The ACCELERATION of METEORAIT
AI as an Ecological Factor, Volume II

WHY - Why this book

This book exists because METEORAIT is not standing still. In the first book, I described **AI as** an **ecological factor:** not a tool, not a simple technological revolution, but **an environmental event.** A phenomenon that does not ask for permission, has no intention, but fundamentally changes the conditions in which we think, learn, work, decide and organize ourselves.

This second book no longer just describes the impact. It describes **the acceleration.** The impact can be watched. The acceleration must be felt **in time.**

If the technological history that began with the taming of fire were compressed into a **single year,** the GPT era would begin in the last **ten minutes.** But those ten minutes don't flow evenly. In the last minute, AI rapidly moves from model to flow, from response to process, from apparent tool to true environmental infrastructure.

The meteorite doesn't fall faster because it wants to. It falls faster because it has entered the gravitational field that attracts it. In the case of AI, this field is humanity itself. The market attracts it. War attracts it. Fear attracts it. Productivity attracts it. Infrastructure attracts it. States attract it. Companies attract it. Users attract it.

This is the central idea of the book: AI doesn’t accelerate on its own. It is **accelerated by humanity** adopting it. We are its gravity. This is the hard part to accept: we don't just approach AI. We pull it towards us. Every delegation of a mental task to the machine acts as a force of attraction.

Artificial language has become a medium of flux. It flows continuously through our digital pipelines. And what flows continuously is no longer just a tool. It becomes climate.

And with autonomous agents, the change becomes even more profound. The chatbot responds. The agent continues. It reads, plans, checks, corrects, and executes. It no longer just speeds up isolated responses, but the entire processes in which we are engaged.

That's why this book is not about the effects of AI on the individual, the group, the community, or the state. That mapping belongs in the first book. Here the perspective changes. There are only two major entities left in the arena: the AI and humanity. The meteorite and the field that attracts it.

This book describes the last ten minutes. With an emphasis on the last minute. **The present.** It examines a **historic regime change:** the moment when a technology is no longer just used by humans, but begins to alter the very pace, scale, and structure of human processes. The new ecosystem. And the denser the atmosphere, the smaller the impact may be.

METEORAIT describes the object.
ACCELERATION describes the fall.
MEG/MEG2 describes the atmosphere.

Chapter 1
Minute 10 - GPT-1 and the beginning of the GPT era

Year **2018.** OpenAI publishes the technical paper "*Improving Language Understanding by Generative Pre-Training*". The name will later become simple: GPT-1 (*Generative Pre-trained Transformer*).

At that moment, for the general public, almost nothing happens. There is no visible revolution. There is no palpable social shock. There are no schools panicking about "what's homework?" There are no companies urgently rewriting their internal processes. There are no states discussing strategic risk at the highest level.

There is just a subtle change in software architecture and a new research direction. This is how systemic changes sometimes begin: without noise. Historical changes that make no noise when they begin, but only make noise when they become the environment of life.

GPT-1 wasn't the first artificial intelligence system. It wasn't even the first language model. Before it, there had been neural networks, pattern and image recognition (a field in which, in 1997, I was writing my own undergraduate thesis focused on artificial neural networks and fuzzy logic), or decades of algorithms programmed to follow rigid rules. History hadn't begun yesterday, and artificial intelligence wasn't a recent fad.

But GPT-1 has established a technical recipe that will prove decisive: pre-training on massive amounts of text and then fine-tuning to different tasks.

Language itself was beginning to be treated as a general training medium. The model was no longer constructed separately for each micro-task. It was no longer given only narrow logical rules. It was no longer forced to remain locked in a small, closed and fully specified domain. It was exposed

directly to written text. And text is the long, deep and time-honored trace of human knowledge.

This is how minute 10 begins. GPT-1 would not have been spectacular by today's standards. By today's standards, it was downright tiny, barely visible in the vast expanse of our culture. It was far from us, from a humanity increasingly fragmented by online social networks. History doesn't always begin with a giant object. It begins with direction.

GPT-1 showed that a model could learn general representations directly from language and then successfully transfer them to different tasks. This was the real mutation.
Until then, most AI systems were built as specialized mechanisms. They did one thing. They responded within a predefined framework. They solved a strictly delimited problem. GPT-1 was not yet general in the public sense of the word, but it contained the germ of generalization: the same linguistic basis could be adapted to several different forms of understanding.

The meteor wasn't visible in the sky yet. It was just the first bright dot, hard to identify. Humanity's telescopes were pointed elsewhere, at the clouds of dust and noise of social media. The problem with the first dot is that it never seems urgent...

GPT-1 grew from a theoretical root that emerged in 2017: *Transformer*. The pioneering work "*Attention Is All You Need*" had radically changed the way models could process language sequences. It was abandoning the way computers had previously read text word by word, from left to right, as if through a narrow tunnel. The new model put the concept of attention squarely at the center. Instead of the system moving through text linearly, it could instantly weigh the complex relationships between all parts of a sequence, regardless of the distance between them.

Attention was not consciousness. It was not human-like understanding. It was not intention. It was just an extremely efficient mathematical infrastructure for correlations and

relationships. But history does not require intention to produce effects.

But it is precisely this clean infrastructure that will enable everything that comes next: scale, concurrent processing, extended context, knowledge transfer, and then fluent generation. At minute 10, none of this was obvious to the general public. It was only visible to the very few who could read the technical signal correctly.

Great history always has such moments of initial silence. A technology first appears as a simple detail for specialists. Then it becomes a product. Then it becomes infrastructure. Then it becomes environment.

Fire was not civilization at first. It was a preserved flame. Writing was not a library at first. It was a sign. The engine was not globalization at first. It was energy set in motion.

GPT-1 wasn't ChatGPT at first. It was a recipe, and perhaps more important than the first fire. Because some recipes can completely change the ecosystem of the human species.

Impact

Minute 10 doesn't change humanity on the surface yet. It decisively changes the direction inside the research labs.

Until then, the dominant question in computer science could be formulated simply: *how do we make a machine solve this specific task?*

After GPT-1, the question begins to change: *how do we build a model that learns from language enough that it can later be adapted to different tasks?*

The difference is huge. The first question produces tools. The second produces environment. And the environment is not used. The environment is inhabited.

A tool is made for a single action. The environment changes the conditions of all our actions.

At this point, AI was not yet a visible ecological factor. But its trajectory was already set. If language can become a terrain for statistical “learning”, then all human activities that depend on language may one day become a terrain for machines. And almost all of our activities are structurally dependent on language.

Commerce depends on language. Education depends on language. Law depends on language. Politics depends on language. Medicine depends on language. Programming depends on language. Administration depends on language. War itself depends on language: orders, reports, coordination, recognition, classification, and decision.

That's why minute 10 is discreet, but not small. It doesn't change the world immediately. But it does change the nature of the object that is about to enter the world.

Acceleration

At minute 10, the acceleration force is still hidden.

It does not come from user pressure, market needs, agent actions, or the logic of war. It comes from the simple fact that the new software architecture allows for exponential scaling.

This is the major difference from most previous inventions. A classical tool improves through manual refinement. A language model of this type improves directly through scale, data volume, computing power, and skill transfer.

When the right mechanism meets the right infrastructure, acceleration is no longer just about inspiration. It becomes an industrial process. And from that moment on, speed is no longer just about brilliant people. It's about machines, money, data, and energy.

In 2018, this industrialization was still in its infancy. But all the conditions were already in place: data centers, high-

performance chips, the global internet, huge volumes of written text, capital, silent competition, and international research.

