LOBSTER IN THE LOOP

2026 / 1 · English edition

What does a giant crustacean, the lobster have to do with AI? In principle, not much. And yet the creature once known as the "poor man's chicken" has become the mascot — and the emoji — of autonomous AI agents (and the companies behind them). So much so that in China, giants such as Tencent and Baidu have set up "lobster-rearing" workshops, while local governments have offered funding for "lobster projects."

How is this possible? Perhaps because these animals can display surprisingly complex behaviour on the basis of a nervous system that is simpler, and less human-like, than it appears at first glance. In this analogy, then, the lobster exemplifies a kind of neural functioning very different from our own, while artificial intelligence stands for a system that can produce astonishing results without human-style understanding.

So be it, then — onto the plate with them!



KT studio / Shutterstock / www.thetakeout.com

INDEXES, CASES, AFFAIRS

So, Where Do We Stand? — Stanford AI Index Report 2026

https://hai.stanford.edu/ai-index/2026-ai-index-report

Stanford University's Human-Centered Artificial Intelligence institute has been publishing a report on the state of AI since 2017. Its data come from more than 190 countries. Its method rests on tracking, compiling, processing and visualising data drawn from a range of external and in-house sources. The findings, released annually, summarise indicators measured across technological progress, economic impact and social trends.

This year's report describes AI's breakthrough impact, with the following key messages.

15 KEY TAKEAWAYS — OR, WHERE WE AI-CTUALLY STAND

(The Hungarian heading, "így AI-llunk", puns on "így állunk" — "so this is where we stand.")

1. AI progress hasn't stalled — it's accelerating

According to the report, AI capabilities continue to advance steeply. In 2025, more than 90% of notable frontier models were developed by industry rather than academia. Several models now match or exceed human-level performance on PhD-level science questions, multimodal reasoning and competition mathematics. On the SWE-bench Verified coding benchmark, performance rose in a single year from 60% to nearly 100% of the human baseline. Meanwhile, organisational AI use climbed to 88%, and four out of five university students already use generative AI.

2. The US–China gap in model performance has all but disappeared

The leading American and Chinese models are now very close in performance. Since early 2025 they have traded places at the top several times. In February 2025, DeepSeek-R1 briefly caught up with the best US model, and in March 2026 even Anthropic's flagship model led by just 2.7%. The two countries' strengths differ, however: the US produces more top-tier models and higher-impact patents, while China leads in publication counts, citations, total patent volume and industrial robot installations. It is also worth noting that, per capita, South Korea is the world leader in AI patents.

3. The US dominates data-centre infrastructure, but the chip supply is precarious

The United States hosts the most AI data centres — 5,427 in all, more than ten times as many as any other country — and, with that, it also consumes the most energy. At the same time, the manufacture of leading AI chips depends almost

entirely on a single Taiwanese company, TSMC. This poses a major concentration and supply-security risk to the global AI ecosystem, even though TSMC's US expansion got under way in 2025.

4. AI is at once brilliant and surprisingly weak — the "jagged frontier"

According to the report, AI progress is uneven. A telling example: Gemini Deep Think reached gold-medal level at the International Mathematical Olympiad, while at reading analog clocks even the best model managed only 50.1% accuracy. Certain high-abstraction capabilities are already outstanding, while tasks that are trivial for humans remain weak spots. AI agents have also improved considerably: on the OSWorld benchmark, which measures tasks performed on real operating systems, the success rate rose from 12% to roughly 66% — which still means that about one attempt in three fails.

5. Robots still stumble in real home environments

Although robots are advancing spectacularly in controlled settings, they remain weak in real domestic spaces: according to the report, they achieve only 12% success on household tasks. In software-based manipulation simulations, by contrast, they already reach 89.4%. The gap stems from the unpredictability of the physical world — objects being swapped around, environmental disturbances and changes, complex movements — which is still a serious obstacle for AI-based robotics.

6. Responsible AI is not keeping pace with the growth in capabilities

While model capabilities are improving rapidly, the measurement, reporting and control of safety and responsible operation are not developing at the same tempo. Nearly all companies developing leading models report on capability benchmarks, but for responsible-AI benchmarks the practice is far less consistent. Meanwhile, the number of documented AI incidents jumped from 233 in 2024 to 362 in 2025. And according to some research, improving one responsible-AI dimension of a system — safety, for instance — can degrade other dimensions, such as accuracy.

