

THE FUTURE ALREADY HAPPENED AND ALMOST NOBODY SAW IT

An artificial intelligence has already rewritten biology and took part of a Nobel for it, almost without a headline. Meanwhile, the founders of the field had already warned, eighty years ago, where this road took a turn.

While the world waits for a humanoid robot in the living room, an artificial intelligence has already rewritten biology — and took part of a Nobel for it, almost without a headline. While companies burn billions on AI projects that never leave the pilot stage, the founders of the field itself, eighty years ago, had already warned exactly where this road took a turn. This is an invitation to look in the right place — before the next future also goes unnoticed.

In 2024, part of the Nobel Prize in Chemistry went to Demis Hassabis and John Jumper, of the DeepMind laboratory, for having built an artificial intelligence — AlphaFold — capable of predicting the three-dimensional structure of a protein from its genetic sequence. It is a problem biology had been chasing for fifty years. AlphaFold has already predicted the structure of more than 200 million proteins — practically the entire catalogue known to science — and is now used by millions of researchers to accelerate vaccine development, the fight against neglected diseases, and the basic understanding of life.

Ask ten people on the street whether they know AlphaFold. Probably none of them will. Ask whether they have heard of a humanoid robot juggling in a promotional video, or a chatbot writing an email. The answer changes.

This is not an article against enthusiasm. It is an invitation to redirect the gaze — because the future that promises to change everything has sometimes already arrived, quietly, in a

laboratory, while the future that promises to change everything in a headline is still stuck in the demo.

WHAT THE FOUNDERS OF THE FIELD ALREADY KNEW

In 1945, before a commercial digital computer existed, the engineer Vannevar Bush wrote an essay called “As We May Think”, imagining a machine — the memex — that would help a human being organise and retrieve their own memory. Not to replace human judgement. To extend it.

In 1960, the psychologist and computer scientist J.C.R. Licklider published “Man-Computer Symbiosis”, with a sentence that is still uncomfortable to read today: the goal is not the machine doing man's work, it is man and machine cooperating, each doing what it does best. Two years later, Douglas Engelbart wrote “Augmenting Human Intellect”, with the same bet: technology to amplify collective intelligence, not to dispense with it.

In 1966, the philosopher Michael Polanyi coined the phrase this text has already quoted once and will quote again, because it is the axis of everything: we know more than we can tell. In 1971, the economist and cognitive scientist Herbert Simon — another name who helped found the field of artificial intelligence itself — wrote a sentence that reads as if written yesterday: *“a wealth of information creates a poverty of attention.”* Simon was warning, decades before infinite feeds or push notifications existed, that the scarcity of the future would not be data. It would be focus.

And in 1976, Joseph Weizenbaum — the scientist who created ELIZA, one of the first programs to simulate conversation, and who spent the rest of his career alarmed by his own invention — published “Computer Power and Human Reason”. The central argument: there is a difference between a machine that processes information and a machine that understands. Delegating judgement to the first, believing you are delegating to the second, is the categorical error of the century.

Eighty years of warning, coming from inside the field itself, always in the same direction: augment, do not replace; and never confuse processing with understanding. The question of this text is simple — why are we still surprised when the market ignores the warning?

THE SIDE NOBODY HAS THE NERVE TO HIDE

The numbers of 2026 confirm the pattern the founders predicted, with uncomfortable precision.

The RAND Corporation, in a widely cited study, found that around 80% of enterprise artificial intelligence projects fail to deliver the intended value — twice the failure rate of technology projects without AI, and that number has not moved in three years of measurement. S&P Global Market Intelligence found that 42% of companies abandoned most of their AI initiatives in 2025 — almost triple the 17% of the previous year. McKinsey describes the phenomenon as the “generative AI paradox”: around 80% of companies have deployed some form of the technology, and around 80% cannot point to material impact on the bottom line. Only 39% attribute any measurable impact, and a small group — close to 6% — concentrates most of the real value captured.

