mar
asset

Letter: November, 2025

When Skeptics Push Geniuses

Around 1920, and for five years in a row, the world went through a revolution in thinking that began with the physicist Niels Bohr, who decided to question the deterministic standards of traditional physics. Until then, any student of physics relied on mathematics, with its proven equations, relationships of cause and effect, and the quest to explain the world in terms of 0s and 1s. Nothing in between those extremes was accepted, and every new theory was quickly discredited by the skeptics on duty.

The beauty of theoretical physics lies precisely in the idea that everything is possible - until proven otherwise. It is easy to picture Bohr, in a moment of simple observation, having an even simpler insight: what if nature does not follow deterministic patterns, but probabilistic ones? What if the act of observing changes the very outcome? What would be left of everything we had learned up to that point?

Quantum physics was born. From that moment on, centered at the Copenhagen Institute, Bohr became a magnet for talents interested in his revolutionary ideas. In just a few years, a generation of scientists, philosophers, and mathematicians completed what would become the greatest revolution in the modern history of physics and mathematics.

An initial group of five scientists - Niels Bohr, Werner Heisenberg, Erwin Schrödinger, Max Born, and Paul Dirac - one skeptic (Albert Einstein), and a mathematical prodigy (John von Neumann) collectively pushed the frontier of knowledge forward. That was one hundred years ago.

Among them, John von Neumann was one of the most emblematic figures. Born in Hungary - then one of the richest breeding grounds of mathematical geniuses in the world - he was a prodigy. At eight years old, he solved complex equations; at twelve, differential calculus. By 23, he was a chemical engineer and a PhD in pure mathematics.

In the second half of the 1920s, von Neumann wrote the mathematical equations consolidated in his seminal book Mathematical Foundations of Quantum Mechanics, which served as the basis for proving the principles of quantum physics.

In the 1930s, he emigrated to the United States and, at Princeton, worked with Einstein and Gödel. Over the following ten years, he broadened his interests and started small revolutions in everything he touched - game theory, hydrodynamics, meteorology, nuclear reactions, and, finally, the early days of modern computing.

During World War II, von Neumann was invited by Oppenheimer to be one of the leaders of the Manhattan Project, having at his disposal what could then be called infinite capital. The project aimed to develop the atomic bomb - the apex of applying the theories formulated by that group of scientists. The direct and nefarious consequences are well known. The indirect one, also linked to von Neumann, was the birth of modern computing.

The calculations required to design and test nuclear weapons involved nonlinear differential equations, independent variables, probabilistic theories about particle behavior and wave propagation - based on chaos theory - as well as the use of game theory for war strategies, all topics mastered by von Neumann. This complexity forced the development of machines able to automate the calculations.

Thus emerged one of von Neumann's main theories: Automata, which dealt with the automation of functions by means of mechanical machines. It was the embryo of the "Digital Brain" - a constant flow of data between memory and processing.

With virtually unlimited resources and access to the first large-scale electronic computer, ENIAC, advances came in quick succession. The needs

of the Manhattan Project spurred the development of computing technologies still used today.

Current technology behind artificial intelligence stems from many of von Neumann's ideas. One of the major advances was overcoming the so-called "von Neumann bottleneck" - the physical speed limit in data transfer between processing and memory, now mitigated by the GPU model and parallel processing.

Einstein was the great skeptic of quantum physics. He ridiculed theories about particles that could exist in more than one place at the same time, in clear defiance of his own theory of time and space. His disbelief, however, ended up pushing Bohr and von Neumann to mathematically prove the improbable.

It was the skeptics pushing the geniuses.

Einstein had an estimated IQ of 170. Von Neumann's was around 200.

A century later, history repeats itself. It was long believed that computing should be linear and that the secret lay only in speeding up processing. Attempts to create parallel computations based on matrix algebra and probabilistic analyses were discredited for decades. Machine learning was relegated to academic curiosity - or ridicule.