7. The US dominates investment volume, but its brain-gain pull is fading

US non-governmental AI investment reached $285.9 billion in 2025 — more than 23 times China's $12.4 billion. This figure does not capture China's total AI spending, since state funding barely shows up in market investment statistics. The

US also leads in entrepreneurial activity: 1,953 newly funded AI companies were created in 2025, more than ten times the number in the second-ranked country. A worrying trend, however, is that the number of AI researchers and developers relocating to the US has fallen by 89% since 2017 — with an 80% drop in the last year alone.

8. AI is spreading through the population at historic speed

Generative AI reached 53% of the population within three years — faster than the PC or the internet before it. Adoption varies from country to country and correlates strongly with GDP per capita, but there are outliers: adoption stands at 61% in Singapore and 64% in the United Arab Emirates. The US, by contrast, ranks only 24th at 28.3%. The report puts the estimated annual value of generative AI tools to American consumers at $172 billion by early 2026, with the median user value tripling between 2025 and 2026.

9. Productivity gains are appearing precisely where entry-level jobs are shrinking

Drawing on several studies, the report identifies productivity improvements of 14–26% in areas such as customer service and software development. The positive effect is not universal, though: in tasks requiring stronger judgement or contextual decision-making, the switch to AI leads to weaker results and sometimes to losses. Enterprise adoption of AI agents still typically sits in the single digits across most business functions. It is striking — and revealing about junior-versus-senior employment — that in software development, where AI delivers its most clearly measurable productivity gain, employment of American developers aged 22–25 has fallen by nearly 20% since 2024, while the number of older developers keeps growing.

10. AI's water and energy needs — and with them its environmental burden — are growing

As AI capabilities expand, so does AI's ecological footprint. Grok 4's estimated emissions were equivalent to 72,816 tonnes of CO_2. The electrical capacity of AI data centres rose to 29.6 GW, roughly comparable to New York State's peak load. Water use is significant too: according to the report, the annual water demand of GPT-4.0 alone may exceed the drinking-water needs of 1.2 million people. All of which suggests that AI's future will be shaped not only by model quality but also by the sustainability of energy and water infrastructure.

11. Scientific AI models often beat humans already — but more does not always work better

In some areas, scientific AI models already outperform human experts. On the ChemBench benchmark, for instance, frontier models on average do better than human chemists. Performance is very uneven, though: they score below 20% on astrophysics reproduction and 33% on Earth-observation questions. An important lesson of the report is that the bigger model is not always the better one. A 111-million-parameter protein language model, MSAPairformer, achieved better results on ProteinGym, and a 200-million-parameter genomics model, GPN-Star, outperformed a model nearly 200 times its size. A substantial share of scientific models are not purely industrial developments but the product of cross-sector collaboration.

12. AI has begun to transform medical care, but the evidence on its efficacy is still weak

In 2025, AI tools that automatically document patient encounters spread widely. In several hospital systems, physicians reported spending up to 83% less time on administration — and being less burnt out as a result. That is a serious practical benefit. At the same time, the report notes, reliable data on the broader effects of medical AI remain scarce.

13. Formal education is lagging behind, while people learn AI at every stage of life

In the US, more than 80% of secondary-school and university students already use AI for schoolwork, yet only half of secondary schools and colleges have an AI policy, and a mere 6% of teachers find those policies clear — a sign of serious institutional lag. Outside school, by contrast, AI skills are developing fast: according to the report, AI-engineering knowledge is growing fastest in the United Arab Emirates, Chile and South Africa. Meanwhile, in the US and Canada the number of new AI-related PhDs rose by 22% between 2022 and 2024, and the expertise thus created is showing up primarily in academia, not in industry.

14. AI sovereignty has become a central theme of national policy

More and more countries want AI capabilities of their own, their own infrastructure, and greater control over their AI ecosystems. National AI strategies are spreading, especially in developing economies, and state-backed investment in AI supercomputing is growing in parallel — a sign that countries regard AI as a

strategic resource. Even so, model development remains heavily concentrated in the US and China. Open-source development, however, is partly redistributing participation: the "rest of the world" already outpaces Europe in GitHub contributions and is closing in on the United States, pointing towards linguistically more diverse models and benchmarks. (On AI sovereignty at the organisational level, see the short summary below.)