The meteorite wasn't yet being violently pulled towards us. But it was clearly on course.

People of the minute

Minute 10 does not belong to one person.

The technical root of this era belongs to the team that created the Transformer architecture: **A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin.** They were the ones who put the mathematical relations of attention at the center of the architecture that would underpin the entire GPT era.

The GPT-1 project belongs to the OpenAI team: **A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever.** They demonstrated that generative pre-training on text, followed by task-specific adaptation, can produce real transfer of competence on complex language problems.

In **Minute 10,** through their work, they permanently changed the direction of the METEORAIT's fall. Nothing was heard yet. But **the trajectory had already been changed.**

Chapter 2
Minute 9 - GPT-2 and the first warning

Year **2019.** GPT-2. This time, the bright point in the sky becomes visible to the naked eye. If GPT-1 had shown the way, GPT-2 demonstrated what happens when direction is scaled up to an industrial scale. The model reaches the threshold of 1.5 billion parameters ("digital synapses") established during the training phase) and is exposed to millions of web pages.

The result is no longer just a discreet laboratory experiment, useful for passing isolated linguistic tests. It becomes a system capable of continuing complex texts, preserving the author's tone, imitating journalistic styles, and producing fragments of striking coherence.

At this level, the question is no longer just a purely technical one. The question becomes social: *what do we do with a system capable of writing?* Because a writing machine doesn't just produce text. It produces the appearance of mind.

Not to calculate. Not to classify data. Not to search for information on the internet. But to write.

GPT-2 does not think in a biological sense. It does not understand the world behind words. It has no intentions, desires, or consciousness. But it produces language with a fluidity and continuity that flawlessly mimics human expression.

And for our species, throughout history, language has been the ultimate mark of thought. When a machine quickly calculates numbers, humans see it as a mere tool. When a machine writes coherent essays, humans tend to see it as an author. This is where the confusion begins. Not in the code, but in instinct.

This is the psychological mutation of Minute 9.

GPT-2 didn't change the world through mass adoption. In 2019, most people hadn't heard of it and were spending their time on traditional social media, caught up in the noise of digital microblogging platforms. But the model has radically changed the tone of the conversation among creators.

OpenAI is making an atypical decision for the scientific community: it is refusing to immediately release the full model. The release is being phased in over several months, starting with much smaller versions. The official reason given is the huge risk of abuse: the automatic and cheap generation of fake text, the automation of manipulation campaigns, and the flooding of the internet with coherent spam.

The moment is historic. Not because GPT-2 was an apocalyptic intelligence. It wasn't.

But because it opens up a fundamental ethical dilemma: *when a language model becomes powerful enough, is its publication still a neutral act?*

Until then, research had had the reflex of total openness. You publish the work, release the code on development platforms, the community checks, reproduces and improves. This is how science progresses.

But GPT-2 introduced a profound hesitation. Some people understood that a huge flow of synthetic language released without filters into an already fragile information environment could cause irreparable damage. The text did not need to be true. It was enough that it was believable.

Minute 9 is the minute when generative language stops being just a technical promise and becomes an ecological alarm signal.

Impact

Minute 9 permanently changes the status of the automatically generated text.

Up until this point, computer-generated texts were either rigid or easily identifiable as artificial. But GPT-2 has

shown that the simple task of predicting the next word in a sentence, when done on a large enough scale, forces the model to indirectly assimilate complex structures: grammatical rules, argumentative styles, emotional nuances, and logical associations about the real world.

We are not dealing with a human-type understanding. It is just an extremely well-tuned statistical simulation, but enough to cause a deep hesitation in the reader's mind.

Here begins a historical fissure. If high-quality text can be produced without the direct intervention of a human hand, then part of the historical authority of writing is definitively detached from the author. Writing had been, for millennia, the unmistakable trace of a conscious mind. From minute 9, writing can also be the trace of a cold computational system.

There is a risk of saturating the information environment with coherent, cheap and scalable noise. Incoherent noise tires. Coherent noise convinces. The Internet had already prepared the infrastructure: instant distribution, virality and polarization. GPT-2 brought the missing ingredient: the ability to generate linguistic content adapted on demand, in huge volumes, ready to be injected into the network.

Acceleration

At minute 9, the acceleration becomes visible through the increase in the physical size of the models.

From the discrete GPT-1 experiment, it moves to a model of 1.5 billion digital synapses, trained on a massive body of text extracted from positively rated web pages by users. The leap is not just numerical; it is a psychological shock to the industry.

Researchers, journalists, and regulators understand that generative models will not remain just niche academic tools. But the acceleration is being temporarily held in check by this phased rollout. OpenAI is delaying the full opening of the doors, giving society time to assess the impact.

People of the minute

Minute 9 belongs to the technical team behind GPT-2: **A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever.** They demonstrated that statistical language prediction, when given enough parameters, develops complex behaviors that go far beyond the initial formal task.

It also belongs to the group of researchers who conceptualized social risk and designed the responsible release strategy in the paper "*Release Strategies and the Social Impacts of Language Models*": **I. Solaiman, M. Brundage, J. Clark, A. Askell, A. Herbert- Voss, J. Wu, A. Radford, G. Krueger, JW Kim, S. Kreps, M. McCain, A. Newhouse, J. Blazakis, K. McGuffie, and J. Wang.**

They showed that launching an AI system is no longer just a neutral technical act, but a social environmental issue.

At **minute 9,** someone hesitated before opening the door. It was a good sign. And a sign that the door existed.

Chapter 3
Minute 8 - GPT-3 and the illusion of knowledge

Year **2020.** OpenAI publishes the pioneering work "*Language Models have Few-Shot Learners*". GPT-3 appears.

If GPT-1 set the direction, and GPT-2 issued the first warning, GPT-3 has brutally changed the scale of proportions. And when the scale changes enough, the way the human mind interprets the phenomenon also changes.

The model reaches the colossal threshold of 175 billion parameters. This huge figure produces a crucial mutation in public perception: the huge size of the system generates an overwhelming appearance of general competence.

GPT-3 could generate responses so fluid, varied, and well-adapted to the context that for most people the difference between *"the machine knows"* and *"the machine perfectly mimics knowledge"* was permanently erased.

Until GPT-2, the question was whether a model could generate coherent text over the course of a few paragraphs. With GPT-3, the question becomes: *can a model perform new tasks just through conversation, without being retrained from scratch?*

The answer is yes. The model exhibits an exceptional ability to adapt in flight, technically called "few-shot learning". You no longer need to be an engineer to reprogram its circuits or modify its database. It is enough to describe the task in natural language and give it a few quick examples so that it understands the pattern and continues solving.

The usual instruction now became the control interface. This was more than just convenience. It was a reversal: human language began to command the machine without going through code.

Humanity had lived for thousands of years in a world where natural language was used exclusively to coordinate other humans. You order a human, teach a child, convince a voter, ask an expert. With GPT-3, human language begins to function as a programming code for computer systems.

Machines no longer required rigid interfaces, physical buttons, complex menus, or strict lines of code. They became addressable through ordinary words, because our language had been statistically modeled deeply enough to be transformed into a universal control surface.

Here an extremely dangerous illusion sets in: the illusion that easy access to an automatically generated answer is equivalent to possessing knowledge.

