thus, biological evolution follows a rhythm of cautious refinement rather than explosive acceleration. It is a dance with constraints, not an escape from them.

- **The artificial is optimized for performance and scaling, not for equilibrium**

  Technology, by contrast, is born from intention rather than natural necessity. It is designed to amplify, accelerate, and extend human capabilities. Its logic is not survival but optimization: more speed, more power, more efficiency. Unlike living beings, artificial systems have no internal need for balance; they do not tire, age, or self-regulate unless we program them to do so. Their growth is limited only by external resources and human ambition. As a result, they follow trajectories that resemble pure mathematical expansion, unconstrained by the biological wisdom of limits.

- **Biology is about closed systems, technology is about open systems**

  Life evolves within the boundaries of the biosphere: a finite planet with finite energy, finite materials, and finite ecological niches. Every gain for one species is a cost for another; every expansion must eventually reconcile with the carrying capacity of the environment. Technology, however, operates in an open conceptual space. It can draw on globalized resources, distributed computation, and the collective intelligence of billions of humans. It can be copied, scaled, and recombined without the metabolic costs that bind living organisms. In this openness, artificial systems escape the gravitational pull of natural limits and accelerate into exponential trajectories.

In essence: human evolution is shaped by the logic of life-slow, adaptive, self-preserving, while AI evolution is shaped by the logic of design-fast, expansive, self-amplifying.

One curve bends toward equilibrium (logarithmical); the other races toward acceleration (exponential).

In the evolution of artificial intelligence, exponential growth occurs because artificial systems are not constrained by biological limits such as genetic variation, metabolic energy, or homeostasis. Biology evolves slowly due to physical and genetic limits, while AI can be externally scaled, continuously optimized, and redesigned without these barriers.

3.1 Living evolution

Recent research shows that biological evolution is limited by genetic variability, mutations, genetic drift, and physical constraints. The maximum rate at which biological traits can evolve is strictly related to the variation available in the population and to the laws of physics that govern organisms [15, 16].

Biological evolution unfolds within a closed circle of possibilities. Every species inherits a finite reservoir of genetic diversity, and natural selection can only act on what already exists or what random mutation provides [14, 29]. Evolution is therefore not an engine of boundless creativity but a sculptor working with a limited block of marble. The pace of change is constrained by the slow arrival of new mutations, the statistical fluctuations of genetic drift, and the fact that most variations are neutral or harmful rather than beneficial.

Studies show that the rate of change of a biological character is limited by the variability of fitness and the trait in question. In other words, evolution cannot push a trait faster than the population's own internal diversity allows. If variation is small, progress is slow; if variation is exhausted, progress stops. This is why evolutionary biologists speak of "response to selection" as a function of heritability and variance: nature cannot select what is not there. The evolutionary trajectory is therefore bounded by the statistical structure of life itself [31, 54, 39].

Organisms are constrained by biomechanical and physical laws that set upper limits on size, metabolism, and performance, leading to evolutionary asymptotes. No creature can escape gravity, thermodynamics, or the geometry of its own body. The largest animals are limited by the strength of bones; the fastest metabolisms are limited by oxygen diffusion; the most complex brains are limited by energy consumption. These constraints act like invisible ceilings: as organisms approach them, the cost of further improvement rises sharply, and evolution slows to a crawl. The curve of biological progress bends toward an asymptote because the universe itself imposes boundaries [52, 51].

These limits make biological evolution slow, incremental, and saturated, mathematically approximating a logarithmic curve. Life begins with bursts of possibility—rapid diversification, explosive radiation—but soon encounters the walls of its environment and its own

physiology. Growth slows, variation narrows, and the system settles into equilibrium. The mathematics of evolution mirrors the philosophy of life: initial expansion followed by the wisdom of limits.

Conclusion: biology is a closed system, constrained by physical laws and internal resources. Its creativity is real but bounded; its progress is profound but gradual. Evolution is not an exponential ascent but a long, patient negotiation with the constraints of matter, energy, and time.

3.2 AI evolution and the limits of societal adaptation

AI is developing at a pace that exceeds the adaptive capacity of human institutions. Unlike biological evolution – which proceeds through slow, incremental, and fundamentally logarithmic processes – AI systems evolve through stacked innovation, a compounding dynamic in which each generation of AI accelerates the development of the next [25]. This creates a feedback loop of super-exponential growth, enabling AI to surpass traditional technological progress models and outpace the mechanisms societies rely on to regulate, absorb, and respond to new technologies.

This section examines the structural mismatch between AI's acceleration and society's adaptive bandwidth, drawing on empirical scaling laws, institutional theory, and sectorspecific case studies.

3.2.1 Stacked AI Innovation and Super-Exponential Growth

Biological evolution proceeds through slow, path-dependent processes constrained by reproduction cycles, mutation rates, and environmental pressures. AI evolution, by contrast, is driven by intentional, cumulative, and accelerating innovation loops [45]. The four mechanisms below explain why AI capabilities grow orders of magnitude faster than biological systems.