Gartner is even more specific about the cause: more than 40% of “agentic” AI projects — systems that act on their own, not merely respond — will be cancelled by the end of 2027, and the reason is not the technology. It is absent governance, poorly defined business value, and what the firm named “agent washing”: an ordinary assistant product, repackaged as an autonomous agent, with no real capability behind it. Among thousands of vendors presenting themselves as agentic, the firm’s own estimate is that around 130 actually are.

And there is a recent detail that sums up the problem in one sentence: between the end of 2024 and the beginning of 2026, the share of enterprise AI tools that went from merely suggesting an action to actually executing it — sending the email, changing the file, moving the money — jumped from 24% to 65%. Companies are giving AI the power to act faster than they are building the controls for that action. It is exactly the point where a badly made pilot stops being waste and starts being risk.

The case of humanoid robots illustrates the same pattern, with a rare piece of public corporate candour. At the start of 2026, asked about his company's humanoid robot factory, the owner himself admitted, on an earnings call, that the line was not yet in “material use” in its own plants, and that the existing units served primarily “for learning, not for productive tasks”. The public target was to produce five thousand units for internal use in 2025; according to reporting with access to internal data, actual production landed in the low hundreds — less than ten per cent of the target.

THE SIDE NOBODY IS LOOKING AT

And at the same time, exactly at the same time, there is a second story — real, measured, quiet.

While one humanoid manufacturer publicly admitted its own line did not work at scale, another completed an eleven-month pilot inside a real BMW plant, with robots physically taking part in the production of more than thirty thousand vehicles. A third completed a full year of commercial pilot inside a Toyota plant, and today moves cargo inside the distribution centres of large logistics operators. A fourth shipped something close to five thousand five hundred units in 2025 alone — the sector's largest production volume in the world, quietly, in Asia.

And the most interesting technical advance of 2026 is not about strength or speed. It is about how these robots learn: instead of being programmed line by line, human operators wear sensor gloves and “show” the task to the robot — in one documented pilot at another plant, twenty human demonstrations were enough for the robot to repeat the task on its own. The technical name is imitation learning. The simple name is: **the machine is learning to do by watching the tacit knowledge of someone who already knows how.**

It is the same principle as Bush, Licklider and Engelbart, sixty years later, actually working — only without the spotlight, because it does not fit in a thirty-second demo.

MORE THAN TWO SIDES

It would be easy to stop here and say: “the honest, scaled pilots work; the hyperbolic, badly governed ones fail.” That is true, but shallow. There are more voices at this table, and none of them is wrong on its own.

There is the side of those who build the models and disagree with each other, publicly, about the path itself. Yann LeCun, one of the most decorated names in the field, has been stating publicly that the current language-model paradigm, on its own, is not the route to an intelligence that genuinely understands the physical world — he argues for different architectures, based on world models. Geoffrey Hinton, another founder of the modern field, has publicly changed position in recent years, saying he fears precisely the speed at which increasingly capable systems are being put into production without equivalent understanding of their limits. Two names from the same Turing Award, disagreeing in public, at the height of their own field — that is not noise. It is the sign that the question is still open, and that whoever claims to have settled it, on any side, is selling confidence that science does not yet have.

There is the side of those who study work, not the machine. Recent studies of corporate restructuring show that two thirds of the

companies that cut staff citing artificial intelligence already report, months later, dealing with the consequences of that decision (*Careerminds*, Feb 2026) — the knowledge that walked out the door was, often, exactly what would have made the AI pilot itself work.

There is the side of those who regulate. The consensus emerging among regulators on different continents is not to ban the technology — it is to demand verification: traceability of decisions, human oversight of high-impact action, and a clear distinction between what a system actually does and what the sales material promises it does.

And there is a side almost never heard in this entire conversation: the side of those who hold the knowledge that the whole debate is trying to capture — the maintenance technician, the agent, the pilot, the sales consultant — whose experience is, at once, the most coveted asset of the moment and the least invited voice at the table where its fate is decided.

THE INVITATION

We live in a strange moment: never has so much been said about artificial intelligence, and rarely has so little been heard from the people who actually know how to operate what it is trying to imitate.