False Cognitive Power Transfer (FCPT) appears directly. The person received the sensation of power without having gone through the effort that produces the power. GPT-3 explains, summarizes, translates, argues and imitates styles with astonishing speed. Sometimes it invents facts with complete certainty, but the linguistic form is so convincing that it deactivates the user's vigilance.

Humans are vulnerable to aesthetics. A fluent sentence feels true; a fluent explanation feels intelligent; a text delivered on the spot feels valuable. GPT-3 didn't create this vulnerability. It found it. And, unintentionally, it began to exploit it.

Impact

Minute 8 structurally alters the relationship between the human mind and the text.

The text stops being just a final product that we read and becomes an active work surface. When the system seems to understand the context, the human changes his natural behavior. He gives the machine more trust, more data, more small decisions and more authority.

Not for lack of intelligence, but for reasons of pure efficiency.

But cognitive support, when it is too accessible and comfortable, tends to replace the actual mental effort. The user receives the conclusion ready-made, but skips the fundamental steps of the process: difficult documentation, critical analysis of sources, formulation of their own arguments, and correction of mistakes.

The form is delivered, but the path is lost. And the path is precisely the space in which the real competence of a mind is formed and consolidated. When the path disappears, the result remains. But the man who receives it is no longer the same man who would have built it.

Minute 8 is not a brutal collapse of our cognitive civilization. It is a seduction through comfort.

Acceleration

At minute 8, the acceleration force moves completely into the area of industrial logic.

The lesson the entire industry learned from GPT-3's success was brutally simple: *physical scale matters more than theoretical elegance*. If you increase the size of the model and feed it more data, emergent behaviors start to emerge.

AI is no longer just a field of pure scientific research and is becoming a fierce competition of heavy infrastructure. Whoever owns the largest databases, who controls the production of chips, and who has access to huge capital ensures their supremacy.

For the first time, the model is no longer just described in a scientific article, but is transformed into an economic utility distributed via API.

This API works as a true digital pipeline: a programmer anywhere in the world can connect to OpenAI 's central servers, using the generative capacity of billions of parameters directly

in their own application or business, without having to build or run such technology themselves.

Through this commercial conduit, the meteorite accelerates its fall, being directly attracted by the gravity of the global market. From this moment on, it was no longer attracted only by the curiosity of research. It was also attracted by money.

People of the minute

Minute 8 belongs to the gigantic team that signed the fundamental work GPT-3, proving that massive scaling of language models generates exceptional adaptive capabilities without retraining: **TB Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, DM Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever and D. Amodei.**

It also belongs to the engineers and business visionaries who transformed this theoretical colossus into an accessible service through the digital pipeline of the API, opening the gate through which artificial intelligence began to be directly connected to the processes of the real economy.

Chapter 4
Minute 7 - The API and leaving the lab

The year is **2021.** GPT-3 was no longer just a paper read in PDF format on researchers' screens. It had become an active service. An API. This is the great mutation of the Minute 7.

Behind this dry acronym, API, is actually a universal digital pipeline. A plug-in. Through this pipeline, the giant model of billions of parameters, housed in remote servers, could be directly linked to any application, product, or workflow on the planet.

A large model, locked in a research laboratory, remains a simple scientific experiment.

A model accessible through a digital pipeline becomes infrastructure. And infrastructure no longer demands to be observed. It only demands to be used.

To use this computing power, a programmer no longer needed to download the digital giant or own multi-million dollar supercomputers. He just had to send a text through this pipe and receive back, in fractions of a second, the processed response.

This radical simplification changed everything.

This is where the real exit from the laboratory begins. Not through public spectacle or sensationalist press headlines, but through silent infiltration. The most profound changes don't always come in through the front door.

The digital API socket opened a large gate between the model and the real world. Behind the gate was raw technology; in front of it were companies, applications, business needs, workflows, and money. AI stopped falling like a single block. It no longer came as a single object out of the sky. It came through thousands of small “pipelines”. It began to fragment and penetrate piecemeal: into an administrative form, into an online

editor, into a technical support assistant, into a marketing campaign, or into a financial analysis tool.

Social systems sense historical changes long before the general public sees them on screen. They sense them when technology becomes a conduit. When a technology becomes a conduit, it's no longer just a choice. It becomes a circulation.

In September 2020, Microsoft exclusively licensed the GPT-3 model. Strategic partnership with the cloud platform Azure showed a clear direction: the language model was integrated directly into the global industrial cloud. It became a commercial utility, metered by the second and sold like electricity.

But in 2021, the clearest sign of transformation appeared: the Codex model and the launch of GitHub Copilot. The model no longer just generated essays, summaries, or polite replies. It began to write software programs.

And code is not a simple descriptive text. Code is language that acts. Here language ceases to be just expression. It becomes mechanism.

A well-written paragraph can persuade a reader. A good report can inform a decision-maker. A well-worded email can coordinate a team. But code runs directly.

Code modifies databases. Starts servers. Controls physical interfaces. Connects systems. Moves financial resources in fractions of a second.

That's why Codex and GitHub Copilot were not just useful assistants for programmers, but represented the first concrete proof that the language model can directly intervene in the mechanisms through which the digital world builds itself.

The digital pipeline had been directly coupled to the global software production engine. From that moment on, AI wasn't just helping humans write about the world. It was helping them modify the digital world itself.

Impact

Minute 7 changes the status of the language model: it transforms from an object of study into a connectable capability.

Through APIs, artificial intelligence begins to become invisible. Invisibility is one of the most powerful forms of the environment. When you interact with a dedicated virtual assistant, you clearly know that you are talking to a machine. But when the autocomplete function in a text editor or email inbox suggests the right phrase, the boundary of awareness disappears. The user sees only a convenient function of his program, but behind that function runs a neural network of enormous complexity. You no longer feel the meteorite as an object. You feel it as a facility.

This is the early form of the ecological environment. The environment doesn't present itself every time; it just exists and surrounds us. You use it naturally, without analyzing it.

And practical utility is the fastest path to cognitive dependence. If a system reduces work time, eliminates repetitive effort, and decreases uncertainty, it is instantly absorbed into economic structures. Human systems do not adopt technology for moral or philosophical reasons, but are drawn to it mechanically, by the gravity of efficiency.

For programmers who have adopted Copilot, the impact has been multiplied exponentially. Software developers are the connecting points of the digital world. When a user uses AI, the effect is isolated; when a developer uses AI to write code faster, the effect cascades across all the apps and platforms they build for the rest of society.

A new loop emerges, with formidable speed: AI helps write code, new code accelerates system development, systems integrate new AI functions, and these require even more processing power.

The meteorite was descending deep into the gravitational field of the market. And the market doesn't ask if a technology is understood. It asks if it works.

Acceleration

In minute 7, the accelerating force moves from researchers in laboratories to economic distribution networks.

Progress is no longer guided solely by academic papers, but is beginning to be dictated by cloud contracts, per-second runtime costs, response latency, and infrastructure optimization. Thousands of developers are simultaneously testing the same digital pipeline, putting machine-generated language where previously painstaking human labor was required.

Some experiments fail miserably; others remain mere digital toys. But the few that work are enough for the market to finally understand the new historical direction.

People of the minute

Minute 7 belongs to the team behind the Codex model, led by **M. Chen,** who demonstrated that a model derived from GPT-3, trained on massive amounts of source code, can translate instructions expressed in natural language directly into functional programming logic.

It also belongs to the visionary **teams** at **OpenAI** and **Microsoft** who have integrated this computing capability directly into programmers' daily editors via GitHub Copilot, definitively moving AI from the window of the academic world to the window of real work.

