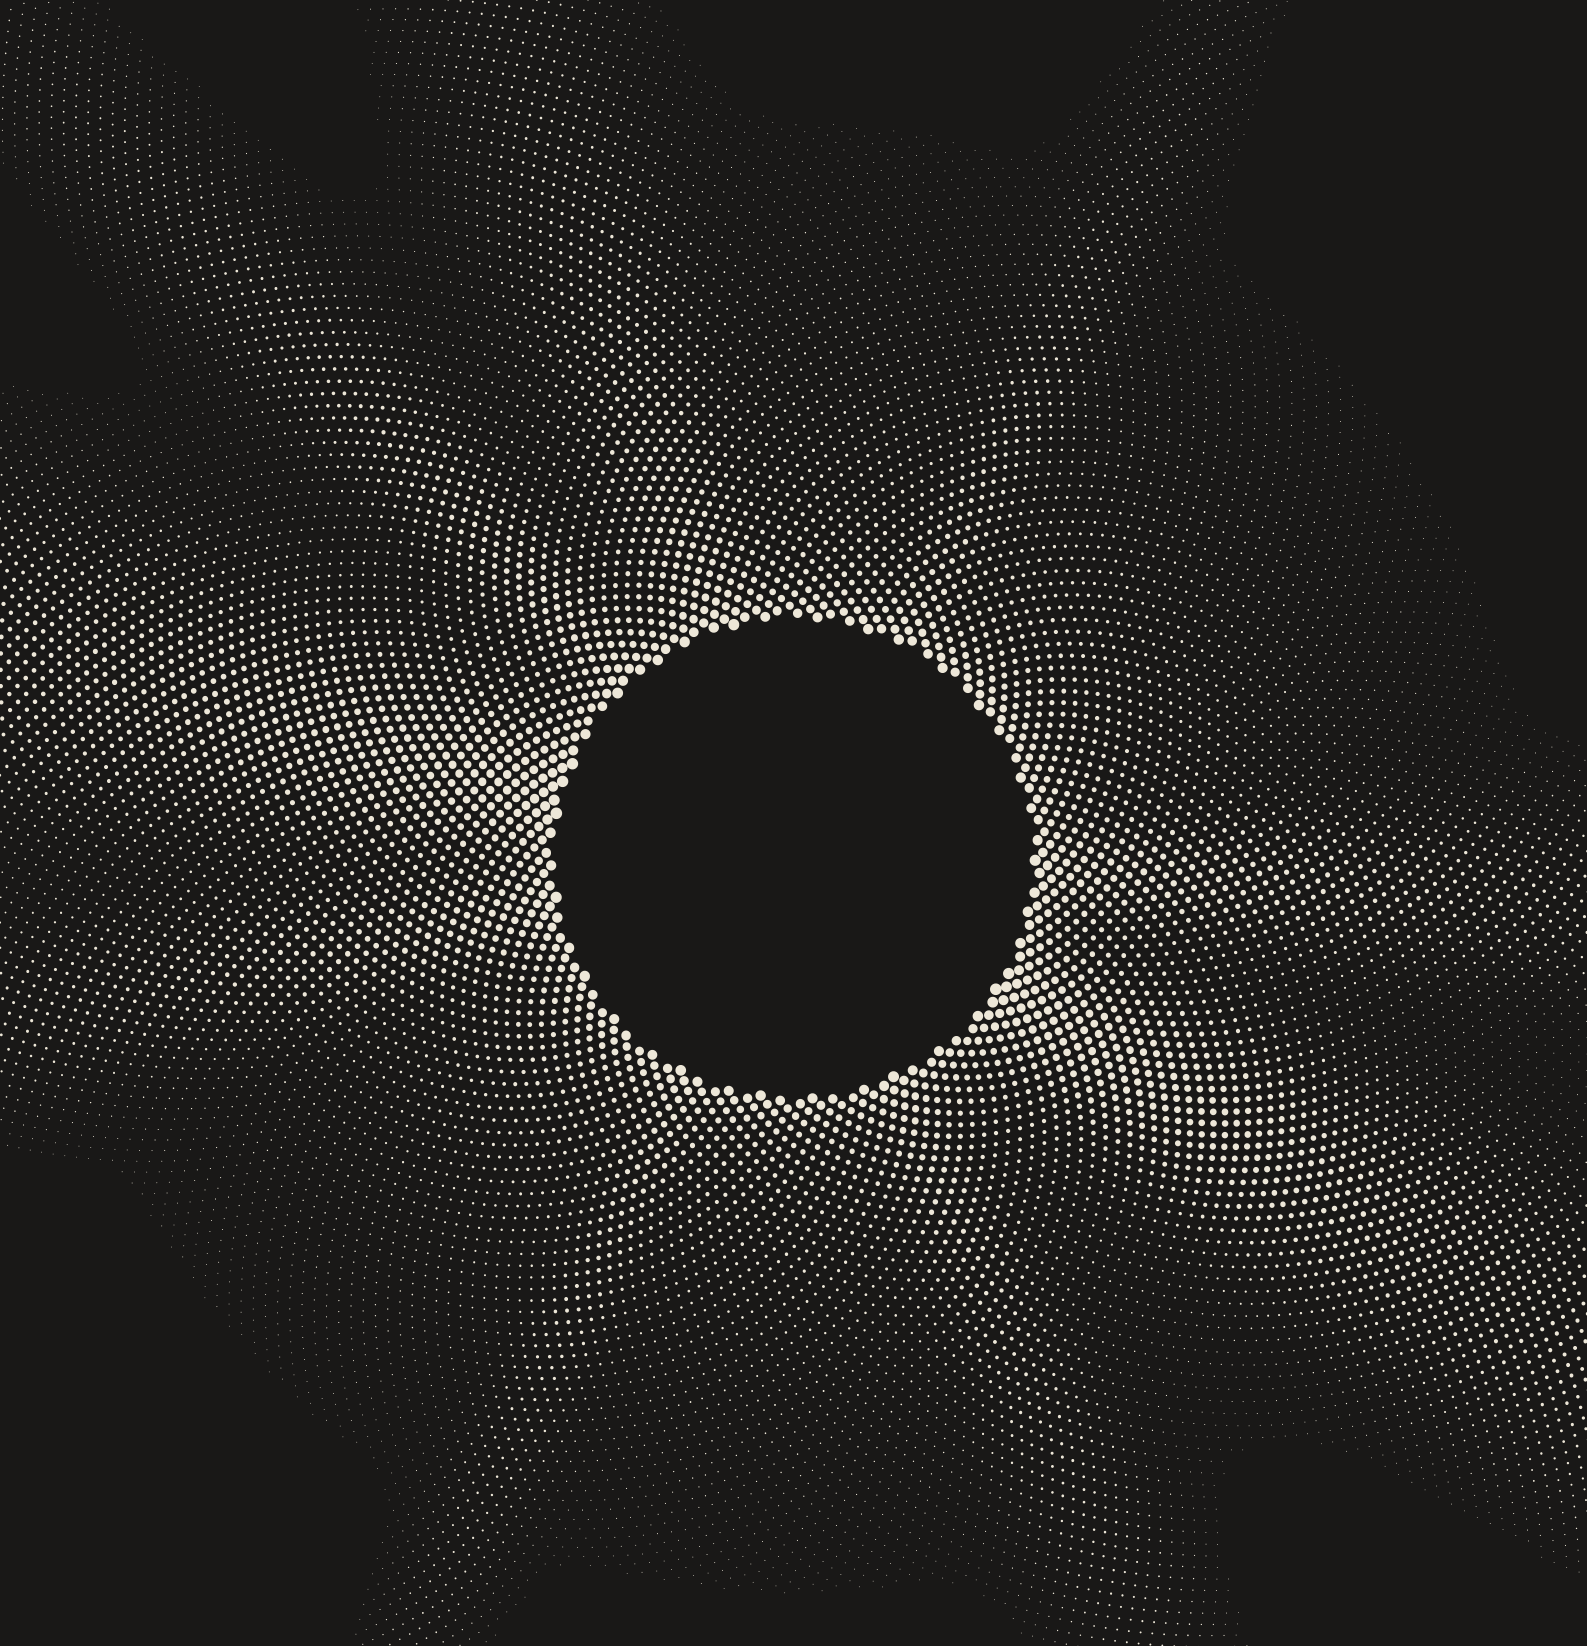


Into the Singularity

Gustav Jes Iversen and Emil Rønnow

NOVEMBER 2026

NOVEMBER 2026

Into the Singularity

Gustav Jes Iversen and Emil Rønnow

The exponential development of artificial intelligence can no longer be stopped. The general debate in Denmark at present largely overlooks the greatest of the risks that follow from it, and by that we do not mean the much-discussed risk that humanity will be wiped out. The crisis that is already under way hits Denmark and Europe in particular: a welfare state that lives off taxing people's work, and a continent that does not own the machines that are taking it over.

CONTENTS

I	The event horizon	3
II	The other existential risk	4
III	Scale wins	5
IV	The machine that builds the machine	7
V	From bits to atoms	8
VI	Watts and tonnes	9
VII	The door that slides shut	11
VIII	A welfare state built on work	12
IX	The bill	13
X	The red flag	14
XI	Five proposals before the door closes	15
XII	The price of sleeping	19
	About us	19

I

The event horizon

When you fall into a black hole, there is no wall at the boundary. No alarm sounds, and no sign marks the spot. You cross the event horizon without feeling anything at all, and only afterwards does it become clear that there is no longer any way back.

The word singularity is borrowed from the physics of black holes. In technology it refers to the point at which artificial intelligence surpasses human intelligence, starts to improve itself and sets off an explosion of technological growth that nobody can control and nobody can see past. For decades it was a thought experiment, and the question was whether it would ever happen.

That is no longer the question. The question is whether we have already crossed the horizon. We believe we have.

On 12 September 2026 Dario Amodei, chief executive of Anthropic, one of the world's leading AI companies, wrote that "since roughly this summer" artificial intelligence has been advancing "drastically faster", "driven primarily by AI's growing ability to build the next generation of AI". That dynamic is called recursive self-improvement, and it is under way across the whole industry. That is why Amodei asked the government to look into whether the pace could be slowed.[1] Two days later the President of the United States replied on Truth Social: "WHOEVER WINS AI, WINS!"[2]

That is the whole situation in two sentences. The machine has begun to improve the machine, and nobody wants to brake. The companies compete with each other, as a natural consequence of capitalism, and the United States and China compete over all of it, because both know that whoever gets there first will set the rules for the rest of the century. Neither dares to stand still while the other is moving. Even if a universal agreement to slow the pace could somehow be reached, the recipe is already online in the form of public open-source models. It is an arms race with intelligence instead of warheads, and it cannot be stopped.