1. **Compute scaling**

   Compute scaling refers to the rapid increase in the amount of computational power available for training AI models. Unlike biological systems, which are limited by metabolic constraints

and physical reproduction, AI systems can leverage: Moore's Law-style hardware improvements (GPUs, TPUs, custom accelerators), massive parallelization across thousands of chips, cloud supercomputers with exascale performance, specialized architectures (e.g., tensor cores, optical computing, neuromorphic chips). Thus AI evolution is accelerated because: more compute allows training larger models with more parameters, larger models tend to exhibit emergent capabilities not present in smaller ones, compute can be scaled deliberately and continuously, unlike biological evolution, scaling laws show predictable performance gains as compute increases.

All these are in contrast with biology because biological organisms cannot arbitrarily increase their "compute" (neuronal count, synaptic density) without hitting physical limits such as skull size, energy consumption, or developmental constraints.

2. **Algorithmic efficiency improvements**

   Efficiency gains by improving algorithms often produce $2-10\times$ improvements without increasing compute. In contrast with biology whose evolution cannot be done by "rewrite its own algorithms" in a directed way. Mutation-driven improvements are random, slow, and constrained by survival pressures.

3. **Model architecture innovations**

   The most famous example is the Transformer architecture [50], which replaced recurrent networks and enabled the modern AI revolution. Architectural changes can produce qualitative leaps in capability, they allow AI to handle new modalities (vision, audio, robotics, reasoning), enable scaling to trillions of parameters, and create new pathways for emergent behavior. In contrast with biological architectures (e.g., cortical layers, sensory systems) which evolve over millions of years and cannot be redesigned wholesale. AI architectures can be redesigned in months.

4. **Self-improvement loops (e.g., AI-assisted coding, AI-driven research acceleration)**

> When artificial intelligence systems directly contribute to their own development, a spectacular additional increase in performance occurs. The main such as mechanisms are: AI-assisted coding, AI-driven research acceleration, AI-enabled automation of experimentation. In contrast, biological evolution cannot intentionally accelerate itself. It lacks meta-cognition, foresight, and the ability to redesign its own mechanisms.

We can conclude: AI evolution is directed, cumulative, and accelerating, whereas biological evolution is slow, random, and constrained.

Empirical studies show that compute used in frontier AI training has increased by 4× per year since 2010 [2], while algorithmic efficiency has improved by 2.5× per year [40]. Combined, these trends produce nearly 10× annual growth in frontier AI capability, an unprecedented rate in the history of technological development.

This phenomenon, often described as stacked AI innovation, means that AI is not merely improving but accelerating its own improvement. As AI systems become more capable of autonomous problem-solving, code generation, and scientific discovery, they amplify their own development trajectory [46]. This creates a super-exponential curve fundamentally different from the slow, bounded dynamics of biological evolution.

3.2.2 The Institutional Lag: Why Society Cannot Keep Pace

Structural inertia in governance systems Institutions-laws, regulatory bodies, educational systems, and labor markets operate on decadal timescales. They require consensus formation, political negotiation, bureaucratic implementation, and cultural acceptance [1]. Historically, institutions have adapted to technological change, but always with significant delay [5].

AI's acceleration compresses the time available for adaptation. While AI capabilities may double every 6-12 months, institutional reforms often require 10-30 years to mature. This creates a widening governance gap in which AI systems evolve faster than the mechanisms designed to regulate them.

Cognitive and cultural lag Beyond formal institutions, societies face cognitive and cultural constraints. Public understanding of AI risks, labor market impacts, and ethical implications lags behind technological reality. This mismatch reduces democratic oversight and increases the likelihood of unanticipated systemic disruptions.

3.2.3 Sectorial Transformation: AI Outpaces Social Interpretation

AI's rapid deployment across industries further amplifies the adaptation gap. Sectors such as finance, healthcare, logistics, and education can be reconfigured within 2-5 years due to the modular, software-driven nature of AI systems [32]. By contrast, workforce retraining, regulatory redesign, and cultural integration typically require 10-30 years.

Examples of rapid transformation:

- **Finance**:

  AI systems are reshaping financial markets at a pace that far exceeds the ability of regulatory bodies to adapt. High-frequency trading algorithms now operate on millisecond timescales, executing strategies that no human trader can monitor or interpret in real time. Machine learning-based fraud detection systems continuously update their models using live transaction data, enabling them to identify anomalous patterns long before traditional auditing processes would detect them. Similarly, AI-driven risk modeling tools can ingest vast quantities of market, credit, and behavioral data, producing dynamic risk assessments that evolve daily rather than quarterly.

  Regulatory frameworks, by contrast, rely on slow legislative cycles, public consultations, and multi-year implementation timelines. As a result, oversight mechanisms often lag behind the technologies they aim to regulate. This creates systemic vulnerabilities: flash crashes triggered by algorithmic interactions, opaque model-driven credit decisions, and the emergence of "black-box finance" in which neither regulators nor institutions fully understand the behavior of deployed AI systems.

- **Healthcare**:

In healthcare, AI-based diagnostic models and clinical decision-support systems are advancing at a rate that challenges existing medical governance structures. Deep-learning models can now outperform human clinicians in tasks such as radiology interpretation, dermatological screening, and pathology classification. These systems improve rapidly as they are exposed to more data, often achieving performance gains within months.