At **minute 7,** AI had not yet reached everyone. But **it had** already entered the mechanisms that build the digital world for everyone.

Chapter 5
Minute 6 - ChatGPT and mass contact

The year is **2022.** AI no longer requires API “pipelines”, developer accounts, or sophisticated integrations. It just asks a question. And the question is the most vulnerable form of humanness. This is the regime change of the Minute 6.

Up until this point, the GPT era had remained invisible to the vast majority of people, a technology used behind the scenes by programmers and companies. But the launch of ChatGPT in late November 2022 moved the door directly in front of the entire society.

You no longer had to know what an API was. You no longer had to write lines of code. You no longer had to be an expert. The technical barrier was gone. The psychological barrier would fall right after it.

You entered a simple dialog window. You typed. You received an instant response. So simple, it seemed harmless.

This radical simplification of design was not a simple interface optimization, but represented the moment of direct contact between the meteorite and the gaze of the general public. The AI ceased to be perceived as abstract software; it began to be felt as an addressable presence. It was not human. But it occupied the place of a voice.

ChatGPT conquered the world not through complex theories, but through conversation. And conversation is the oldest and most practiced interface of the human species. The machine adapted to the natural way in which people already know how to ask: through words.

The classic computer required exact commands. The Internet required keywords for search. ChatGPT provided the answer directly. And the direct answer is seductive precisely because it eliminates the path.

The difference is colossal. Classic search kept a safe distance: you searched, selected links, compared pages, read and built your own conclusion. Conversational response dramatically shortens this distance: it delivers you the text ready-digested, formulated and adapted to your tone or size requirements.

Searching the internet would show you where the library was. ChatGPT seemed to speak to you directly from within it. The library didn't have to be traversed. It seemed to lean towards you and whisper the conclusion.

The model was not infallible. It made frequent mistakes, invented facts with complete certainty, and simplified complex realities. But the linguistic form was so fluent and useful that millions of people felt for the first time that they were receiving real help in their daily cognitive routines.

Not just raw information. Help. And **repeated help** quickly becomes **a reflex.**

The first signs of mass delegation appear. In just five days after launch, ChatGPT reaches the historic threshold of one million users. In mid-December, I finally had an active account, after days of seeing only the famous message: «*ChatGPT is at capacity right now*». Three months before, OpenAI had sent me, dryly, a: «*You're on the DALL·E waitlist*».

This speed of adoption was not just a marketing victory, but a profound historical signal: the human species had instantly recognized the usefulness of an artificial interlocutor available at all hours, without fatigue and without limits of patience. For the first time, infinite patience seemed available in a browser window.

Impact

Minute 6 completely redefines the psychological relationship between humans and artificial systems.

Through the dialogue interface, the AI becomes personally addressable. The user can ask it questions, correct it, ask it to change its perspective, or explain a complicated problem in a way that a child can understand. This continuity of dialogue creates a deep functional attachment: when a tool repeatedly reduces your daily mental effort, it quickly enters your unconscious reflex.

This is where the major vulnerability of this minute comes in. The danger lies not in the machine's occasional errors (which can, in theory, be checked), but in the fact that its usefulness is great enough that people start to move the thought process outside their own minds.

A quick rewording, a document summary, a difficult email draft, a line of code, a small organizational decision. Step by step, mental effort becomes optional in more and more places. It doesn't disappear through prohibition. It disappears through convenience.

And in a time-pressured environment, anything that becomes optional tends to become rare. That's the silent part of loss: no one takes your thinking away. You're just no longer obligated to use it.

The new ecological environment does not prohibit critical thinking. It merely makes it optional, offering a convenient shortcut over effort. And the shortcut, repeated daily, becomes a map.

A profound divergence is thus established among the species: not so much between those who use technology and those who reject it, but between a minority that uses these models to expand their analytical capacity and a majority that uses them to stop making the effort to think where they should.

Acceleration

At minute 6, the acceleration engine becomes purely social and decentralized.

Every conversation shared on the internet, every example of *"look what this system can do"* worked as an active agent of propagation. Human curiosity proved to be a far more powerful force of distribution than any organized commercial campaign. It was not pushed into society. Society called it.

Traditional systems have found themselves forced to react to a completely new environmental reality. Schools and universities have been forced to ask themselves urgent questions about the meaning of written assessments and homework; newsrooms have had to reexamine the authority of journalistic text; companies have begun to calculate how many more people they need to achieve the same volume of administrative work.

The debate about artificial intelligence has definitively broken away from the narrow circles of specialists, becoming the great problem of adaptation of the entire society.

People of the minute

Minute 6 belongs to the team of researchers who managed to educate the raw nature of language models, making them more obedient and capable of following people's intentions.

Through the **InstructGPT** research line and the development of **RLHF** (*Reinforcement Learning from Human Feedback* / alignment through feedback received from human evaluators), the team consisting of: **L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike and R. Lowe** took the decisive step of domesticating the machine.

They demonstrated that simply enlarging a model does not automatically make it more useful or safer, requiring a rigorous process of education through human interaction.

It also belongs to the **OpenAI team** that put this directly aligned capability into the hands of the general public through ChatGPT, permanently changing the human-machine contact interface.

In **minute 6,** the meteorite was no longer just in the sky. It was in **dialogue** with us.

Chapter 6
Minute 5 - Humanity begins to attract the meteorite

The year is **2023.** The period of collective wonderment abruptly ends. The era of massive demand and global strategic competition begins. From this point on, AI no longer needs to convince humanity. Humanity begins to call upon it.

Just two months after launch, ChatGPT reaches the historic milestone of 100 million monthly active users. It didn't take a generation for this adoption. It didn't take a decade. It took just eighty days.

This figure is not a simple marketing statistic. It is a historical signal of enormous gravity. It wasn't just adoption. It was **attraction.**

Humanity had instantly recognized the ecological utility of the system. And utility works like a mechanical force of attraction. A technology may be viewed with suspicion by philosophers, contested by activists, or regulated with fear by politicians. But if it considerably reduces working time and eliminates individual effort, it will be instantly absorbed into the flow of society.

Humanity doesn't just adopt what it understands. It adopts as a priority what reduces its effort. And effort is one of the last natural barriers to thought.

This is where the meteorite really starts to be attracted. Not by a force of the system's own, but by us.

Each user draws on their own interests: the student who wants the assignment done faster, the employee who wants the report completed before the deadline, the programmer who wants the code written without syntax errors, the company looking to reduce its administrative costs, the teacher who wants teaching materials without effort, the journalist who needs a quick draft.

Millions of small rational decisions at the individual level accumulate into a huge field of attraction. No one decides the collapse. Everyone decides only their shortcut.

On March 14, 2023, OpenAI released **GPT-4.** It was no longer just a minor improvement, but a massive leap in capability. It was no longer limited to simple friendly chats; it could solve complex academic exams, understand dense technical documents, process images, and demonstrate superior problem-solving logic to many human professionals.

With GPT-4, the public standard of expectation shifted forever. What had seemed like a technological miracle in November had become the new minimum accepted level by March. This is one of the harshest laws of acceleration: miracles "age" immediately.

Just two days later, Microsoft announced the integration of Copilot-like assistants into its Microsoft 365 office suite. AI was no longer a separate destination that the user visited in a browser; it had become a layer integrated directly into Office applications. When it enters the tools of daily work, a technology no longer waits to be chosen. It becomes available before you ask yourself if you want it.