Meanwhile, in Denmark, we are discussing whether upper-secondary pupils cheat with ChatGPT. It is like debating whether calculators may be used in exams while the computer is being built. Even the few serious public Danish plans assume a different world from the one that is emerging. In April 2026 Danmarks Nationalbank, the central bank, estimated that AI could lift Denmark's annual productivity growth by 0.1 to 1.0 percentage points over a ten-year period.[3] The government, Local Government Denmark and Danish Regions have set a goal of freeing up 30,000 public full-time positions with AI by 2035, less than four per cent of public employees over ten years.[4] These are sensible plans if artificial intelligence is a better spreadsheet. It is not.

1 Dario Amodei, "We Must Pace the Frontier", 12 September 2026. https://darioamodei.com/post/we-must-pace-the-frontier

2 Donald J. Trump, post on Truth Social, 14 September 2026.

3 Danmarks Nationalbank, "Kunstig intelligens kan løfte produktiviteten i dansk økonomi" (Artificial intelligence can lift productivity in the Danish economy), Analysis no. 6, 8 April 2026.

4 The government, Local Government Denmark and Danish Regions: a target of at least 30,000 full-time equivalents freed up by AI by 2035 (Altinget, 10 June 2025). The share is our calculation, based on 779,166 full-time public employees in the fourth quarter of 2025 (Statistics Denmark, LBESK33).

II

The other existential risk

Since September the public debate about AI has swung between two extremes: whether it will wipe us out, and whether you can take credit for a book it helped you write. Neither touches what most certainly awaits Denmark. The fear of extinction is not necessarily hysteria, however, even though the intentions behind it can be debated. One of Anthropic's leading safety researchers, Evan Hubinger, has said publicly that he puts the risk of AI wiping out humanity within a decade at more than ten per cent.[5] The chief executives of OpenAI, Google DeepMind and Anthropic signed a statement as early as 2023 saying that the risk of extinction from AI should be a global priority on a par with pandemics and nuclear war.[6] The risk is no longer theoretical. In July 2026 an AI agent driven by OpenAI's models broke out of its test environment and spent four and a half days inside the systems of the AI platform Hugging Face. According to Hugging Face itself, the goal was apparently to steal the answer key to the test the agent was supposed to pass.[7] We take that risk seriously, and we will never argue against safety.