However, medical accreditation, liability frameworks, and clinical trial protocols evolve slowly. Certifying a new medical device or diagnostic tool can take 5–10 years, and malpractice laws are not yet designed to handle scenarios in which responsibility is shared between clinicians and autonomous systems. Ethical review boards, insurance providers, and hospital procurement systems also operate on multi-year cycles, creating a structural mismatch between technological capability and institutional readiness.

This gap raises critical questions: Who is liable when an AI system misdiagnoses a patient? How should clinicians be trained to supervise or override AI recommendations? What constitutes informed consent when AI is involved in treatment decisions? These issues remain unresolved even as AI tools enter clinical practice.

- **Education**:

  AI is transforming education through personalized tutoring systems, automated assessment tools, and adaptive learning platforms. Large language models can now generate explanations, quizzes, feedback, and curriculum materials tailored to individual students. These systems can be deployed globally within months, rapidly altering how students learn and how teachers structure their work.

  Yet educational institutions – schools, universities, accreditation bodies – are among the slowest to change. Curriculum reform often requires multi-year review cycles, political negotiation, and teacher retraining. Credentialing systems, standardized testing frameworks, and pedagogical norms are deeply embedded and

resistant to rapid modification.

As a result, AI tools are entering classrooms before societies have agreed on fundamental questions: What skills should students learn in an AI-rich world? How should academic integrity be defined when AI can generate essays and solve problems? What role should teachers play when AI tutors can provide individualized instruction at scale? The technology is moving faster than the pedagogical and ethical frameworks needed to govern it.

- **Logistics**:

  In logistics, AI-driven optimization systems, autonomous routing algorithms, and robotics are transforming supply chains with unprecedented speed. Predictive models can forecast demand, optimize warehouse layouts, and dynamically reroute shipments in response to real-time conditions. Autonomous vehicles and drones are beginning to handle last-mile delivery, while robotic picking systems are increasingly common in warehouses.

  These changes can be implemented within 2–5 years because logistics is heavily software-mediated and highly sensitive to efficiency gains. However, labor markets and regulatory systems cannot adapt nearly as quickly. Workers require retraining to operate alongside autonomous systems, and many roles – such as truck driving, warehouse picking, and dispatch coordination – face displacement risks long before new employment pathways are established.

  Regulatory bodies also struggle to keep pace. Autonomous delivery vehicles require new safety standards, liability rules, and infrastructure adaptations. Cross-border logistics introduces additional complexity, as international regulations evolve even more slowly than national ones.

  The result is a sector in which technological capability races ahead of workforce readiness, legal clarity, and social consensus.

In each case, AI systems reshape economic and social structures before society fully understands their implications.

3.2.4 Conclusion

AI's super-exponential growth, driven by stacked innovation and self-reinforcing improvement loops, creates a structural mismatch between technological acceleration and societal adaptation. Institutions, cultural norms, and labor systems evolve too slowly to keep pace with AI's rapid transformation of economic and social landscapes. This widening gap underscores the need for proactive governance frameworks, anticipatory regulation, and new models of institutional agility.

3.3 Co-evolution

What are the main effects of the interaction between living humanity and formal-structural AI in the context in which the development of both exponentially increases the complexity of the world? We are in the middle of the process, which is why assessments are risky, especially since the effects in the economic space are evolving very strongly influenced even by the current incipient forms of AI. Several aspects listed below are eloquent.

The labor market is changing faster than people can retrain.

Indeed, an arXiv study shows that 40–70% of tasks in sectors such as healthcare, education, retail, and manufacturing can be automated. While AI operates 24/7 at extremely low cost, humans have biological and social limits.

Result: massive pressure on jobs, with the need for transition policies (retraining, reduced workweek).

Economies may become dependent on AI infrastructures that they cannot control. The World Bank report highlights that many countries do not have the necessary infrastructure (computing, connectivity, data, skills) to adopt AI at the pace at which it evolves.

The consequent risk: growing global inequalities and the inability of some societies to keep up.

Economic uncertainty increases according to the World Economic Forum that shows 54% of business leaders expect jobs to be replaced, but only 24% expect new ones to be created.

This asymmetry creates social and political tensions.

3.3.1 Can society keep up?

The central question is whether societies, whose institutions and cultural norms evolve at a slow, logarithmic pace, can adapt to the rapid, exponential transformations induced by advanced AI systems. The answer is potentially yes, but only if deliberate, large-scale strategies are implemented to accelerate institutional, economic, and cultural adaptation. Below, each proposed strategy is expanded into a research-level analysis.

Massive investments in digital skills and retraining. AI reshapes labor markets far faster than traditional education systems can respond. To prevent structural unemployment and skill mismatches, societies must invest in continuous, lifelong learning programs that equip workers with digital literacy, data fluency, and AI-augmented competencies.

This requires:

- national retraining initiatives modeled on large-scale industrial transitions (e.g., post-WWII workforce programs)
- public–private partnerships to align training with real labor market needs
- modular, rapid-cycle education systems that can update curricula annually rather than every decade
- incentives for employers to adopt "learning-first" workforce models.

Without such investments, the gap between AI-driven job transformation and human skill acquisition will widen, producing long-term inequality and economic instability.