That was, until a group of scientists, led by Yann LeCun and Geoffrey Hinton, challenged orthodoxy and dove into the development of deep learning. The belief was simple and powerful: the more data you feed an algorithm, the higher the probability that it will learn patterns.

Like the visionaries of the 1920s, they too were discredited - until the results began to appear. Today, we are living through the artificial intelligence revolution, largely based on the work of three figures: Hinton (University of Toronto and Google), LeCun (Meta AI), and Demis Hassabis (Google DeepMind).

Hassabis is the emblematic figure of this new era. Born in Britain in 1975, he became a chess master at 13 and the second best player in the world. He graduated in Computer Science at Cambridge with top marks at 21, but his curiosity led him to cognitive neuroscience. He obtained his PhD at 33, studying the neural mechanisms of memory, and the following year

founded DeepMind, with the goal of “solving intelligence” by means of computation.

In 2014, Google acquired DeepMind. Two years later, his AlphaGo project defeated the world’s best Go player, using a computer trained with von Neumann’s game theory principles and neural networks. Today, Hassabis leads Google’s artificial intelligence projects.

From roughly the same generation came Sam Altman (OpenAI), Dario Amodei (Anthropic), and the omnipresent Elon Musk (X.ai), completing the group of 21st-century prodigies who are transforming computing. As in the Manhattan Project, they have virtually unlimited capital at their disposal. If von Neumann unified physics, mathematics, and computing a hundred years ago, this group is now trying to merge neuroscience and computing in the search for intelligence.

Just as quantum physics redefined our understanding of matter, artificial intelligence is redefining the infrastructure of information - and facing, in its own way, a new universe of skeptics.

Demis Hassabis has an estimated (not measured) IQ of around 170.

Crossing the frontier of genius

In the same year Demis Hassabis was born, 1975, Terence Tao was born in Australia, now a professor at UCLA. At seven, he was already solving calculus problems; at twelve, he scored 760 on the math section of the SAT; at 24, he became a university professor. His career has been devoted to mathematics - prime numbers, differential equations, and probability theory.

Seven years later, in 1982, Christopher Hirata was born in the United States. At 13, he won a gold medal at the International Physics Olympiad, the youngest ever. At 14, he entered Caltech; at 16, he was working with NASA scientists on a project for colonizing Mars; at 22, he finished his PhD in cosmology at Princeton.

What do the two have in common? A measured IQ of around 225, according to the Mensa/Stanford–Binet methodology - the highest ever recorded.

Standard IQ tests have an upper scale of 180; any value above that depends on statistical extrapolation. Tao and Hirata literally **broke the scale**.

Someone is considered "superintelligent" with an IQ above 160. The metric has limitations, but it helps illustrate the current stage of evolution of artificial intelligence.

ChatGPT-1, released in 2018, was trained with 120 million parameters and would have the equivalent IQ of an eight-year-old child (in a loose analogy, since measuring the IQ of algorithms is conceptually debatable).

GPT-3, with 1.5 billion parameters and ten thousand times more computing power, would reach something close to an IQ of 120 - the IQ of an average university student.

GPT-5, recently launched, was trained with 1.5 trillion parameters and one hundred times more capacity than GPT-3. In comparative terms, it has already reached the level of **Einstein**, approaching **von Neumann**.

Today, any user of a frontier LLM - such as GPT-5 - interacts with an intelligence that contains aspects found in the minds of the greatest geniuses in history.

Up to GPT-4, AI was just a support tool - useful for drafts, research, and simple tasks. A good intern, at best. GPT-5 changed that: it already thinks, synthesizes, and learns with the fluidity of a specialist.

Projections suggest that **GPT-6** will be trained with **ten times more computing capacity** and **more than ten trillion parameters**, in addition to new post-training techniques, real-time image integration, and advances in matrix computation. Experts estimate that its "intelligence" will double relative to the current model - a **hypothetical IQ of 300**.

Whether this will amount to an Artificial General Intelligence (AGI) or some form of superintelligence, we still do not know. But it seems clear that we are about to meet the most intelligent entity ever created by humankind.