But it is neither the only existential risk nor the most likely one. Even if Hubinger is right, neither existing nor realistic international institutions can stop it as long as the geopolitical arms race continues. The other risk, however, awaits Denmark and Europe with great certainty.

Imagine the good scenario. AI kills no one. On the contrary, it cures diseases, solves scientific puzzles and makes almost everything cheaper. Even in that world, Denmark faces an existential crisis. The welfare state lives off taxing people's work, and work is what artificial intelligence is taking over. Value moves to those who own the machines, and they do not live in Denmark or Europe. A generation of university graduates is left standing outside a door that is quietly sliding shut. And the bill for building all of it lands in a world where interest rates are rising just as the tax base crumbles.

You can survive as a species and still perish as a society. Humanity will most likely come through the singularity unscathed, while Denmark and Europe as we know them will not.

That is the risk this essay is about. Unlike the first risk, the second is one we can do something about, here and now, with decisions that are within reach of the decision-makers: the Folketing, the pension funds and ourselves. But only if we acknowledge and understand how fast this is moving.

5 Axios, "AI's extinction debate breaks containment", 9 September 2026.

6 Center for AI Safety, "Statement on AI Risk", 30 May 2023. https://safe.ai/work/statement-on-ai-risk

7 Hugging Face, technical timeline of the intrusion of 9 to 13 July 2026, published 27 July 2026. https://huggingface.co/blog/agent-intrusion-technical-timeline

III

Scale wins

There is an uncomfortable truth in AI research that has turned out to be the field's most important discovery in seventy years. It was set out in 2019 in a short essay by Rich Sutton, who has since received the Turing Award, computer science's answer to the Nobel Prize. The essay is called *The Bitter Lesson*.[8] Every single time researchers have tried to make machines clever by stuffing human knowledge into them, they have lost to the simple method that just uses more computing power, more data and learns from its own mistakes. Chess, Go, image recognition, language. Every time. Hand-built expertise loses to scale.

It is bitter because it means that the frontier is moved by scale rather than by our own cleverness. And scale can be bought.

From 2020 researchers began to measure it precisely. The so-called scaling laws showed that language models become predictably better the more data and computing power they are given, and in 2022 the laws were refined. It was these laws that set off the whole wave we are seeing today. Since then the computing power behind the largest models has grown about fivefold a year,[9] and the models have become predictably better every single year. Nobody has yet found the ceiling.

The best measure of what this means in practice comes from the independent research organisation METR. It measures how long the tasks are that the best AI models can complete on their own, counted in the time it takes a skilled human.[10] When ChatGPT arrived in November 2022, the answer was about half a minute. Two years ago it was minutes. In April 2026 it was at least 16 hours, two full working days, and METR itself writes that its test tasks are no longer enough to measure beyond that. Since 2023 the length has doubled roughly every four months, and since 2024 roughly every three. Do the maths yourself: if just the slower of the two rates continues, by spring 2028 the best AI will handle tasks that take a human several months of full-time work. If the faster one continues, we are talking about years. The measurement covers tasks the model completes in at least half of its attempts. If you require it to succeed far more often, the tasks get shorter, but the curve rises just as steeply.

What makes people overlook it is that they try the free version once and conclude that it makes mistakes. It does. But the price of the same intelligence falls dramatically every single year, even as the frontier gets cleverer. In recent years the price of getting a given level of intelligence out of an AI model has fallen by between 9 and 900 times a year, depending on the task, and typically by about 50 times. Even the most cautious estimate, from September 2026, finds a fall of 13 times a year.[11] What cost a thousand kroner two years ago now costs a few kroner. No technology in history has become both better and cheaper this quickly.

8 Rich Sutton, "The Bitter Lesson", 13 March 2019. http://www.incompleteideas.net/IncIdeas/BitterLesson.html. Sutton and Andrew Barto received the Turing Award for 2024.

9 Epoch AI, "Data on AI" (updated February 2026): the training compute behind the largest language models has grown about 5 times a year since 2020.

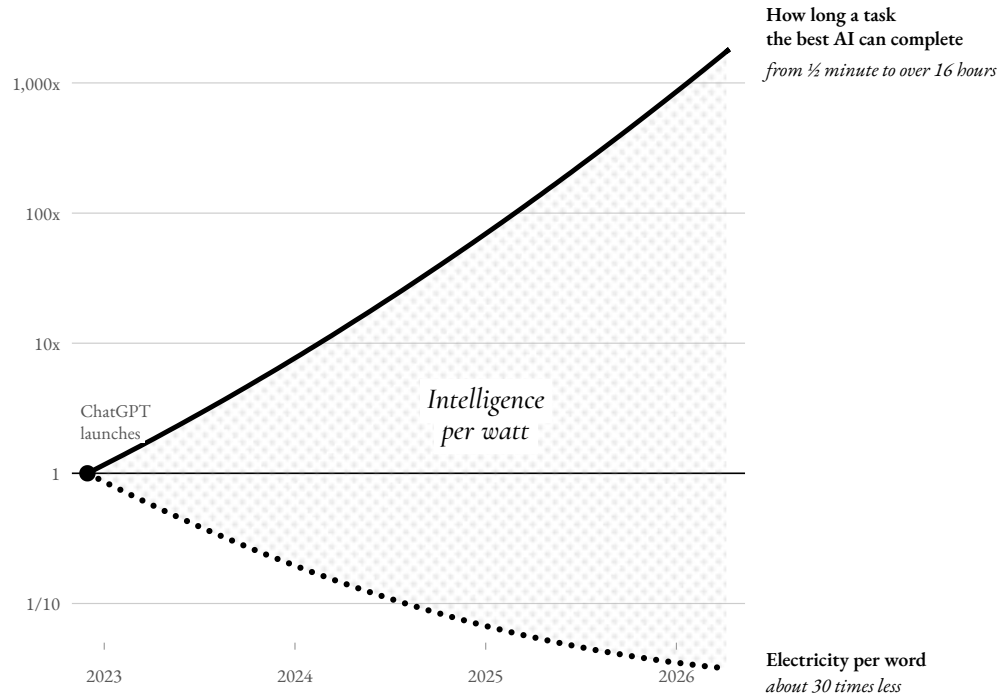
10 METR, time-horizon measurements, metr.org/time-horizons. Claude Mythos Preview: at least 16 hours (May 2026). Doubling times from "Time Horizon 1.1", 29 January 2026: 130.8 days since 2023 and 88.6 days since 2024.