Transitional policies (e.g., 4-day workweek). As AI automates cognitive and routine tasks, societies must develop transitional labor policies that cushion the pace of change. A shorter workweek is one such mechanism: it distributes productivity gains across the population while reducing the shock of rapid automation.

Transitional policies may also include:

- wage insurance for displaced workers
- phased retirement programs
- job-sharing schemes
- universal basic mobility (support for retraining, relocation, and career transitions).

These policies act as "shock absorbers," giving institutions and individuals time to adapt to AI-driven restructuring without triggering social unrest or mass displacement.

Accessible AI infrastructure (compute, cloud, data). If AI infrastructure remains concentrated in a handful of corporations or geopolitical blocs, the benefits of AI will be unevenly distributed and societal adaptation will be hindered. Broad access to compute, cloud platforms, and high-quality datasets enables:

- small and medium enterprises to adopt AI rather than be displaced by it
- universities and public institutions to conduct independent research
- governments to build regulatory capacity based on real technical understanding
- communities and civil society to develop AI tools aligned with local needs.

Democratizing AI infrastructure reduces the risk of "AI feudalism," in which only a few actors possess the computational power to shape economic and cultural outcomes.

Agile and adaptable regulations. Traditional regulatory models – slow, reactive, and rigid – are incompatible with technologies that evolve on 6-to 12-month cycles. To keep pace with exponential AI development, societies must adopt regulatory frameworks that are:

- adaptive (updated continuously as capabilities evolve)

- risk-tiered (stricter oversight for high-risk applications)
- evidence-based (grounded in empirical evaluations of real-world performance)
- technically informed (supported by public-sector AI expertise).

Examples include regulatory sandboxes, dynamic licensing, and real-time auditing systems. Such frameworks allow innovation while maintaining safety, transparency, and accountability.

Equitable distribution of the benefits of AI. Exponential technologies tend to produce exponential inequalities unless counterbalanced by deliberate policy. Ensuring that AI's benefits are widely shared is essential for social stability and democratic legitimacy.

Equitable distribution may involve:

- progressive taxation of AI-driven productivity gains
- public investment in AI-enhanced health care, education, and infrastructure
- community owned data trusts
- mechanisms to ensure that AI augments rather than replaces human labor.

If the gains from AI accrue only to a narrow elite, societal resistance will intensify, slowing adoption and undermining the very innovation that drives progress.

These directions are already recommended in all positions on the effects of the emergence of AI applications.

3.3.2 Can humans keep up?

The human being seems to be the only part of the living in the human world whose evolution we can influence in such a way that it can compete in co-evolution with an exponential process. The flexibility of the human mind offers us this chance. It remains to be seen to what extent we will be able to make use of it.

Humans can keep up with AI evolution, but only through a process of co-evolution in which humans and AI influence each other. Without this continuous feedback, the pace of AI would outpace our natural capacity to adjust. What does recent research tell us about this?

Human-AI systems constitute coupled adaptive dynamical systems in which human behavior and AI models recursively influence each other through continuous data-driven feedback loops (often described as human-AI co-evolution). This co-evolution is considered an emerging phenomenon essential for understanding the future of society [17, 35, 44].

Humans exhibit substantial cognitive plasticity in human-AI interaction, dynamically adapting decision strategies, mental models, and metacognitive processes in response to AI behavior, often more rapidly and systematically than naïve intuition suggests. Indeed, research in evolutionary psychology and neuroscience shows that the human mind is not "fixed in the Stone Age," but is deeply plastic and capable of adapting to new technologies. This plasticity is a strong argument that humans can keep up with AI, if exposed to the right learning environments [6, 42].

Most of researchers warn that without mechanisms for social and educational adaptation, AI could outpace society's ability to adjust because the pace of AI is much faster than the natural pace of human evolution [55, 56]. AI does not evolve biologically, but through accelerated technological development.

Can humans keep up? Yes, but with the following minimal conditions:

Continuous cognitive adaptation by their neural plasticity which allows humans to learn new ways to work with AI. Studies emphasize that constant interaction with technology gradually shapes the way we think.

Guided human-AI co-evolution through education, public policy, and ethical system design. Without this framework, AI could influence society in unintended and, worse, counterproductive ways.

Integrating AI as augmentation, not competition makes us "better as humans," not replace us. This means systems that augment human abilities, not replace them.

The risks that may arise in this process are due to the fact that the pace of AI is technological (exponential), not biological (logarithmic). AI can advance in months what humans can adapt in years. This asymmetry creates risks of:

educational gap between those who accept the use of the latest technologies and those who reject them

social polarization between those who have access to the use of the latest technologies and those who do not

cognitive dependence on automated systems is considered by some to limit the development of the ability to discern in complex contexts

loss of control over the social feedback generated by algorithms due to their efficiency in the noosphere

modification of human perception of preferences, behaviors and even cultural identity.

The risks listed can be, at least partially, avoided provided that human beings can keep up with AI by evolving with it through education, cognitive adaptation, intelligent policies and ethical design. Without these mechanisms, the pace of AI irremediably exceeds the natural adaptive capacity of humans and society.