Your own personal genius

But what good is all this intelligence if the interface is not good? What would be the point of talking to Einstein if you did not speak the same language?

For decades, anyone who wanted to interact with computing power had to master a programming language. Between humans and computers there was a gatekeeper who only handed over the key to the initiated. That is why software engineers and programmers dominated technology for so long.

We, ordinary users, were hostages to interfaces created by third parties - or to the effort of learning a new "language." And these languages followed one another at a frenetic pace: Assembly, Fortran, Cobol, Pascal, C, Python. Imagine having to relearn how to communicate every few years..

LLMs changed that. They eliminated the gatekeeper. Now anyone can talk to the machine in natural language - and in their own language.

We finally have an interface of superintelligence that communicates effortlessly. And in it, something almost magical happens, with echoes of quantum physics: while an LLM is being trained, its tokens - pieces of words - float in a cloud of computation, connected by probability matrices. It is a chaotic dance in search of order, similar to the quantum wave function. Just like particles that need to be observed in order to "collapse" into a defined state, the prompt organizes the chaos and generates a coherent answer - in a process that not even its creators fully understand.

The great products and services of recent decades - iPod, iPhone, Netflix, Spotify, Instagram, WhatsApp, Google Search, Google Maps - share common traits: ease of use, removal of bottlenecks, an element of surprise, learning about the user, and a price that is low relative to the value delivered. The result is almost addictive retention.

LLMs, especially ChatGPT, have all these characteristics. **More than 2.5 billion prompts** are written per day, and the retention rate is around 60% of the user base.

The next step is to measure how quickly this scales.

Is AI profitable?

Since the launch of GPT-5 - initially met with frustration - the model has evolved in an impressive way. OpenAI has announced dozens of integrations and partnerships, new image-generation versions, and a clearer path to monetization.

Within a few weeks, it entered into computing-supply agreements with chip, memory, and datacenter companies. The sequence of announcements moved markets significantly, but also raised questions about cross-references among partners.

Almost at the same time, two events ilustrated where the AI business model is headed. The **OpenAI Dev Day**, in October, presented ChatGPT as an "intelligence portal," integrating:

- **e-commerce tools (PayPal, Salesforce, Stripe, Shopify, Booking, Walmart), enabling purchases directly in the interface;**
- **corporate platforms (Slack, Canva, Figma, DocuSign, and Codex);**
- **and the new Atlas browser, a potential competitor to Google.**

OpenAI reported having reached **800 million users**, an increase of 100 million in a month, and projected an addressable market of **8 billion people** (or 9 billion smartphones).

Its revenue model rests on four pillars:

1. **Subscription plans;**
2. **E-commerce take rate;**
3. **Integration commissions from other platforms;**
4. **Advertising and search revenues.**

A fifth pillar, still incipient, is the **patents market**. If LLMs are great problem solvers, the leap in computing capacity will open space for a corporate "factory" of patents. Today there are about **200 thousand patents related to AI**, versus **20 thousand related to cancer**. JAMA Oncology estimates that the global cost of cancer will reach **US$ 25 trillion** over the

next 30 years - giving a sense of the potential profit pools in areas such as nuclear fusion and biotechnology.

Weeks later, Oracle announced the full integration of its main product - Oracle Database - with any LLM model, branding it AI Database. The significance is considerable: the main resistance from companies has always been the risk of exposing proprietary data. And one of the last great frontiers of data for training models is private data. Oracle seems to have solved this at scale, with an architecture built for security..

The practical consequence is straightforward: corporations will be able to converse with all their historical data in natural language, applying LLM productivity tools. It is AI leaving the slides and entering measurable use cases.