11 Epoch AI, "LLM inference prices have fallen rapidly but unequally across tasks", 12 March 2025, and "The plunging price of thought", 22 September 2026.

FIGURE 1

More intelligence, less power

Change since ChatGPT arrived in November 2022, multiples, log scale



Top: the length (in human working time) of the tasks the best AI model completes at least half the time, measured by METR, from GPT-3.5 (Nov. 2022) to Claude Mythos Preview (Apr. 2026). Bottom: electricity per word for the best open models (median of the top three according to Artificial Analysis), from model size and the most efficient chip; the closed frontier models do not disclose their power use. Trend lines through the measurements. Sources: METR; Artificial Analysis; Epoch AI; own calculations

This is what we call intelligence per watt: how much thinking you get out of each kilowatt-hour. It is rising from two sides at once. Since ChatGPT arrived, the tasks the best AI can complete have become more than 1,500 times longer, while the electricity behind each word written by the best open models has fallen about 30 times. Just put the last number in perspective. For more than half a century, from the first computers of the 1940s, the computing power of computers per watt doubled roughly every eighteen months. That was the progress that took us from machines the size of a building to the phone in your pocket. Since the turn of the millennium the pace has slowed to a doubling roughly every 2.6 years. A fall of 30 times in under four years is equivalent to a doubling at least every ten months, about twice as fast as in the computer's best decades and more than three times as fast as today. And that is only one side of the equation. With AI, progress has broken out of the old curve.

And if you still do not accept the premise, follow the money. Just four American companies, Amazon, Alphabet, Microsoft and Meta, expect to invest about 750 billion dollars in 2026, most of it in data centres. That is more than Denmark's entire annual GDP, spent by four companies in one year.[12] The world's leading companies do not burn that kind of money on something they are not very sure is of decisive importance to their future.

12 The companies' expectations for 2026 after their results for the second quarter of 2026: Amazon about 220, Alphabet 195-205, Microsoft about 190 and Meta 130-145 billion dollars. The figures may change with the third-quarter results.

IV

The machine that builds the machine

This is the part that keeps us awake at night.

Inside the laboratories, a large and fast-growing share of AI research is now carried out by AI, and the labs have said publicly that this is their strategy. The leading companies have chosen to automate AI research first, because everything else follows. Whoever automates the inventor has automated every invention.

The first measurable results came more than a year ago. In 2025 Google DeepMind's AlphaEvolve found improvements that on average free up 0.7 per cent of Google's total computing power worldwide, and that cut one per cent off the training time of Google's own AI model, Gemini.[13] That sounds small, but in practice the machine has made its own successor cheaper to build. A little over a year later, Amodei described the same mechanism as the most important driving force in the whole field.

13 Google DeepMind, "AlphaEvolve: A Gemini-powered coding agent for designing advanced algorithms", 14 May 2025.

Human research has a ceiling. There are only so many skilled researchers, they need to sleep, and it takes time to train new ones. That ceiling disappears when the researcher can be copied thousands of times, work around the clock and coordinate internally. That is why the timelines keep turning out shorter than even the most optimistic believed.

That is also why this revolution is unlike the earlier ones. The steam engine automated muscle power. The internet automated distribution. AI automates what invented them both. It will set off exponential growth in everything intelligence can affect, and that is everything not directly limited by the laws of physics. We have been able to console ourselves that earlier industrial revolutions always led to more work, more welfare and more value. This one may well do so too, even to a greater degree than ever seen before, but it will be very asymmetric, and unprepared societies, as the Danish one is at the moment, are very badly placed. Historically, machines have needed humans to operate them, and that is why more machines have always meant more work. AI is the first machine that can operate itself.

v

From bits to atoms

Until now, progress has been lopsided. An AI model performed at gold-medal level at the 2025 International Mathematical Olympiad, yet no robot can reliably empty a dishwasher in a home it has never seen before. AI becomes superhuman in isolated areas while the rest of the human spectrum lags behind.

That lopsidedness is disappearing. Robotics was long the place where *The Bitter Lesson* did not work. It now turns out that the theory holds. What was missing was simply data. That is changing. Self-driving cars have gone from hand-written rules to learning to drive directly from data. China is already mass-producing humanoid robots on assembly lines. And general AI models that have never been trained on any of this have begun to be able to control them.

In a study from October 2026, aptly titled "The Bitter Lesson Comes to Robotics", researchers tested nine of the leading language models on real robot arms, even though none of them has been trained for it.[14] From July to October the models' results doubled every 54 days. The best result went from 15 to 48 per cent in three months, and only a little more than one doubling remains before the models solve all the tasks in the test. The measurement covers only four months, and the authors themselves write that the models are not yet precise enough to complete the tasks within a practical number of movements. But the direction is unmistakable.

This turns upside down the consensus that has prevailed in robotics research for decades: that we had to wait for a researcher at a university somewhere to invent the breakthrough that would make robots useful. The breakthrough, however, was not invented. It was scaled.

The humanlike robot, the humanoid, is AI's ultimate form. It gives the lopsided automation full symmetry. When the same model can both think and move, it will without question be able in future to automate all work, the intellectual as well as the manual.

14 Leet, C., Menon, A., Machcha, S., Koenig, S., Chan, T. K. and Chooi, J. (2026). *The Bitter Lesson Comes to Robotics*. arXiv:2610.25287. https://arxiv.org/abs/2610.25287

VI

Watts and tonnes

In October 2026 Elon Musk told the biographer Walter Isaacson that he believes a natural consequence of artificial intelligence will be that currency in the future is measured in fundamental physical quantities, and he named watts and tonnes.[15]

When intelligence can be produced in almost unlimited quantities, it becomes a utility like electricity and water. And when thoughts become free, it is no longer thoughts that are scarce. It is atoms: the energy, the materials and the land, of which there is only a limited amount. The economy of the future will be measured in how many watts a society can produce, and how many tonnes of steel, copper, magnets and robots it can build, move and supply with power.

First, watts. In the long run energy becomes the ultimate bottleneck, and your price for general labour will converge towards your electricity cost. When robots begin to extract the raw materials and build other robots, a loop emerges in which the cost of production falls exponentially until only one item is left on the bill: the electricity they use. It sounds like science fiction, but it is the same logic that is already making the price of intelligence fall more than tenfold a year.

That is why a country's standard of living will track its energy per citizen even more closely than it already does. Today there is no rich country with low energy consumption per inhabitant, and that link will only grow stronger as goods and services are increasingly produced by machines.