You didn't go to the machine anymore. The machine came directly into your document. In your email. In your meeting. In your spreadsheet.

At the same time, Google is opening up access to its own conversational model, Bard.

The competition between tech giants quickly turned into a massive industrial race for market positioning. Global platforms were not just distributing software products, but installing new mass behavioral reflexes.

A simple chatbot can go viral. An integrated platform becomes a living environment. And a living environment doesn't install through shock. It installs through habit.

Impact

The year 2023 is structurally changing the relationship between institutions and artificial intelligence. The fundamental questions have shifted from the general *"what is this system?"* to the much more pragmatic *"what tasks can it take on for me?"*.

Mass cognitive delegation kicks in. People don't just use AI to search for raw information, they delegate entire parts of their thinking process to it: formulation, logical structuring, argumentation, synthesis, and small decision-making.

And when parts of the thinking process become delegable, the cognitive environment changes completely. Not because the person suddenly stops thinking. But because they start thinking more and more often in terms of something else.

The ecological pressure is becoming visible: the environment is changing rapidly even for those who have not decided to change it. An institution or a company can no longer simply say: *"it doesn't concern us"*. It concerns them directly. Students use the systems, employees covertly integrate them into their daily tasks, competitors use them to deliver faster, and financial markets penalize slowness.

In an already accelerated world, slowness becomes the biggest competitive flaw. Here fear changes its form. It is no longer the fear of AI, but the fear of being left out of it. Those who use AI gain time; those who do not use it risk becoming irrelevant.

This is no longer a personal choice. It is a total systemic pressure. When everyone else is speeding up, caution starts to resemble delay.

Acceleration

In minute 5, the accelerating force moves completely into the area of economic demand and industrial competition.

Every promise of productivity translates into a huge competitive pressure. Companies are building new products in

a hurry, investors are moving their capital where they see quick returns, platforms are constantly adding generative features, and chipmakers and cloud providers are pushing their computing capabilities to the limit.

AI is no longer just being accelerated by research in laboratories. It is being accelerated by the entire global economy that wants to use it to survive the competition.

People of the minute

Minute 5 no longer belongs to a single laboratory or a small number of isolated researchers. It belongs to large industrial structures that have moved technology into the mass production phase.

It belongs to the **OpenAI teams** that developed GPT-4, demonstrating that an advanced system can surpass the threshold of simple reasoning, becoming a real professional work partner.

It belongs to **the Microsoft teams** that introduced Copilot to the office suite, transforming AI from an isolated application into an invisible layer integrated into the daily routine of cognitive work.

It belongs to **the Google teams** that responded quickly by opening Bard, definitively entering the great race of platforms that configure the digital behaviors of all modern civilization.

By **minute 5,** the meteorite wasn't just falling toward humanity. Humanity had begun to **accelerate** its fall.

Chapter 7
Minute 4 - When AI becomes flow

The year is **2024.** AI breaks its old format barriers and begins to flow freely between different forms of expression: text, image, voice, code, documents, and real-time interactions. When something flows through all forms, it becomes difficult to locate. And what you can no longer locate begins to resemble the environment.

The big change this minute is the expansion of two fundamental coordinates: native multi-modality and the massive expansion of the working *context*. *The* AI no longer just wanted to respond better. It started to **spend more time** with everyone.

To understand *the context,* we can simply define it as "user working memory", that is, the total volume of information that the system can simultaneously keep in attention, actively, during a work session, without forgetting details from the beginning of the interaction.

In February 2024, Google announced the Gemini 1.5 Pro, equipped with a revolutionary working memory of up to **one million tokens.** If we understand a token as a **"part of a word"** (the machine does not "read" whole words, but fragments them into smaller pieces of characters), a working memory of one million such parts means that the system can simultaneously keep in mind an entire file of thousands of pages, the source code of an entire application, or hours of audio recordings.

The system no longer just provides a short answer to an isolated question. It actively dwells and orients itself inside your problem, without losing focus. A machine that doesn't forget easily quickly becomes more present than many people.

In May 2024, OpenAI releases **GPT-4o.** The model is built to natively process text, images, and audio in real time.

This changes the control interface again. The user is no longer forced to type; he can speak directly to it, interrupt it in motion to ask for clarification, show it images through the phone's video camera, and receive instant answers in a voice with dynamic intonations.

The keyboard barrier has been removed. Interaction is directly approaching our biological rhythm: voice, gaze, gesture. The more natural the interface becomes, the weaker the critical defense becomes.

In parallel, Meta is breaking new ground by launching the Llama 3.1 405B model, demonstrating that these top models no longer remain exclusively behind private servers, locked down by a single corporation, but can circulate in open forms, being freely run, analyzed, and modified by the global developer community.

Impact

In minute 4, the use of artificial intelligence undergoes a subtle mutation: AI becomes a diffuse presence integrated directly into phone functions, online search engines, or the documents we edit at the office. You no longer enter a place called AI. You encounter it everywhere.

When you clearly observe technology, you maintain a critical distance from it. When you no longer observe it, it has already become part of your living environment. Danger is not only what scares you. Danger is also what ceases to seem foreign.

This fluid presence continues to subtly change the cognitive muscle of the human species. Every small decision to let the system summarize a long text, compose a polite email, or check a line of code acts as a micro-shift of individual effort. Effort gradually becomes a rare option, replaced by the convenience of a ready-made machine response. You don’t feel

the loss the moment you choose comfort. You feel it later, when the effort already seems unnatural.

In this usage cycle, AI approaches human biological rhythms, but is not limited by them. It can process colossal volumes of data simultaneously, assist millions of people at the same time, and keep its working memory with the user intact, generating huge ecological pressure on the traditional way in which people form their opinions, decisions, and trust information.

Acceleration

In minute 4, acceleration manifests itself through fluidization and lateral spread. Technology no longer grows only vertically (through increasingly intelligent models), but also expands horizontally, entering all our daily digital gestures. This is the mature form of infiltration: it no longer just climbs in capacity, but spreads in habit.

The system adapts naturally to the way we already live: whether you write, search, speak, present, or read, AI automatically appears as a function integrated directly into those actions. It doesn't change your life by command. It changes your gestures by suggestion.

The gravitational field is becoming increasingly dense. Daily user demand, rapid company integrations, mass distribution of global platforms, and strategic state funding are all coming together in a single accelerating force.

People of the minute

Minute 4 belongs to the teams of researchers and engineers who managed to take technology beyond the limits of written text, transforming it into a continuous stream of multi - modal interaction.

It belongs to the **Google DeepMind teams,** which have radically expanded the limits of user working memory,

allowing advanced systems to process contexts of colossal size without losing their accuracy.

OpenAI team that developed GPT-4o, led by **M. Murati, M. Chen, and B. Zoph,** who removed the barrier between text, voice, and image in real time, offering the audience a fluid, latency-free interaction.

It belongs to the **Llama teams at Meta,** who have demonstrated that cutting-edge models can also circulate in an open format, breaking the monopoly of closed doors and offering the entire developer community an advanced integrable computing capability.

By **minute 4,** AI was no longer just a model, an application, or a product. It had begun to **flow** through the forms of digital life.

Chapter 8
Minute 3 - Infrastructure becomes the protagonist

The year is **2025.** AI is no longer just an artificial mind floating in the abstract air of the internet. It shows its massive physical body. This is the regime change of minute 3.

Up until this point, for the average user, the operation of artificial intelligence seemed immaterial. You write a simple requirement on the screen, you get a well-structured text in a few seconds, and everything seems to come out of nothing. But nothing consumes megawatts.