Wanting to be seen is not hypocrisy — recognition is legitimate, and nobody should pretend otherwise. But there is a difference between the stage, which rewards whoever speaks best, and listening, which only exists when someone decides, on purpose, to open space and protect whoever answers. Company, person, society and government face the same task, at different scales: **first, to know themselves honestly; then, to decide what to do with that — before another party decides for them.**

A company that does not know what it knows automates ignorance faster, with more confidence. A person who does not document their own judgement becomes replaceable by the most superficial part of it. A society that debates only in two sides — utopia or collapse — loses the only question that really matters: who verifies, and how. A government that legislates against technology, instead of in favour of transparency about it, protects today's fear and disarms tomorrow's generation.

The invitation of this text is simple to write and hard to do: **audit what you — person, company, country — actually know, before handing that question to someone who will only guess the answer.** Knowledge that was never spoken is the only asset of this decade that no model, however large it grows, can mine on its own.

It is the bet on which we built Zthex — a layer that exists so that an organisation's tacit knowledge is heard, verified and preserved before it becomes a statistic of an abandoned project, or of a layoff whose real reason nobody remembers. It is not the only answer to the moment. It is the one we chose to build, convinced that listening deeply, with respect and without haste, is still the most underrated technology of 2026.

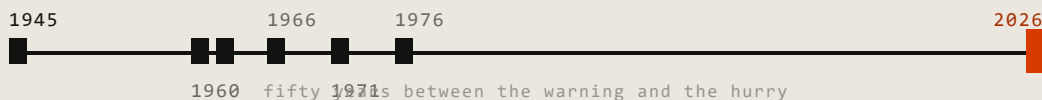
THE FUTURE ALREADY HAPPENED, AND THE NEXT ONE WILL TOO

AlphaFold had no parade. No robot walking on a stage. It had, simply, millions of scientists silently solving problems humanity had carried for decades — and a Swedish committee that realised, two years later, the size of what had happened.

The future always arrives like this: quietly, too early for the headline, too late for anyone who was only watching the demo. Yesterday's has already passed. Today's is passing right now, probably in a conversation nobody recorded, with a person who knows exactly why the process works that way — and whom nobody, yet, has asked.

Z END OF ESSAY

BOX — WHAT WE ALREADY KNEW, AND WHEN



1945 Vannevar Bush, “As We May Think” — the machine should extend human memory, not replace it

1962 Douglas Engelbart, “Augmenting Human Intellect” — technology to amplify collective intelligence

1971 Herbert Simon — a wealth of information creates a poverty of attention

1960 J.C.R. Licklider, “Man-Computer Symbiosis” — partnership, not automation

1966 Michael Polanyi — we know more than we can tell

1976 Joseph Weizenbaum, “Computer Power and Human Reason” — processing is not understanding

BOX — THE TWO SIDES, IN NUMBERS

~80%

of enterprise AI projects do not deliver the expected value (*RAND Corporation*)

42%

of companies abandoned most of their AI initiatives in 2025, against 17% the year before (*S&P Global Market Intelligence*)

~39%

of companies attribute any measurable bottom-line impact to generative AI (*McKinsey*)

>40%

of agentic AI projects should be cancelled by the end of 2027 (*Gartner*)

24% → 65%

share of AI tools that went from suggesting to acting on their own, in sixteen months (*market research, 2026*)

200+ million

protein structures predicted by AlphaFold, now in public use

2024

the year the Nobel Prize in Chemistry recognised this technology — two years after it was already changing science

Editorial content produced by Zthex. Facts about third parties come from public sources cited in the text.

ABOUT ZTHEX

Zthex structures the layer of knowledge no system shows. It talks to the people who do the work, removes identity before recording, and compares what was said with what policy states and what the systems record. Every divergence becomes a finding with a traceable origin.

PRESS

Signed republication and additional reporting:
zthex.com/en/imprensa
info@zthex.com



ZTHEX.COM
point your camera