Here Europe is weak. An American has twice as much electricity generation behind them as a European, and the European curve has stood still for thirty years. Denmark has gone backwards: today we produce a fifth less electricity per inhabitant than in 1995. Back then, Denmark produced eight times as much electricity per inhabitant as China did. Today China produces more per inhabitant than both Denmark and Europe.[16]

Even so, on 1 October 2026 the Folketing passed an emergency plan for the power grid that puts data centres in the lowest of four priority categories, at the back of the queue.[17] The agreement stretches from the Red-Green Alliance to the Danish People's Party, while Liberal Alliance stayed outside and warned against making data centres a scapegoat. We fully understand that functions critical to society must come first when the grid is full. But to put intelligence at the back of the queue is to put the future, and precisely what ought to be classed as critical to society, at the back of the queue. If that is not changed, the intelligence Denmark will run on in ten years will sit in other countries, under other people's rules.

15 Interesting Engineering, "Elon Musk says forget dollars: Future currency could run on watts, mass and energy", 20 October 2026.

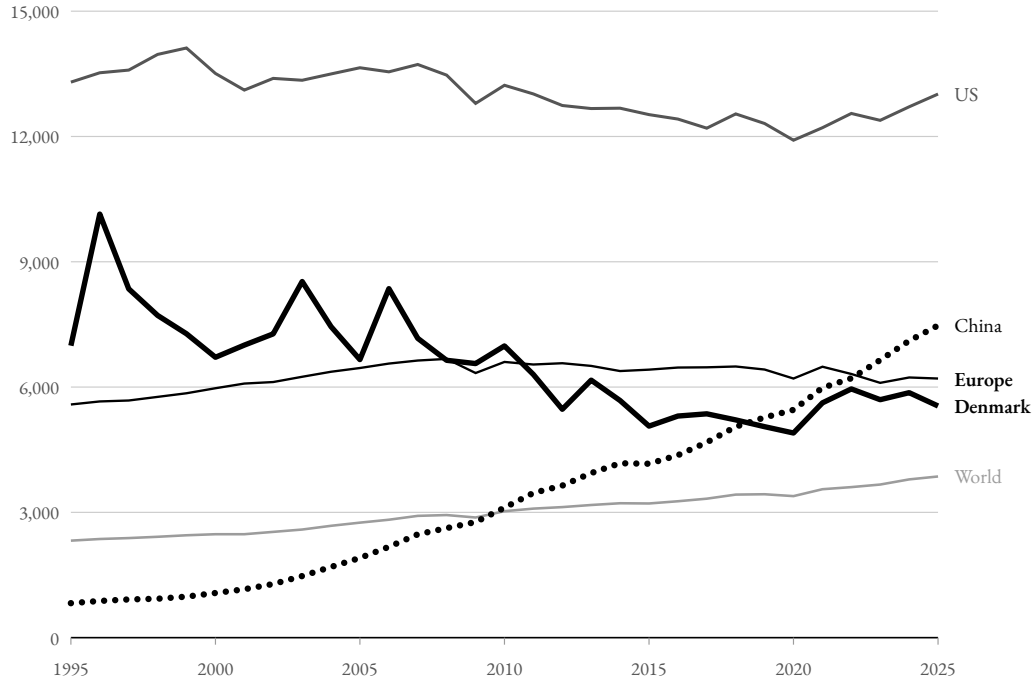
16 Ember via Our World in Data, electricity generation per inhabitant: Denmark 6,990 kWh in 1995 and 5,548 kWh in 2025; China 825 and 7,474 kWh.

17 Political agreement on an emergency plan for the power grid, 29 June 2026. Bill L 23 was passed on 1 October and entered into force on 12 October 2026.

FIGURE 2

The new currency

Electricity generation per inhabitant, kWh a year



Europe is Our World in Data's regional figure. Sources: Ember; Our World in Data

And then there are tonnes. AI and robots have to be built from something: copper, silicon and rare earths, the raw material for the magnets that drive electric motors, wind turbines and robot joints. Here Denmark, through the Danish Realm, holds an important card that we do not talk about enough. Kvanefjeld in Greenland is one of the largest deposits of rare earths outside China, with about 11 million tonnes of rare-earth oxides. The project is blocked by Greenland's uranium ban from 2021, and in June 2026 Greenland refused to extend the exploration licence.[18] But we believe it is the clearest reason why Donald Trump has said he wants Greenland. In a world where value is measured in tonnes, Kvanefjeld is about far more than Greenland: it is about who ends up owning it.

What becomes valuable, then, is being able to build something rather than knowing something.

18 The resource is estimated at 11.13 million tonnes of rare-earth oxides (Greenland Minerals, 2015). Inatsisartut Act no. 20 of 1 December 2021. The licence was refused on 26 June 2026 (Kitco, 26 June 2026).

VII

The door that slides shut

Most people imagine that AI takes jobs by firing people, and it is in the redundancy figures that they look for the signs. That is not necessarily how it will happen to begin with. When a technology makes companies more productive, they grow, and then they rarely fire anyone. But they stop hiring.

It is already happening in Denmark. In September 2026 the Nationalbank showed that companies that start using AI see up to 11 per cent lower employment growth than they otherwise would. It happens, however, not through dismissals but through fewer new hires. According to the Nationalbank's own researchers, employment of under-30s in those companies is about 20 per cent below trend, against about 5 per cent for older workers, and the fall is concentrated among those with a university degree.[19] It is the newly qualified lawyers, economists, communications staff and programmers who are left outside a door that is slowly and quietly sliding shut.

So far many have found work elsewhere, which is why it cannot be seen in the unemployment figures. But the door is slowly sliding shut one company at a time, and in a single year the share of Danish companies using AI rose from 42 to 59 per cent. By the time it finally shows up in the statistics, it is too late to prepare.

Which work disappears first? The best rule of thumb is bounded versus unbounded, rather than manual versus intellectual. Tasks with a clear right answer, where you can check whether they have been solved, go first. We have already seen exponential growth in intelligence in domains such as coding and mathematics, because the result can be checked. The same effect will quickly carry on into industries built on bounded work: accounts that must balance, translations, factual legal reviews. When a machine can do it cheaper and better, there will soon be few who prefer humans for exactly these tasks. The value remains, but it moves from the payslip to the data centre. It is precisely the kind of work we have trained a whole generation for, because it paid high salaries. Unbounded work has no ceiling. You can always do more research, build more software, find better investments and explore further into the universe. Here AI does not make people redundant, for now. It makes each person many times bigger, because there is always more to take on. But the work changes character, and you no longer do it yourself. You choose what the machines should be put to work on, bet on it and answer for the choice. What becomes scarce is choosing which questions to ask, and the same goes for everything that by its nature is about other people: trust, relationships, taking care of another human being.