The movement is already visible in earnings reports: 287 S&P 500 companies mentioned AI in 2Q25. Examples:

- **Anthropic & Novo Nordisk: clinical reports prepared in 5 minutes, versus up to 9 weeks before, with human review and identical results.**
- **C.H. Robinson: 50% of orders automated by AI; 75% on short routes, with a 35% gain in productivity.**
- **Citibank: 180 thousand employees using AI; 100 thousand hours saved; intelligent agents in credit cards.**
- **J.P. Morgan: US$ 2 billion invested in AI; an equivalent return in savings; 20% more software productivity; 450 use cases identified, a goal of 1,000 in one year, and an additional US$ 1.5 billion in savings.**

These CEO statements tend to accelerate corporate adoption: the more evidence, the greater the demand.

In education, California is leading:

- **OpenAI + Caltech: ChatGPT for 500 thousand students;**
- **Google + community colleges: Gemini for 2.1 million students and teachers in the state;**
- **Microsoft: US$ 4 billion for AI programs in schools and universities.**

It does not seem far off that every educational institution will have its own official AI tool. In the U.S. alone, there are 70 million students; globally, 1.5 billion. The user base of the four main models already reaches 1.2 billion.

Winner takes all?

A recurring question was whether there would be room for more than one frontier LLM. Today, the main ones - disregarding open source - are **GPT (OpenAI)**, **Gemini (Google)**, **Grok (xAI)**, and **Claude (Anthropic)**, with a few smaller models running on the sidelines.

What is starting to emerge is a scenario of **segmentation, not a single dominant winner**:

- **ChatGPT seeks to be everything to everyone, with a stronger focus on personal use and on an individual assistant model, including inside companies.**
- **Gemini is focused on enhancing Google's internal ecosystem and on solving highly complex problems - from cellular diagnostics to advances in nuclear fusion.**
- **Grok, from xAI, explores real-world simulations, leveraging Tesla's massive real-time image database.**
- **Claude (Anthropic) is becoming the platform of choice for companies that depend on intensive programming. Although not exclusively targeted at programmers, adoption by this group - who usually sit at the innovation frontier - has been so strong that Anthropic has posted the fastest revenue acceleration in modern history, going from zero to US$ 9 billion per year in two years. (If you have not yet tried Claude, ask it to build a complex Excel spreadsheet using messy data. And hold your jaw!)**

What once seemed like **redundancy or a waste** of capital and computing capacity is proving to be complementary. Each model occupies its own niche, and competition - not only accelerating the mutual innovation curve - is giving way to functional specialization - closer to an ecosystem than to a winner-takes-all dynamic.

(Seemingly) insurmountable obstacles?

The three main obstacles to the global expansion of computing and the diffusion of AI platforms are capital, **energy**, and **supply chain**.

1. Capital.

So far, the advance of AI has been funded by big techs' cash flow and by venture capital and private equity money. These resources are expensive and limited - they require high returns and come with significant risk.

The new development is the growing use of the **credit market**. **NVIDIA** has announced a GPU leasing model linked to long-term contracts with technology companies. This structure allows securitization of the contracts, with the GPUs themselves as collateral - which reduces risk and enables financing costs far lower than those required by shareholders[1].

Without the combined use of equity and structured debt, it would be impossible to raise the US$ 2 to 3 trillion that **Morgan Stanley** estimates will be needed over the next five years to rebuild the global computing infrastructure.

The trend is for new structures of this kind to emerge, loosening the knot created by cross-partnerships along the value chain - still with little clarity about the ultimate source of capital. In **Blackstone**'s latest earnings, for instance, AI infrastructure financing was repeatedly cited as a strategic focus. We are, indeed, entering the **project-finance phase of AI**.

2. Energy.

The energy bottleneck is even more complex. Recent projections from Morgan Stanley point to a nearly 50 GW gap between the expected demand for computing and the available supply of energy in the United States by 2028 - the equivalent of tens of billions of dollars in new investments.

1 In the last month, two large transactions were structured under this model - one with xAI (Elon Musk's platform and owner of the Grok LLM) and another with Meta - in which business risk is split between equity and debt holders, with contracts and chips as collateral.