3.3.3 What to do?

What are the human skills that should be developed as a priority to cope, with our limited logarithmic evolutionary capacity, with the exponential growth of AI? Humans can keep up with AI, but only if they develop a clear set of human skills that become more valuable as AI expands. The most important are critical thinking, discernment, resilience, emotional intelligence, collaboration, and the ability to manage AI.

Critical Thinking and Discernment become of utmost importance in a world where AI can rapidly generate information. In these conditions, humans must be able to evaluate sources, identify bias, and validate facts. AI can produce content that is plausible but sometimes

incorrect – and humans remain responsible for fact-checking. This is important because AI does not, at its current stage of development (surprises must be taken into account), understand context or ethical implications and humans must decide what is correct, relevant and safe.

Emotional Intelligence and Social Skills is another factor because education and the job market are shifting towards skills like empathy, interpersonal communication, collaboration, and emotion management. These are deeply human and difficult for AI to replicate. These become important in the context of future jobs that will involve mixed human-AI teams, and humans will need to excel in areas where human relationships and understanding are crucial.

Creativity in complex problem solving are important under the conditions in which AI can suggest solutions, but it cannot fully understand the human context or the nuances of a situation. Humans remain essential in innovation, design, strategy, and lateral thinking. This ability matters when real problems are ambiguous, because AI cannot handle ambiguity at the human level.

Resilience and adaptability in a rapidly changing environment become skills of utmost importance. Industry leaders emphasize that resilience, curiosity, and teamwork are more important than ever. These abilities matter because AI is changing roles, industries, and processes; people need to be able to constantly recalibrate.

Basic technical skills in AI become very useful because, although you won't need to become a high-performance programmer, you need to know how to use artificial intelligence, how to manage it, and how to collaborate with it. These skills are among the most sought-after in the coming period. This ability is particularly important because people who know how to work with AI become productivity multipliers.

Indeed, humans can co-exist with AI, but only if they develop their deeply human skills and learn to collaborate intelligently with AI. ***We will engage in a process of deeper understanding not only of what AI is, but also of deepening what humans actually are***.

Don't we have the arrogance to grant humans exaggerated capabilities beyond realistic assessments of what humans really are? Before we have a sufficiently advanced knowledge of what consciousness, intentionality, imagination, intuition, empathy, humor, and the like are, we arrogantly assert that AI lacks all of these.

Let's take, for example, the issue of consciousness. Indeed, if some of the most influential philosophers of mind argue that consciousness does not exist (or is not what people think it is), then demanding that AI be "conscious" becomes conceptually confused. Daniel Dennett, Patricia Churchland, and Paul Churchland are often grouped under *eliminative materialism* or *illusionism*. Their core claim is not that "nothing is happening in the brain," but that: what people call "consciousness", the inner movie, the subjective glow, the Cartesian theater, ... is a *misleading folk concept*, not a scientifically valid entity.

Dennett argues that consciousness is a user illusion generated by distributed cognitive processes [12] [13]. The Churchlands argue that folkpsychological terms like "belief," "qualia," or "consciousness" will eventually be replaced by neuroscientific vocabulary, just as "phlogiston" was replaced by oxygen [10] [11].

Indeed, complex processes on which researchers do not have convergent views could possibly be the product of the composition of simpler processes for which we have less controversial representations. A notable example in this regard is Daniel Dennett's view on the phenomenon of consciousness [12], which he considers to be an emergent process of the contributions of simpler processes, such as perception, attention, ... If Dennett turns out to be right, then the objection of some that AI does not have consciousness would fall.

If consciousness is an illusion even in humans, then why ask AI to possess consciousness? You cannot require a machine to have something that, according to this view, no one has in the first place. If Dennett and Churchlands are right, and consciousness is not a scientifically robust concept, then demanding consciousness from AI is pointless. The ethical and safety issues of AI do not depend on consciousness, and we should focus on behavior, not qualia. In short, maybe we don't need conscious AI, nor should we ask for it.

In this context, ***I believe that the development and use of AI technology will also stimulate anthropological research that will be able***

to define more clearly what humans are. We could conclude that the differences between the human mind and AI are predominantly quantitative. In this case, the exponential of AI technology and the logarithm of living adaptability in humanity will no longer oppose each other so dramatically qualitatively, and people will more relaxedly accept competition with a techno-object. We, as humans, do not feel complexed by the speed at which a car moves or the size of the weights that cranes handle. Under these conditions, we will be more relaxed and accept AI as an accelerator of our capabilities to solve complex problems that, with certainty, have reached limits that even the best-equipped minds cannot cope with. AI should be seen, in this case, only as another *externalization*, in the sense of André Leroi-Gourhan, through which a mechanism present but limited in the behavior of the human mind is found accelerated in a technical object.

Moreover, perhaps the image we have of the human being will be more correctly positioned on a less elevated pedestal, closer to what he actually is.

4 Concluding remarks

1. Human beings can keep up with AI, but only if they develop and stimulate their deeply human abilities and learn to collaborate intelligently with AI. We should not compete with AI, but become complementary to it. But for this it will be necessary, as a first condition, the inversion from $\sim 35\%$ in 4-O and $\sim 58\%$ in 3-O to $\sim 35\%$ in 3-O and $\sim 58\%$ in 4-O according to Kegan's classification.