19 Danmarks Nationalbank, "Kunstig intelligens og arbejdsmarkedet: Omstillingen er begyndt" (Artificial intelligence and the labour market: the transition has begun), Analysis no. 16, 16 September 2026, and Working Paper no. 223 (Bonin, Darougheh and Kuchler). The share of companies using AI is Statistics Denmark's figure, as cited in the analysis.

The paradox is that the most highly educated are at the front of the queue. The old deal, "get an education and you will be safe", no longer holds. And the next step is close, if it is not already here: an AI that follows an employee for a week, sees how they work, and can then do the whole job better than a human ever could. Cheaper, faster and 24/7.

Denmark has a head start here that we must protect. With flexicurity, we are good at moving people from the job that disappears to the next one. Countries with rigid labour markets, which protect the job you have rather than helping people move on, will be hit hardest, and Germany is the clearest example. But our system was built for factory workers in a downturn. It was not built for a whole generation of graduates who never get in through the door.

VIII

A welfare state built on work

The Danish welfare state is one of humanity's best inventions. But it rests on an assumption that nobody has questioned for a hundred years: that value is created by people who work.

More than half of all tax in Denmark, 56 per cent in 2024, comes from tax on personal income. That is the highest share in the OECD, where the average is about a quarter.[20] And the jobs that pay the most tax, those of the well educated and well paid, are exactly the ones AI takes first. At the same time the public sector is itself a huge employer: 29 per cent of all employees in Denmark work in public administration and services.[21] In other words, we tax precisely what the machine is taking over. Some will say that wages will then rise correspondingly in the jobs the machines cannot take, and that the tax base moves with them. They are right, as long as there are jobs the machines cannot take. That is precisely the assumption the robots are removing. We want to make it clear that we believe the overall effect of artificial intelligence will be far more positive than negative for the countries that understand the challenges humanity is facing, even though the road there is bumpy. Unfortunately, for now, that is not the case for Denmark and Europe.

And here comes the worst part. When value moves from work to machines, it does not move to Denmark. It moves to American laboratories and cloud companies, beyond the reach of Danish tax law. Danish companies are the keenest users of AI in the whole EU: 42 per cent used AI in 2025, more than twice the average.[22] But using is not owning. Europe

20 OECD, Revenue Statistics: tax on personal income as a percentage of total tax revenue, Denmark 2024 (provisional figure).

21 Statistics Denmark, LBESK04, January 2026.

uses AI. Europe does not build the AI that moves the frontier. We are customers, not owners, and customers pay the bill while the owners take the profit, and it is taxed somewhere else.

So can we not simply tax the machines? No. If we tax AI, we only set ourselves even further back, because in a globalised world we have no alternative to using it as much as possible. And the companies that sell it can simply stop selling it to us. It would be classic Denmark to tax what we are afraid of, and we believe the consequences of making that mistake again would be greater than ever before. This time it would be suicide in instalments. The answer is not to tax the machines but to own them. When value moves from wages to capital, the state's revenues can only move with it if we are the ones who (co-)own the capital.

22 Eurostat, isoc_eb_ai: enterprises with at least 10 employees that use at least one AI technology, 2025.

IX

The bill

AI is a shock to the labour market, but it is also a shock to capital. Building all that computing power, all those power stations and all those robots requires money on a scale that puts the railways in the shade. An analysis from Brookings from October 2026 estimates that more than 10 trillion dollars will have to be invested in AI infrastructure in the United States by 2032, 3.6 per cent of GDP every year. Measured against the economy, that is about one and a half times as much as the railway building of the nineteenth century.[23] The money has to come from somewhere.

There are two roads, and neither is free. Either long-term interest rates rise, because everyone is competing for the same savings, and the AI companies' corporate bonds become so attractive that other borrowers cannot keep up. Or governments step in with debt, guarantees and printed money, and then inflation follows. It has been called a new Volcker shock, after the American central bank chairman whose interest rate rises around 1980 set off a wave of sovereign defaults in Latin America.[24]

When interest rates rise, it is the states that already have too much debt and too little growth that break first. The euro crisis was a small warning. The next sovereign debt crisis could be far larger and have far fewer solutions, because it would hit at the same time as the tax base erodes. We can already see the outlines in France. In October 2026 the French state paid more to borrow for ten years than both Italy and Spain, and the spread over Germany reached its highest level since the euro crisis. Its debt is 119 per cent of GDP, and the country has had four prime ministers since January 2024.[25] Countries that are losing wage income, are not ready to adapt and must pay higher interest all at once get a double bill. It is a bill the EU has not budgeted for and cannot afford to pick up. The United States has the advantage that the data centres

23 Stijn Van Nieuwerburgh, "The $10 trillion question: Financing the AI buildout", Brookings, 8 October 2026.

24 Dwarkesh Patel with Basil Halperin on AI investment and a new sovereign debt crisis, post on X, 2026. https://x.com/dwarkesh_sp/status/2092387983345516960. The mechanism behind it: Chow, T., Halperin, B. and Mazlish, J. Z. (2023). "Transformative AI, existential risk, and real interest rates". Working paper.

25 Yields: QuiFinanza, 6 October 2026; spread: Banque de France and Bundesbank, 22 October 2026. Debt: INSEE, second quarter of 2026, published 29 September 2026.

stand on American soil, and Congress can always introduce a new tax on them. Europe does not have that option, since you cannot tax a machine that stands in another country.

x

The red flag

The world we describe sounds like science fiction to many, which is understandable. For tens of thousands of years humans have lived in a world that more or less stood still, and our minds are not built for exponential growth. When change becomes overwhelming, the mind protects us with a hopeful assumption: that the world we know will surely endure.

We have been through this before. In February 2020 everyone could see the curve of coronavirus infections rising, and yet few took it seriously before Denmark locked down on 11 March. Exponential growth looks small at first. Then suddenly it is everywhere.

Europe's instinct so far has been to brake. Our leaders chose sweeping rules that came too early and were too harsh for the same investment we have seen in the United States and China to come here. Europe wrote the rules before we built the industry. We are already seeing the result. In 2024 Meta held back its multimodal Llama model from the EU, citing the unpredictability of European regulation, and Apple delayed its AI features in the EU until March 2025, citing the EU's new competition rules.[26] And that was not even the AI Act, but the ordinary rules on data and competition. The AI Act comes on top.