The physical concentration of datacenters on U.S. soil has created an intense race for uncontracted power, megawatt by megawatt, pushing prices up and generating political noise over its impact on costs for the rest of the economy.

The positive side is that this scarcity is driving a new **Schumpeterian** cycle of energy innovation. **Nuclear fusion companies, mini-reactors, batteries, modular thermal plants, and fuel cells** have been receiving substantial venture capital funding. Former bitcoin miners are being converted into energy suppliers for datacenters, leveraging long-term, low-cost contracts signed before the AI boom.

Among the emblematic projects is **Helion Energy**, backed by Sam Altman (OpenAI), focused on nuclear fusion. One of humankind's greatest scientific challenges may be solved precisely with advanced AI models - based on chaos theory and the engineering of new materials.

The **Trump** administration has also stepped in. Power supply for AI has been classified as a matter of **national security**, with **US$ 80 billion** earmarked for building new nuclear reactors and measures to speed up grid interconnection for new sources. This is, in part, a response to China's advances in generation capacity.

The solution to the energy imbalance will involve **new technologies and greater geographical dispersion of computing**, reducing dependence on the U.S. - especially in the inference phase, which will demand much more energy and processing.

Google has announced **US$ 10 billion** in investments in India; **Microsoft**, multi-billion-dollar contracts with European neoclouds powered by clean energy; and the **Stargate project** - a US$ 500 billion consortium led by SoftBank and Abu Dhabi's sovereign fund - plans to build global datacenters. The move is coherent: Japan is restarting nuclear plants, and **Abu Dhabi** combines gas reserves with **4 GW** of solar generation.

3. Supply chain.

Even with capital and energy, the final bottleneck is **equipment** Semiconductor production remains constrained by the dependence on

rare earths and the concentration of manufacturing in Taiwan. Turbines, transformers, network components, memory, and datacenter infrastructure are also in short supply. **Backlogs** already span **2 to 4 years**.

Moreover, the installed base of traditional datacenters still runs on **CPUs** and **x86 architecture**, requiring migration to networks compatible with the parallel processing of **GPUs**. The **inference** phase - when AI begins operating at scale - will demand another major **capex** cycle.

A fourth potential bottleneck is data to train new model versions. Part of this bottleneck has been addressed with post-training tools and reinforcement learning, processing of real-time data streams, and the creation of synthetic data by the models themselves.

If **OpenAI** is right and LLMs are used by every smartphone and company in the world, the current computing capacity will be only the starting point for a new global industrial cycle.

Market skepticism

Over the next 6 to 12 months, we expect a series of relevant announcements that should **unlock the entire AI value chain** – from advances in applicability to new revenue models and solutions for the bottlenecks in capital, energy, and equipment. This process should make the **transformational potential of AI** more evident, with direct implications for the prices of assets exposed to it.

Much has been said about a possible **AI bubble**, with parallels to the final phase of the internet and telecom boom between 1999 and 2000. The market, indeed, has shown signs of excess in a few peripheral niches:

- the return of SPACs;
- IPOs of companies that store cryptoassets and are worth multiples of their own reserves;
- quantum computing firms with no revenue but billion-dollar market caps;
- the reappearance of meme stocks;
- and the speculative craze around 1-day options.

These are typical symptoms of an environment in search of the “new new thing” and quick gains.

But **bubbles rarely start with euphoria**. The trigger usually comes from an acceleration in credit, which supports system-wide leverage and masks diminishing marginal returns.

Between 1997 and 2000, telecom companies led the construction of the global fiber-optic infrastructure that enabled the internet - but without associated revenues until the fiber was “lit.” The network still exists, but almost all original shareholders disappeared. A large part of these assets now belongs to Google, which controls the largest private fiber-optic mesh in the world. Between 2005 and 2008, mortgage credit was the fuel, culminating in the global financial crisis.

The current cycle is different. **There is still little evidence of companies “selling air.”** Most AI investments come from big techs’ own cash generation and from venture capital and private equity funds. The direct involvement of the U.S. government reinforces the structural nature of this movement, and sector indebtedness is still in its early stages.