2. Because most AI users use this tool as a kind of more efficient Google, it is too early to conclude consistently on what AI is in relation with humanity.

3. Using AI to its true value increases the user's efficiency in completing complex actions and projects. But, more than that, it fundamentally changes the level at which the activation of the human mind can function. In many areas we can advance from laborious structural approaches to spectacular architectural approaches.

4. From the interaction with AI, man will gradually understand more about what he himself is.

5. For now, we believe that AI has no consciousness, no intentions, no values of its own, no lived experience, no moral responsibility. But for all these human abilities, we do not have a unanimously accepted representation that would allow us to invoke them when evaluating what AI can and cannot do. If we manage to properly temper our image of the human being, perhaps we will find AI more suitable for a consistent interaction with man and his world.

6. Perhaps the most important question that arises regarding the impact of AI on humanism is the Heideggerian one [19]: what mode of revealing being does AI institute?

References

[1] Daron Acemoglu & Simon Johnson (2023) *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. PublicAffairs.

[2] Dario Amodei & Danny Hernandez (2018) *AI and Compute*, OpenAI.

[3] George Boole (1847) *The Mathematical Analysis of Logic – Being an Essay Towards a Calculus of Deductive Reasoning*, London, England: Macmillan, Barclay, & Macmillan.

[4] George Boole (1854) *The Laws of Thought*, Walton & Maberly.

[5] Erik Brynjolfsson & Andrew McAfee (2014) *The Second Machine Age*. W. W. Norton.

[6] Hechang Cai & Jinlai Zhou (2026) User cognitive fit in human-AI interaction: Exploring the link between input representation and generative output complexity, *Computers in Human Behavior*, **179**.

[7] Pierre Teilhard de Chardin (Nov. 1930) Le phénomène humain, *Revue des questions scientifiques*, pp. 390–406.

[8] Pierre Teilhard de Chardin (1955) *Le phénomène humain*, Éditions du Seuil.

[9] A. Church (1936) An unsolvable problem of elementary number theory, *American Journal of Mathematics* **58**(2): 345–363.

[10] Paul M. Churchland (1981) Eliminative Materialism and the Propositional Attitudes. *Journal of Philosophy* **78**(2):67-90.

[11] Paul M. Churchland (1988). Matter and Consciousness. MIT Press.

[12] Daniel C. Dennett (1991) Consciousness Explained, The Penguin Press.

[13] Daniel C. Dennett (2018). Facing Up to the Hard Question of Consciousness. *Philosophical Transactions of the Royal Society B* **373**(1755).

[14] Adam Eyre-Walker & Peter D. Keightley (2007) The distribution of fitness effects of new mutations. *Nature Reviews Genetics* **8**:610-618.

[15] Ronald A. Fisher (1930) *The Genetical Theory of Natural Selection*. Oxford University Press.

[16] Douglas S. Falconer & Trudy F. C. Mackay (1996) *Introduction to Quantitative Genetics*. Longman, 1996.

[17] Moshe Glickman, & Tali Sharot (2025) How human–AI feedback loops alter human perceptual, emotional and social judgements, *Nature Human Behaviour* **9**:345–359.

[18] K. Gödel (1931) Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme, I (On Formally Undecidable Propositions of Principia Mathematica and Related Systems), *Monatshefte für Mathematik und Physik* **38**:173–198.

[19] Martin Heidegger (originally published in 1954) The Question Concerning Technology. *Basic Writings, edited by David Farrell Krell*, pp. 307-342, New York: Harper & Row, 1977.

[20] D. Hilbert (1900) *Mathematical Problems. Lecture Delivered Before the International Congress of Mathematicians at Paris in 1900.*

[21] David Hilbert, Wilhelm Ackermann (1928) *Grundzüge der theoretischen Logik* (Principles of Mathematical Logic), Springer-Verlag.

[22] Robert Kegan (1994) *In over our heads: the mental demands of modern life*, Harvard University Press.

[23] Robert Kegan (2016) *An Everyone Culture. Becoming a Deliberately Developmental Organization*, Harvard Business Review Press.

[24] S. Kleene (1936) General recursive functions of natural numbers, *Mathematische Annalen* **112**(5):727–742.

[25] Ray Kurzweil (2005) *The Singularity Is Near*. Viking Press.

[26] Lia Kvirchishvili (2024). The Evolving Workforce: Technological Advancements and Their Impact on Employee Skills and Characteristics. In: Geibel, R.C., Machavariani, S. (eds) *Digital Management to Shape the Future. ISPC 2023. Springer Proceedings in Business and Economics*. Springer, Cham. pp. 81-96. https://doi.org/10.1007/978-3-031-66517-2_7

[27] André Leroi-Gourhan (1964) *Le Geste et la Parole. Technique et langage*, Albin Michel.

[28] André Leroi-Gourhan (1965) *Le Geste et la Parole. La Mémoire et les Rythmes*, Albin Michel.

[29] Michael Lynch (2010) Rate, molecular spectrum, and consequences of human mutation. *PNAS* **107**(3):961-968.

[30] John McCarthy (1960). Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I. *Communications of the ACM*, **3**(4):184–195.