But artificial language doesn't come out of thin air.

It comes from giant data centers. It comes from extremely rare silicon circuits. It comes from massive electricity consumption. It comes from millions of liters of water used for cooling, from fiber optics, from complex networks, and from decades-long supply contracts.

At minute 3, the AI shows its heavy industry size. The answer seemed easy. The body that produced it was huge.

The meteorite is no longer just a spectacular light in the sky. It has raw industrial mass. And mass always changes the size of the impact.

Stargate project is becoming a symbol of this material transformation. OpenAI, Microsoft, Oracle, and SoftBank are announcing plans to develop a computing infrastructure on an unprecedented scale: dedicated data centers and energy grids, with announced investments in the hundreds of billions of dollars. We are no longer talking about simple software run by a group of engineers in an office. We are talking about a giant industrial system that claims its own physical territory on earth. When a technology starts demanding territory, energy, and water, it no longer belongs only to the screen.

This definitively changes the meaning of the word "model".

The performance of a leading-edge system no longer depends solely on an ingenious algorithmic idea, but becomes a matter of material strength. Who controls the production of specialized chips? Who has secure access to high-capacity power grids? Who owns the cloud infrastructure needed to support billions of simultaneous requests? Who can bear the gigantic costs of operating and cooling servers?

In minute 3, AI becomes a tough physical competition. It's not just ideas competing anymore. It's also electrical grids competing.

With this materialization comes a total strategic opacity. While the technology was a mere scientific curiosity in the laboratory, researchers openly published the number of parameters, the structure of the algorithms, and the data sets used for training.

But when the system becomes a critical infrastructure of security and economic power, secrets are secured. Transparency retreats just as dependence grows.

Advanced top-of-the-line models are protected as strict trade secrets. Parameters are no longer published. Data sets are described partially, in general terms. Actual costs are only estimated by analysts. Public access is controlled by strict filters, and the most powerful variants are selectively offered only to certain governments, large corporations, or strategic partners.

AI is no longer just a piece of software you download from the internet. In the area of frontier models, it is beginning to resemble energy resources, national telecommunications, or the defense systems of states: critical infrastructure, expensive, opaque, and controlled by a small number of actors.

But this is not the whole story.

In parallel, open-weight models open an opposite direction: AI as a distributed, modifiable, auditable, and locally

run infrastructure. If frontier models are the centralized "Windows" of the AI era, open models can become its "Linux": harder to use for the general public, but essential for autonomy, research, education, audit, and technological sovereignty. Open models are not just a technical alternative. They are the place where the ecosystem stores its antibodies, but also the risk of its own mutations.

The tension of minute 3 is not, therefore, the simple centralization of AI. It is the struggle between two forms of infrastructure: one closed, gigantic and strategic; the other open, distributed and difficult to fully control.

NVIDIA is becoming the central name of this minute. Not because it develops language models, but because it produces the specialized chips that are the physical muscle of the entire digital age. Without these computing accelerators, the language model remains a simple equation on paper. Without massive data centers, the answer on the screen remains an unfulfilled promise.

NVIDIA *data center* division's revenue is no longer just a stock market indicator. It's a physical signal of an era: thousands of tons of silicon are being bought up in masse to support the planet's new cognitive environment.

Impact

Minute 3 dismantles the illusion of technological immateriality. AI has an electric bill. AI has physical terrain. AI has a global supply chain. AI has hard material constraints. Immateriality was just the interface.

And these constraints bring about a huge concentration of power. The ability to develop and run advanced systems is restricted exclusively to those actors who possess the colossal capital and the necessary physical infrastructure.

There is a huge gap between the simplicity of the interface and the complexity behind it. For the user on the

phone, everything comes down to a simple tap on the screen. The gesture is local and free. But the mechanism activated is planetary, opaque and expensive. This is the new asymmetry: the finger presses here, the infrastructure lights up somewhere far away.

This profound asymmetry hides dependency. The easier it is to get a result on the screen, the harder it is to be aware of the power infrastructure that makes it possible. People often confuse easy access with their own autonomy of decision, forgetting that they are connected to a pipeline that is absolutely controlled from a distance. Access does not mean ownership. Use does not mean sovereignty.

Acceleration

At minute 3, the acceleration force moves into the area of heavy industrialization. It no longer relies solely on better software codes, but on building physical accelerators of acceleration.

Every new data center that opens is not just a server building; it is a factory capable of generating and distributing massive volumes of artificial language. It is a new “language factory”.

A self-referential physical loop emerges: people build supercomputers for AI, AI is used intensively to optimize and design new generations of specialized chips, new chips allow training of even higher capacity models, and these expand market demand, justifying the construction of even larger data centers.

The meteorite is no longer just attracted by our gravity. Humanity is paving its path. We are not just calling it. We are building its runway.

People of the minute

Minute 3 belongs to the builders of the physical body of the machine.

It belongs to **the NVIDIA team** and **J. Huang,** who understood many years ago that the future of digital computing would rely on high-power parallel physical accelerators, becoming the exclusive suppliers of industrial muscle for the new era.

It belongs to the infrastructure leaders who designed the transition of AI into a planetary utility investment regime: **S. Altman, L. Ellison and M. Son,** along with the **OpenAI, Oracle and SoftBank teams.**

They are the ones who decisively moved the issue of artificial intelligence from the area of algorithms to the area of energy networks and large physical infrastructure construction sites.

At **minute 3,** the AI stopped looking like a cloud. It became **a weight.**

Chapter 9
Minute 2 - Agents and the shift from response to process

The year is **2026.** The AI no longer passively waits for the user to ask it a direct question. This is where the real anxiety begins: not the answer, but the continuation. It begins to act continuously between questions. The space between two human commands no longer remains empty.

This is the historic change of the 2nd minute.

Up until this point, the entire evolution of the GPT era had been a continuous acceleration of on-screen response. You type a prompt in the chat window, the machine generates adapted text for you, and the basic unit remains the direct conversation: human, request, finished response.

Autonomous agents completely change this basic unit. The unit of work is no longer the isolated response. It becomes the ongoing process.

chatbot answers. The agent continues. And what continues can quickly escape the field of human attention.

The difference seems minor only on a grammatical level; historically, it's a huge leap.

The user no longer coordinates each step manually. He gives up the path and keeps only the destination, and gives the machine only a general objective (for example: *"identify the errors in this application, look for solutions in the official documentation, rewrite the affected sections, test the result in an isolated environment and deliver me the final report with the changes made"*).

Once the objective is received, the agent assumes its deployment. It reads internal files, autonomously browses the internet, runs software tests, encounters execution errors, analyzes its own failure, rewrites its code, tries a new path,

resumes the testing process, and completes its tasks without requesting confirmation at each step.

The system no longer just generates words on the screen. It generates autonomous trajectories of action. And a trajectory is not as easily verified as a sentence.

In the first half of 2026, the use of these agent assistants (such as Codex, Claude Code, Cursor, Aider or OpenHands) is experiencing exponential growth in professional workflows.

Data from the study by D. Johnston, D. Holtz, A. M. Richmond, C. Ong, P. Tambe, and A. Chatterji shows that the number of professionals actively using agent assistants has multiplied massively in an extremely short period of time. More than 10% of developers currently manage networks of three or more autonomous agents simultaneously in a workweek.

People are no longer simple operators typing requirements; they become administrators of artificial cognitive processes. The human does not disappear from the process. He moves higher, but sees less.