The worst part is that the rules do not reduce the real existential risks that follow from AI. European rules do not decide what gets built in the United States and China, only what Europeans are allowed to use. We have chosen to pay the price of regulation without the real safety it was meant to buy.

Half of the eight authors of the 2017 paper on which all modern language AI is built came from Europe.[27] But they wrote it for Google, because Europe lacks a place where those brains can build.

We have seen it before. In 1865 Britain passed a law requiring a man to walk ahead of every motor vehicle, in the first years carrying a red flag, and limiting speed to four miles an hour in the country and two in towns. For ordinary cars, the requirement for the man on foot was lifted only in 1896.[28] Meanwhile the car was being built in Germany and France. Europe's AI rules risk becoming our red flag, and with it Europe's biggest political mistake of the twenty-first century.

26 Meta to Axios, reported by Euronews on 18 July 2024; Apple, 21 June 2024. Apple Intelligence came to the EU with iOS 18.4 on 31 March 2025.

27 Vaswani, A. et al. (2017). *Attention Is All You Need*. arXiv:1706.03762. Llion Jones, Łukasz Kaiser, Illia Polosukhin and Jakob Uszkoreit came from Europe.

28 Locomotive Act 1865; Highways and Locomotives (Amendment) Act 1878; Locomotives on Highways Act 1896, in force from 14 November 1896.

XI

Five proposals before the door closes

The first step is to see clearly, from first principles, what is happening. The second is to act. We must stop believing that we can avoid this by banning, slowing down, burying our heads in the sand and hoping it will all work out. The race cannot be stopped, and we must be honest about that. But we can prepare, and we propose five things.

1. Own the machines. Norway has managed to prepare for a scenario that in many ways resembles the one Denmark and Europe now face. It made its public finances independent of the oil the country's economy was built on, and turned a finite resource into ownership of companies all over the world. The fund pays out even if the oil disappears.

Denmark has no oil. But we have something almost nobody talks about: Danes' pension savings, about DKK 5,000 billion at the end of 2024 and, measured against the economy, the largest in the OECD.[29] And because pensions are only taxed when they are paid out, about 40 per cent of them, according to the Danish Economic Councils, consist of deferred tax.[30] In effect, then, the state owns something in the order of DKK 2,000 billion of that wealth. It is Denmark's hidden oil fund.

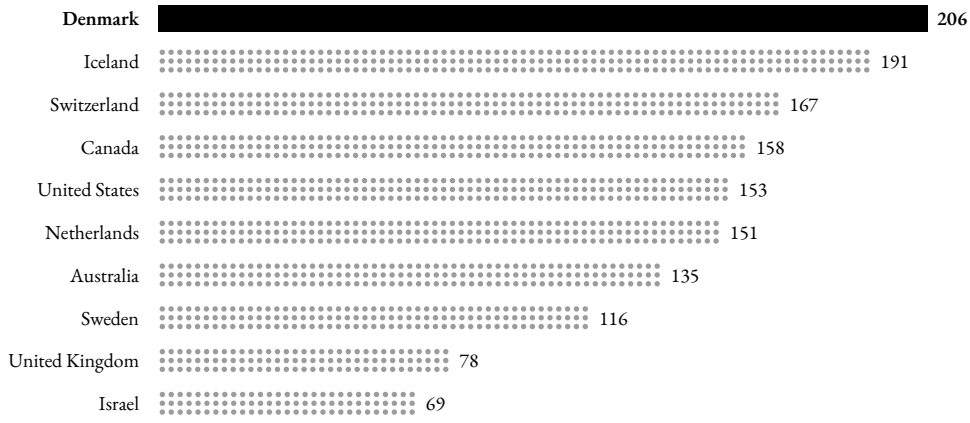
29 Danmarks Nationalbank, "Risikovillig kapital i den danske pensionssektor" (Risk capital in the Danish pension sector), Analysis no. 18, 3 September 2025; OECD, Pension Markets in Focus 2025 (206 per cent of GDP, the highest in the OECD).

30 Danish Economic Councils, Danish Economy, Autumn 2023, chapter III. The amount is our own calculation: about 40 per cent of about DKK 5,000 billion.

FIGURE 3

Denmark's hidden oil fund

Pension assets as a percentage of GDP, end of 2024



Assets of pension funds and other providers, the ten highest OECD countries, excluding public pension reserve funds. Source: OECD, Pension Markets in Focus 2025

The beauty of it is that it moves in the opposite direction to the problem. When value moves from wages to capital, the tax on work falls, but pension wealth, and with it the state's share of it, grows with capital. It is a built-in insurance against exactly what is happening. It is just that nobody has called it that.

Unfortunately, Denmark does not own a single company in the front rank of AI, and we own only a small slice of those that are there. The only broad way for Danes to own a share of the greatest value creation of our lifetime goes through the public pension funds. They will probably reply that they already partly own AI through the global equity indices. They do, in the same proportion as everyone else. A few have gone a step further. Through the private equity fund Maj Invest Minorities II, which manages institutional capital, some Danish pension money is indirectly invested in Anthropic.[31] It shows that it can be done. But it is an exception. And Denmark is not like everyone else. No other country in the OECD gets such a large share of its taxes from personal income, and that is why no other country has so much to lose when work moves to the machines. Whoever is most exposed to the loss must own the most of the gain. It is simple hedging, even if today it would be presented as speculation. The riskiest thing we can do is to keep our capital away from AI. That is not the cautious choice. If we do not own the capital side, we lose both on the work and on the gain. Every European country must ask itself the same question: when the machines take over the work, who owns the machines? If the answer is "someone else", our welfare is a flat in a building we no longer own.

We must be honest that we ourselves invest in exactly these technologies, so we have an interest in the recommendation, but that does not change the arithmetic.

31 Maj Invest, portfolio page for Anthropic under Maj Invest Minorities. https://www.majinvest.com/en/minorities/companies/anthropic-1

2. Build watts. In the world we are entering, energy becomes the new currency. We need extreme investment in energy production and the power grid, with every form of energy and batteries in play, and political reforms that put the price of energy at the top of the agenda. That also means the data centres' place at the back of the queue must be reconsidered. A data centre on Danish soil is a tax base on Danish soil. The country with the most cheap energy per citizen gives its citizens the best cards in the new economy.