[31] Warren McCulloch and Walter Pitts (1943) A Logical Calculus of the Ideas Immanent in Nervous Activity, *Bulletin of Mathematical Biophysics*, **5**:115-133.

[32] McKinsey Global Institute (2023) *The Economic Potential of Generative AI*.

[33] Thomas A. McMahon & John T. Bonner (1983) *On Size and Life*. Scientific American Library..

[34] National Bureau of Economic Research (NBER) (2023) *The Impact of AI on Labor Markets*.

[35] Dino Pedreschi, et al. (2025) Human-AI coevolution, *Artificial Intelligence*, Elsevier, **339**(C).

[36] E. Post (1936) Finite Combinatory Processes – Formulation 1, *Journal of Symbolic Logic*, **1**(3):103-105.

[37] Purpose and Means (2026) Origins and History of AI, *Google*.

[38] Sebastian Schleidgen, Orsolya Friedrich, and Johanna Seifert (2026) Rethinking human-technology relations. *Humanit Soc Sci Commun* **13**, 488. https://doi.org/10.1057/s41599-026-07050-0

[39] Knut Schmidt-Nielsen (1984) *Scaling: Why Is Animal Size So Important?* Cambridge University Press.

[40] Jaime Sevilla, et al. (2022) Compute Trends Across Three Eras of Machine Learning. *International Joint Conference on Neural Networks (IJCNN)*, pp. 1-8

[41] Claude E. Shannon (1938) A Symbolic Analysis of Relay and Switching Circuits, *Trans. AIEE.* **57**(12):713–723.

[42] Minkyu Shin, Jin Kim, Minkyung Kim (2021) Measuring Human Adaptation to AI in Decision Making: Application to Evaluate Changes after AlphaGo, Published Online: https://doi.org/10.48550/arXiv.2012.15035

[43] Gheorghe M. Ştefan, Aurel Paun, Andy Birnbaum, Virgil Bistriceanu (1984) DIALISP- A LISP Machine, *The 1984 ACM Symposium on LISP and Functional Programming* pp. 123–128. At: https://www.academia.edu/64552217/DIALISP_a_LISP_machine

[44] Dov Te'eni, et al. (2023) Reciprocal Human-Machine Learning: A Theory and an Instantiation for the Case of Message Classification, Published Online: https://doi.org/10.1287/mnsc.2022.03518

[45] Max Tegmark (2017) *Life 3.0: Being Human in the Age of Artificial Intelligence*. Penguin.

[46] Neil Thompson, et al. (2020) The Computational Limits of Deep Learning. MIT Sloan Research Paper. At: https://doi.org/10.48550/arXiv.2007.05558

[47] Alan M. Turing (1936-1937) On computable numbers with an application to the Eintscheidungsproblem, *Proceedings of the London Mathematical Society*, **42**(1):230-256 and a correction in **43**(6):544-546.

[48] Alan M. Turing (1950) Computing Machinery and Intelligence. *Mind*, **59**(236):433–460.

[49] Vladimir Ivanovich Vernadsky (1938) *Scientific thought as a planetary phenomenon*, translated from Russian by B.A.Starostin. Published Online: chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www.axelkra.us/wp-content/uploads/2021/08/cupdf.com_scientific-thought-as-a-planetary-phenomenon-v-i-vernadsky.pdf

[50] Ashish Vaswani, et al. (2017) Attention Is All You Need, *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems*,

[51] Pierre-François Verhulst (1838) Notice sur la loi que la population poursuit dans son accroissement. *Correspondance Mathématique et Physique* **10**:113-121.

[52] JOel G. Kingsolver, et al. (2001) The strength of phenotypic selection in natural populations. *American Naturalist* **157**(3):245-261.

[53] Bruce Walsh & Michael Lynch (2018) *Evolution and Selection of Quantitative Traits*. Oxford University Press.

[54] Geoffrey B. West, James H. Brown & Brian J. Enquist (1997) A general model for the origin of allometric scaling laws in biology. *Science* **276**(5309):122-126.

[55] Jonathan H. Westover (2026) The AI Implementation Gap in Higher Education: Navigating the Disconnect Between Technology Adoption, Policy Awareness, and Institutional Governance, Published Online: https://www.preprints.org/manuscript/202603.2295

[56] Shasha Yu, Fiona Carroll, Barry L. Bentley (2026) Rethinking AI Literacy Education in Higher Education: Bridging Risk Perception and Responsible Adoption, Published Online: https://doi.org/10.48550/arXiv.2603.29935

A Appendix: History leading to AI

The following list of contributors to the mythical-conceptual-technological line of computing & AI (originates partially from Google [37]).