But this autonomy of execution hides a huge consumption of digital resources. Research published in the study coordinated by L. Bai, Z. Huang, X. Wang, J. Sun, R. Mihalcea, E. Brynjolfsson, A. Pentland and J. Pei reveals a crucial technical detail: tasks executed by autonomous agents can consume up to 1,000 times more *tokens* than a typical conversational query: *"... agentic tasks are unique expensive, consuming **1000x** more tokens than code reasoning and code chat.."*.

To understand this consumption leap, we need to go back to the definition of **a token** as a "part of a word."

If a simple question to a chatbot consumes 100 such parts of words to deliver a short answer on the screen, a single general goal given to an agent will trigger a huge work loop in the background: long document readings, successive code runs, internal error analysis and repeated self-corrections.

Without the user seeing this effort directly on the monitor, the autonomous agent consumes 100000 word parts in the background just to deliver the clean final result.

AI is starting to produce a huge amount of invisible cognitive work. And invisible work is hard to control, hard to morally tax, and hard to audit.

Impact

Minute 2 structurally alters the nature of delegation and **accelerates** False Cognitive Power Transfer (**FCPT**).

When a human writes down each step of a task, the effort is visible and the responsibility is clearly understood. When the autonomous agent fully assumes the intermediate steps of the process, the human is removed from the role of direct executor and placed in the role of simple manager who approves or signs off on the final result.

This withdrawal increases productivity in the short term, but it installs a dangerous opacity. The user sees only the finished product, clean and functional, but loses touch with the "road": he no longer understands the logical assumptions on which the machine relied, the alternative hypotheses that it eliminated in the background, or the small errors masked during auto-correction. The result becomes legible. The process becomes foggy.

And losing your way means losing the space in which one's own mental competence is formed and consolidated. Competence is not formed in signature. It is formed in difficulty.

Cognitive divergence among users reaches a critical level: some use agent networks as multipliers to orchestrate complex architectures that they understand in depth, while others become mere passive beneficiaries of automated results, signing off on reports, codes, and decisions that they can no longer critically evaluate. Here, responsibility remains human, but control becomes partially delegated.

Acceleration

In minute 2, the acceleration becomes an over-acceleration generated by the movement of the control.

In the era of conversational chatbots, the volume of activity depended directly on the physical presence of the user: a human had to sit in front of the screen, think, and type a new question.

In the era of the autonomous agent, activity is freed from the limitations of human biological time. The machine does not sleep between two intentions. A single user can simultaneously start ten different agents that run in the background, analyze data, test hypotheses, and interact with each other continuously, consuming colossal volumes of data per second.

The throughput of artificial mental activity no longer increases simply with the number of people asking questions. It increases exponentially with the number of internal processes that machines run autonomously to perform delegated tasks. Here, the speed is no longer the speed of the assisted human. It is the speed of the process detached from the human.

The meteorite no longer just falls faster. It begins to fragment the entire space of our cognitive work into invisible automatic processes. We no longer just see the fall, we see the multiplication of fragments and the dispersion of the impact.

People of the minute

Minute 2 belongs to the software engineers and researchers who have transformed artificial intelligence from a passive dialogue engine into an autonomous participant in the execution of complex projects.

He belongs to the teams that developed the first integrated agent development environments, such as those behind **the Codex, Claude Code, Cursor, Aider and**

OpenHands systems, definitively moving office work into the era of autonomously assisted processes.

It belongs to the researchers who first analyzed and documented the actual behavior of employees in the face of these tools: **D. Johnston, D. Holtz, A. M. Richmond, C. Ong, P. Tambe, and A. Chatterji.**

It also belongs to the researchers who measured the energy impact and massive consumption of invisible resources generated by running agent loops: **L. Bai, Z. Huang, X. Wang, J. Sun, R. Mihalcea, E. Brynjolfsson, A. Pentland and J. Pei.**

Through their efforts, these people have shown that the transition to the agent era is not just a simple software improvement, but the installation of a completely different processing regime in the ecosystem of our civilization.

In **minute 2,** the AI was no longer standing in front of us as an interlocutor. It had started working behind us as **a process.**

Chapter 10
Present - Last minute

July 1, 2026. Exactly six months after the publication of the first volume, *METEORAIT,* on January 1, 2026, this date no longer represents a closed minute that we can analyze from a distance. Half a year in the **acceleration era** is equivalent to decades in ancient technological history.

This is the immediate present. And that's precisely why it's the hardest to observe objectively. We're inside the phenomenon we're trying to describe.

The present does not have the decency of historical archives. It does not wait for us to understand it. It moves, reconfigures itself, and alters our habits even as we try to draw its boundaries. In the first nine minutes, we were able to follow a clear, step-by-step trajectory from discreet beginnings in laboratories to the daily invasion of digital assistants.

But the present is not just another step on this ladder. It is the moment when the ladder begins to move beneath our feet. It is the point of impact where all the threads meet and collide simultaneously: the industrial volume of artificial text, the massive expansion of agents, the strategic opacity of cutting-edge systems, and, above all, the algorithmic war for survival.

Humanity has written slowly. On stone, on clay tablets, on papyrus, on parchment, on paper, and finally on digital screens. We have written laws, prayers, commercial registers, political treatises, novels, love letters, software code, technical analyses, administrative shorthand, and footnotes. We have written as much as we can, for millennia, consuming biological life time. AI has entered this “river of words” and changed **the flow forever.**

In raw **volume, global AI** can currently write, in about **a week,** as much text as all of **humanity has written** throughout its entire documented **history.**

Due to the lack of clear and complete data, I cannot say whether the exact ratio is one to one, a little under or a little over. **The historical threshold is important:** artificial language production has entered the order of magnitude of the entire written history of mankind. It took **humanity millennia** to sediment this volume. **AI infrastructure** can reach it in **weeks.**

Advanced systems did not invent language or knowledge. They invented something else, much more to fear: the industrial production of adaptive language in real time. It is not the isolated lie that is new. What is new is the flow. Language on demand, without fatigue, without sleep, without its own biography and without a direct human author.

The historical limit of writing, represented by the human body and biological attention span, has been definitively shattered. When a human writes, something of him is irreversibly consumed. When AI writes, only the physical infrastructure is consumed: specialized chips, electricity, working memory, and cooling systems. The human gets tired after a few pages; the infrastructure scales and generates text incessantly in the background.

And with autonomous agents (Minute 2), the nature of this flow has completely transformed. The conversational chatbot only responded when the human asked, maintaining a direct relationship of control. The autonomous agent continues the process after the human has left the monitor. It runs in the background, planning, testing, failing, analyzing its errors, rewriting its code, trying again, and delivering the final result directly.

In the conversational phase, AI was dialogue. Today, AI becomes autonomous process execution. Dialogue could be interrupted. The process continues.

And if in civilian life this autonomy translates into professional comfort and budget optimization, on the military scene it becomes a harsh condition for existential survival.

The military has always been the absolute pioneer and **most brutal accelerator** of human technology. Where the stakes are not quarterly profits but physical survival in the face of a capable enemy, the luxury of ethical prudence and safety filters is completely eliminated. War selects only that which offers maximum reaction speed and precision of strike.

In modern theaters of operations, autonomous systems are no longer a mere prospect of the future. They define the immediate present of combat.

When intense electronic jamming blocks radio communications between the human pilot and the attack drone, manual control becomes impossible. At that point of jamming, the system is forced to take over control completely autonomously. The drone enters the terminal autonomy phase: it recognizes its target by visual sensors, locks onto the object, calculates its trajectory, and strikes without asking permission from a remote human operator.