Furthermore, **earnings revisions have tracked the rise in stock prices.** At the end of 2022, **NVIDIA’s** estimated earnings per share for 2026 were below US$ 1 (adjusted for the split). Today, they are headed toward **US$ 8 to 10.** The company’s market cap has risen from **US$ 400 billion to US$ 5 trillion**, and the stock trades at **25x the estimated 2026 earnings** - a high multiple, but far from the irrationality of past bubbles.

At the turn of 2000, Cisco, then the most valuable company in the world, traded at **120x earnings,** with projections that “only” doubled between 1997 and the burst of the bubble. Actual earnings delivered in 2001 were **half of what had been forecasted**.

Today, the picture is different. There are localized imbalances of supply and demand, but no clear signs of a blow-off top. Even so, we see two potential **sources of fragility**:

1. **U.S. retail investors, who already account for about 30% of stock-market volume, operating with leverage of close to US$ 1 trillion - significant, but less than 1% of total household wealth.**

2. **The beginning of leverage in the AI infrastructure cycle. If Morgan Stanley's projected US$ 2 to 4 trillion in capex is confirmed, a large share will be financed through traditional and structured debt.**

For now, this places us closer to 1997 than to 2000 – midway through the acceleration phase. These phases are spectacular by nature: they bring almost daily innovations, a strong concentration of leadership, and a temporary mismatch between financial projections and concrete results.

We will keep a close eye on this transition, in one of the most fascinating periods of the modern era - when, once again, skeptics push geniuses, and perhaps push all of humanity toward the next technological leap.

But, recalling Niels Bohr's line:

> *"Prediction is very difficult,*
> *especially about the future."*

MAR Asset

Bibliographic References:

Bhattacharya, Ananyo. *The Man from the Future: The Visionary Life of John von Neumann*. 2022.

Hao, Karen. *Empire of AI*. 2024.

Tegmark, Max. *Life 3.0: Being Human in the Age of Artificial Intelligence*. 2017.

Mitchell, Melanie. *Artificial Intelligence: A Guide for Thinking Humans*. 2019.

Waldrop, M. Mitchell. *The Dream Machine: J.C.R. Licklider and the Revolution That Made Computing Personal*. 2001.

Olson, Parmy. *Supremacy: AI, ChatGPT and the Race for Artificial Intelligence*. 2023.

Kelly, Kevin. *Out of Control: The New Biology of Machines, Social Systems, and the Economic World*. 1994.

Kelly, Kevin. *The Inevitable: Understanding the 12 Technological Forces That Will Shape Our Future*. 2016.

Kim, Tae. *The Nvidia Way*. 2024.

Becker, Adam. *What Is Real?: The Unfinished Quest for the Meaning of Quantum Physics*. 2018.

Labatut, Benjamin. *Maniac: The Brilliant Mind of John von Neumann*. 2023.

Labatut, Benjamin. *When We Cease to Understand the World*. 2020.

Isaacson, Walter. *Einstein: His Life and Universe*. 2007.

Maxwell, Andrew. *Demis Hassabis: The Mind Behind DeepMind*. 2024.

Mallaby, Sebastian. *The Infinity Machine: How OpenAI Changed Everything*. 2024.

Metz, Cade. *Genius Makers: The Mavericks Who Brought AI to Google, Facebook, and the World*. 2021.

Suleyman, Mustafa. *The Coming Wave: Technology, Power, and the Twenty-First Century's Greatest Dilemma*. 2023.

Kissinger, Henry; Schmidt, Eric; Huttenlocher, Daniel. *The Age of AI: And Our Human Future*. 2021.

Mollick, Ethan. *Co-Intelligence: Living and Working with AI*. 2024.

mar
asset

Investor Relations
Igor Galvão
55 21 99462 3359
contato@marasset.com.br

rio de janeiro – rj – av. ataulfo de paiva 1165 / 801, leblon – 22440 034
marasset.com.br