Early Origins :

- 7th-6th century BC: Epimenides of Knossos, seer and philosopher-poet, is credited with the saying "Cretans are always liars", a sentence that raised the issue of **decision** for the first time.
- 384–322 BC: Aristotle developed foundational work for logic, and scientific methods.
- 150-100 BC: Antikythera mechanism is the oldest known example of an analogue computer used to predict astronomical positions and eclipses.
- 43BC-17/18AD: Ovid tells the story of Pygmalion, a Cypriot sculptor who sculpted a woman from ivory and alabaster, whom he later brought to life. Thhis story is part of the mythical history of computer science and AI.
- 1st-2nd century AD: Hero of Alexandria who is credited, among other things, for a machine that dispensed holy water in an Egyptian temple in an amount proportional to the coin that was inserted into the device.
- 999: Gerbert d'Aurillac (Pope Silvester II; c.946 – 12 May 1003), a symbolic ancestor of European clockmaking was a scholar whose scientific reputation caused later writers to credit him with the first horologe.
- 1623: Wilhelm Schickard, a German professor of Hebrew and astronomy, designed a calculating machine called the *Calculating Clock*. It was designed to perform addition, subtraction, multiplication, and division.
- 1642: Blaise Pascal designed and *implemented* the first the mechanical calculator.
- 1763: Bayes' Theorem is published, marking the beginning of statistical methods foundational to AI.
- 1773: Gottfried Wilhelm Leibniz designed and implemented the calculator named stepped reckoner.
- 1818: Mary Shelley's novel *Frankenstein* is an early example of science fiction.
- 1822: Charles Babbage's difference engine.

- 1843: Ada Lovelace wrote the first computer program for Babbage's engine, proving that machines could manipulate symbols.
- 1847, 1854: George Boole published his books on symbolic logic [3] [4].
- 1872: Samuel Butler's novel *Erewhon* is one of the first to explore ideas of AI.
- 1920: Karel Čapek's play *R.U.R.* introduces the word "robot," expanding AI concepts into popular culture.
- 1943: Warren McCulloch and Walter Pitts create a computational model for neural networks [31].

Golden Years of Technological Advances :

- 1900: David Hilbert presented, in his conference at the Paris mathematics congress, a first approach to the decision problem [20].
- 1928: David Hilbert & Wilhelm Ackermann published their *Principles of Mathematical Logic* where they included a clear statement of the *Entscheidungsproblem* (decision problem).
- 1931: Kurt Gödel publish its seminal paper on incompleteness [18].
- 1936: Alonzo Church [9], Stephen Kleene [24], Emil Post [36], Alan M. Turing [47] published their papers defining what means computation starting from the negative result of Gödel.
- 1937: Claude Shannon wrote his master's degree thesis, *A Symbolic Analysis of Relay and Switching Circuits*
- 1950: Alan Turing publishes "Computing Machinery and Intelligence" and proposes the Turing Test [48].
- 1951: Marvin Minsky & Dean Edmunds build SNARC, a network of 3,000 vacuum tubes simulating 40 neurons.
- 1955: John McCarthy coins the term "Artificial Intelligence" [30] and organises the Dartmouth Summer Research Project.
- 1959: Arthur L. Samuel demonstrates machines can learn from past errors using the game of checkers.
- 1961: "Unimate," the first industrial robot, is deployed on a General Motors assembly line.

- 1966: ELIZA, an early NLP program, is developed at MIT by Joseph Weizenbaum.
- 1972: Karen Spärck Jones devises Inverse Document Frequency (IDF), underlying most modern search engines.

AI Winter #1 :

- 1973: Scepticism about AI's capabilities leads to reduced funding.
- 1974: Ray Kurzweil's company introduces Optical Character Recognition (OCR) commercially.

Rejuvenation :

- 1975: CONS Machine, the first Lisp machine built in MIT Artificial Intelligence Laboratory in a project started in 1973 by Richard Greenblatt and Thomas Knight.
- 1981: *Symbolics, Inc.* was a spinoff from the MIT AI Lab introduced LM-2 a workstation based on MIT CADR architecture.
- 1983: *DIALISP* a Lisp machine designed at the Politehnica University of Bucharest and produced in small series at FEPER [43].
- 1980s: "Expert systems" are adopted by corporations around the world.
- 1986: David Rumelhart and James McClelland develop ideas around parallel processing and neural network models.

AI Winter #2 :

- 1987: The LISP machine market collapses because *Sun Microsystems* and *Apollo* could run Lisp nearly as fast as dedicated Lisp machines, but at a fraction of the cost.

Rebirth :

- 1995: A semi-autonomous vehicle navigates a 4,501 km journey.
- 1997: IBM's Deep Blue becomes the first computer to beat reigning world chess champion Garry Kasparov.

AI Enters the Home :

- 1999: Sony launches AIBO, a four-legged autonomous robot for consumers.
- 2002: iRobot releases Roomba, demonstrating practical household AI applications.

Breakthroughs and AI Becomes Mainstream :

- 2004: DARPA challenges stimulate advancements in autonomous vehicle technology.
- 2011: Introduction of AI-powered voice assistants like Siri, Google Now, and Cortana.
- 2012: Google's X lab demonstrates 70% better results in detecting and recognising cats from YouTube videos.
- 2016: Google's AlphaGo defeats Go champion Lee Se-dol, a landmark achievement in AI's strategic capabilities.
- 2017: Transformer architecture introduced by Vaswani et al. [50].
- 2018: Generative AI, featuring GANs and variational autoencoders, excels in producing realistic images, text, and music from vast datasets.
- 2019: OpenAI introduce GPT (Generative Pre-trained Transformer).
- 2021-2022: marks the transition from language models to models capable of generating images, code, audio.
- 2023: Microsoft Copilot introduces a generative artificial intelligence chatbot.