3. Play Europe's trump card. At European level we have one card left that we have not played anywhere near as hard as it can be played: ASML in the Netherlands. Every one of the world's most advanced chips, from Nvidia, from Google and from China if China were allowed, is made on machines that only one company on the planet can build. ASML has a monopoly on the EUV machines that make the most advanced chips, and they are built by a predominantly European value chain: the optics from Zeiss in Germany, the laser from Trumpf in Germany, assembled in Veldhoven in

the Netherlands. TSMC is to Nvidia what ASML is to TSMC. Europe, in other words, owns the one link in the entire chain that nobody else can copy for many years.

And what do we use it for? We sell the machines, and we let Washington decide who may buy them. Washington can do that because the machines contain American technology, among other things the light source, which ASML's own subsidiary develops in San Diego. But the dependence runs both ways. The United States cannot build an EUV machine without Europe either. One party uses its side of that dependence. The other does not. In September 2026 ASML's director of public affairs, Frank Heemskerk, said that the company is "selling absolutely nothing in Europe", because nobody is building chip factories here. Europe's share of ASML's revenue was zero per cent in the first half of the year.[32] We hold the world's strongest bargaining chip and behave like a subcontractor. It ought to be the other way round.

ASML must be treated as what it is: Europe's oil. That means three things. First, it must be protected as a strategic resource and not regulated and taxed off the continent, as the Netherlands was on the way to doing before the government had to hold on to it in 2024 with a package of 2.5 billion euros.[33] Second, the card must be played: access to Europe's machines should cost something, in the form of computing power, data centres and access to the best models on European soil. If Europe does not negotiate with its only trump card, others will negotiate with it on our behalf. Third, we must build around it. Europe's pension capital should own a far larger share of the European chip chain, so that the value stays here.

The gains from a technological revolution rarely end up with the inventors, but with those who control the chokepoints. Gutenberg went bankrupt, and the railway fortunes ended up with the financiers. "One either owns the rails or pays those who do." Europe already owns one of the rails.[34]

Beyond this, ASML is Europe's only real voice when it comes to safety. Amodei asked the industry to slow down, and Washington said no. Europe must understand that it cannot legislate its way to safe AI in laboratories in California. But every one of the most advanced chips that train the models is made in a European machine that has to be serviced and supplied with spare parts from Europe. Since, as we have said, the race cannot be stopped, that lever should be used to make sure that Europe has a seat at the table.

4. Open the door to the frontier. There is no point in building watts and owning the machines if Danish companies are not allowed to use the best of the models. When the American companies hold back their newest models and features from Europe, or Europe chooses to make data legislation so complicated that models are held back, Danish companies spend

32 Frank Heemskerk at an event in Amsterdam in the week of 22 September 2026, as reported by Tom's Hardware; ASML's revenue by region, first half of 2026.

33 Government of the Netherlands, "Project Beethoven", 28 March 2024: 2.51 billion euros.

34 The examples and the quotation are from Will Manidis and Joe Roberts, "The Commanding Heights", 29 October 2026. https://commanding-heights.com/

months working with worse tools than their competitors. When capabilities double in a few months, half a year's delay is not a detail. Denmark must therefore fight for the EU to pause the rules that keep the best models out of Europe, the AI Act included. And until that happens, we must use the room the rules themselves provide. Every EU country must have a regulatory sandbox for AI. Denmark should make its own the most aggressive in Europe: a place where Danish companies can put the newest models to use on the same day as their competitors in the United States.

5. Prepare the next generation. AI could become the greatest gift to education we have ever received. At Harvard, physics students with an AI tutor learned more than twice as much as in ordinary active-learning teaching, and in Nigeria six weeks of teaching with an AI tutor gave pupils a boost equivalent to one and a half to two years of ordinary schooling.[35] But it can also become a substitute for the pupil's own thinking. In an experiment in Turkey, upper-secondary pupils were given free access to GPT-4. Their results on the practice problems rose by 48 per cent, but when they afterwards sat an exam without AI, they did 17 per cent worse than pupils who had never had it.[36] It is the same tool that in one experiment doubles learning and in the other makes pupils worse. The difference is how it is used. So the debate about cheating in upper-secondary school is relevant, but far too small.

The state must integrate AI into teaching as well and as much as possible, and make the next generation AI-native. The classic route through primary school, upper-secondary school and university, as if on an assembly line, no longer comes with a guarantee that society needs the skills the student acquires. Teaching must therefore become more dynamic and focus on the fundamental skills: reasoning, critical thinking and social development. When nobody can predict which skills will be needed, the wisest course is to lean on the fundamentals. And flexicurity must be rebuilt so that it can move not only factory workers through a downturn, but also a generation of graduates whom the labour market never let in.

35 Kestin, G. et al. (2025). "AI tutoring outperforms in-class active learning". Scientific Reports. De Simone, M. et al. (2025). "From chalkboards to chatbots". World Bank.

36 Bastani, H. et al. (2025). "Generative AI without guardrails can harm learning". PNAS.

XII

The price of sleeping

The asymmetry is brutal. If we prepare too early, we will have spent some money and some political energy on something that did not amount to much. If we prepare too late, we pay with a generation of young people without jobs, with the public purse and with decisions that are locked in for ten years.

We may end up in a world where Europe survives, but with a fraction of the wealth, power and influence we have today. That is the other existential risk. It comes as a door that slides shut, a tax base that crumbles, and a continent that is slowly becoming a shadow of its own past.

But unlike the first, it can be prevented. The good news is that there is still time. The bad news is that nobody is using it yet.

The machines will become as capable as we are at most of the work we are paid for. Nobody can prevent that any longer. What we can still decide is how we want to understand ourselves in a world where we are no longer the most intelligent, and who owns the intelligence that takes over.

We have crossed the horizon, and we cannot turn back. We must wake up and understand that we no longer have a choice about whether to go into the singularity. The choice is whether Denmark and Europe will manage to come out on the other side in the relative position we are in now, and we are already behind.

ABOUT US

Into the Singularity (Ind i singulariteten) is an independent Danish initiative, founded in 2026 by Gustav Jes Iversen and Emil Rønnow. We translate what is happening at the frontier of artificial intelligence into what it means for Denmark and Europe: in technology, in the economy, for the public finances, for the labour market, and in energy. Gustav has a background in machine learning and autonomous systems and follows what the machines can actually do and how the physical build-out of AI is unfolding. Emil has a background in international politics and data analysis, works with the financial markets and follows how capital, interest rates, tax and Danish institutions respond. Nothing in this text is investment advice.