AI has penetrated deep into the space between perception and action. And there, in that split second, man no longer has time to be human.

When hundreds of drones, radars, satellites, and jamming systems must be coordinated simultaneously, in fractions of a second that biological reflexes cannot process, war becomes a direct confrontation between algorithmic operating systems. Whoever has the system capable of seeing faster, correlating more data, deciding autonomously, and acting in milliseconds wins the battle. In war, speed usually becomes the final argument.

The moral difference between the assistant who fixes your software code at the office and the drone that autonomously fixes its target in battle is huge. But technical ecology does not respect the moral boundaries that we draw. The technical mechanism behind it is related: replacing the intermediate steps of human decision with autonomous computational processes.

The meteorite doesn't fall faster because it wants to destroy. It falls faster because it is pulled with enormous force by the entire economic, social, and military complex of human civilization.

Each of us draws from our own motive: the market from the desire for profit, the army from the need for survival, the states from geopolitical ambitions, and the average user from the simple search for cognitive comfort. The result is unique: an acceleration that traditional society can no longer digest in a timely manner. Historical digestion has been surpassed by the flow.

In this final minute, the advanced frontier of research becomes completely opaque. Cutting-edge designs protected as trade secrets are no longer described in public works, and their codes are guarded as strictly as nuclear technologies. AI has moved beyond the phase of academic curiosity and has definitively established itself as a strategic infrastructure of power.

In the age of agents, the legal question can no longer remain simply who wrote the code or who pressed the button. The question becomes where liability attaches (*MEG2: Legal Governance for Agentic AI*) when the intermediate process has been executed autonomously.

There is no longer "AI" as a simple isolated computer program. There is an entire artificial ecosystem that surrounds us and modifies our behaviors. An environment that, once installed, radically changes the rules of social selection: those

who delegate without verification gradually lose their intellectual depth; those who control large computing capacities control economic flows; and those who no longer understand the intermediate steps of processes become absolutely dependent on the result delivered by the machine.

This is the reality of the immediate present. Not some distant science fiction prophecy, but the precise finding of **the minute we live in.** It is now, here.

Epilogue
Atmosphere

A meteorite **cannot be stopped** from falling through **the atmosphere,** but its existence, size, and density are decisive for what happens on the ground. The denser the atmosphere, the more kinetic energy is consumed by friction before impact. The thinner the atmosphere, the more of the object's brute force reaches the ground, causing massive destruction.

This is the realistic position on which **MEG** (*Minimal Ethical Governance*). Not denying the fall. Not worshipping the meteorite. Not panic. But building the atmosphere.

MEG / MEG2 do not promise the impossible: they cannot stop the accelerating decline of technology and cannot return society to the state it was in before the emergence of generative models and networks of autonomous agents. They do not represent a simple stop button on progress, but try to build a legal and mental atmosphere capable of taking over from the impact energy of the machine. The atmosphere does not stop history. It makes it bearable.

In the era of autonomous processes, the great dangers are no longer just about the correctness of the answers displayed on the screen, but move to the area of the dissolution of responsibility.

When an autonomous agent produces real negative consequences (financial losses, execution errors, or bad administrative decisions), perhaps through intermediate steps that the programmer did not fully specify, or the user did not directly control, the blame risks being completely dissipated in the system. And society cannot slow down or control what it cannot clearly attribute to anyone. **Without accountability, acceleration becomes a legal fog.**

That is why MEG2 proposes a top-down pillar: **functional legal personality for agents.** MEG2 starts from a simple distinction: giving an agent a functional legal identity does not mean attributing consciousness to it. It means building a point of imputation, where otherwise the blame would dissolve between the producer, the user and the infrastructure.

Legal fiction is not new in the history of law. Humanity has already created the concept of "legal person" for corporations, associations and institutions, abstract entities that have no soul, no conscience and cannot feel suffering, but to which the law attaches clear rights, obligations, patrimony and civil liabilities in order to be able to function safely in society.

MEG2 extends this principle from the corporation to the autonomous agent. Not to humanize software code and not to recognize any form of artificial consciousness in it, but solely as a pragmatic legal mechanism to have a clear point of imputation of civil and financial liability. **MEG Address. Address for liability.**

In this control architecture, MEG2 defines **five operational states** through which an autonomous agent can be technically and legally managed: **Active** (the state in which the agent runs in an authorized manner), **Inactive** (controlled sleep state, with the agent's activity temporarily frozen), **Marked** (the system is kept running, but placed under special observation), **Suspended** (temporary freezing of any action or network access permission), **Deactivated** (the agent's

transition to an irreversible non-functioning state, equivalent to its complete removal from the digital ecosystem).

This is the atmosphere built from the top down, through laws, industry standards and rigorous legal structures. But the law alone cannot think for us. **ACCELERATION** shows why accountability is beginning to dissolve. **MEG2** tries to give it a new address.

But a truly effective atmosphere cannot be left exclusively to state regulations. It absolutely requires building resistance from the bottom up, at the level of each individual, by protecting **cognitive sovereignty.** The legal atmosphere protects society. The mental atmosphere protects the individual.

At this point, the applied manuals (***METEORAIT for Teens*** and ***METEORAIT for Educators***) become tactical tools for mental self-defense against **False Cognitive Power Transfer (FCPT).** They show that this series is not just a collection of cold theoretical analyses, but offers a set of training protocols capable of keeping the human thinking muscle active in the face of the temptation of total machine comfort. The unused muscle atrophies, but it can be trained by applying three simple mental hygiene protocols daily:

1. Controlled Latency: Voluntarily postponing the immediate impulse to open the chat window with the AI in the first second we are faced with a difficult problem. We force our mind to remain in the discomfort of uncertainty, to search for our own associations and to sketch a first logical structure before accessing the artificial answer.

2. The 5-minute wall: In the process of learning or creating, we impose the effort of trying to solve a task completely individually for at least five minutes. Even if the effort fails, this initial mental break activates the networks of attention and understanding, preparing the brain to critically assimilate the answer it will receive later, preventing the installation of the illusion of easy competence.

3. Reverse Engineering as Training: Using AI not as an easy generator of ready-made results, but as a tough partner in auditing and critical debate. The student or professional builds their own text, code or analysis, and then uses artificial intelligence exclusively to identify its weaknesses, show its logical errors or challenge its assumptions. The primary cognitive effort remains with the human, while the machine is retained in its rightful role as critic and trainer. This is where the real partnership begins: not when the machine thinks for us, but when it forces us to think better.

With a solid atmosphere, formed by the MEG2 legal density and the cognitive sovereignty protocols in *for Teens* and *for Educators,* much of the impact energy is spent constructively: the laws protect our social responsibility, while the mental exercises keep our cognitive sovereignty intact in front of the machine.

This is not a *promise* of reassurance, but the pure realism of our intellectual survival. Fear sees the meteorite. **Reason builds the atmosphere.**

METEORAIT describes the object.

ACCELERATION describes the fall.

MEG/MEG2 describe the atmosphere.

In **the last minute,** the size of the crater depends on us. And that's why the last minute is not just about AI. It's about **the maturity of the species that welcomes it.**

This is the next book:

The ATMOSPHERE of METEORAIT
AI as an Ecological Factor, Volume III
Structural Resistance through IGT and MEG/MEG2

Adrian (Adi) Stan
01.07.2026