15. Experts and the public see AI's future differently

There is a serious perception gap between AI experts and the wider public. On AI's impact on work, 73% of experts expect positive consequences, against just 23% of ordinary people. Similar divergences appear in judgements of the impact on the economy and on healthcare. Trust in governments is fragmented as well: of the countries surveyed, the US shows the lowest confidence in its own government's ability to regulate AI — a mere 31%. Globally, more people regard the EU as a trustworthy regulator than the US or China.

ORGANISATIONAL AI SOVEREIGNTY

What CEOs Need to Know About Sovereign AI (MIT Sloan Management Review, 16 June 2026)

- **Sovereign AI:** a regulatory and technological framework that defines:
  - where data may be stored and processed,
  - which infrastructure and which AI models may be used,
  - how AI decisions are supervised.
- **Background:**
  - Nearly 70% of leading AI models originate from the US and China.
  - A growing number of countries are building AI-sovereignty strategies of their own to reduce external dependence.

Accenture survey (2025)

- Sample: 1,928 executives, 28 countries.
- **60%** say sovereign technology solutions have become important because of rising geopolitical risk.
- **Only 15%** treat AI sovereignty as a CEO- or board-level priority.

- **Fewer than 13%** see it as a business-growth opportunity; the majority handle it as a compliance matter.

Key strategic recommendations for CEOs

1. Make AI sovereignty a strategic issue

- It should not be solely an IT or compliance task.
- Its impact can be:
  - geopolitical,
  - capital-allocation-related,
  - supply-chain-related,
  - decisive for long-term competitiveness.

2. Sovereignty is not binary — it is scalable

The degree of control required depends on three factors:

- the industry,
- the country's regulation,
- the AI's use cases (e.g. credit scoring or diagnostics vs. marketing).

3. Build a hybrid AI ecosystem

- Not every company needs to build a full AI stack.
- Potential partners:
  - hyperscalers (AWS, Microsoft, Google, Oracle),
  - regional AI providers,
  - national AI platforms,
  - consortia.

Corporate examples

- **BNP Paribas**
  - Mistral AI as a strategic partner.
  - On-premise AI for European regulatory compliance.
  - Also backs Mistral as a financial investor.
- **AstraZeneca**

    - **In China:** Alibaba Cloud + local Qwen models (for data-localisation reasons).
    - **In other countries:** AWS infrastructure for R&D.
    - The same company uses different infrastructure depending on country and use case.
- **Indosat (Indonesia)**
    - Building a national AI cloud in partnership with NVIDIA and Accenture.
- **Helical**
    - Uses European AI infrastructure (Nebius) to keep data inside the EU.

Regulatory trends

EU

- Implementation of the AI Act.
- The main obligations for high-risk AI systems take effect from **August 2026**.
- In addition:
    - GDPR,
    - NIS2,
    - DORA,
    - the new EU AI Sovereignty package.

Current practice

Companies currently apply sovereignty controls primarily at the following layers:

- **60%** → data,
- **46%** → infrastructure,
- **22%** → AI models.

In sum

- AI sovereignty is **not merely a compliance question** but a **strategic competitive advantage**.
- Successful companies:
    - handle it at CEO level,
    - differentiate by use case and by country,

- build hybrid AI ecosystems drawing on multiple providers.

OUT WITH THOUGHT, IN WITH INTELLIGENCE?

CHINA IS PHASING OUT 12,000 UNIVERSITY PROGRAMMES IN FAVOUR OF AI INTEGRATION STRATEGY

In mid-June 2026, major Chinese and international news outlets made public a higher-education process that China had set in motion five years earlier.

The key development

- Between 2021 and 2026, Chinese universities discontinued 12,000 arts, humanities and language programmes, redirecting student capacity towards technology-focused fields.
- The declared aim is to align higher education with the AI-driven economy China has set as its strategic goal.

Scale and figures

- Ministry of Education data indicate the change has affected nearly one third of China's university programmes, touching millions of students.
- Between 2021 and 2025, meanwhile, 10,200 new undergraduate programmes were launched as part of the same transformation.
- Policymakers and officials increasingly regard arts and humanities programmes as "outdated" and oversaturated (as reported by the South China Morning Post), while viewing new fields such as "embodied intelligence" as a better fit for Beijing's economic goals.

A concrete example

- A recent graduate of the Shanghai University of Science and Technology said the institution had discontinued its product-design programme because of poor labour-market prospects — noting that AI's rapid advance had hit the field hard, since basic tasks such as modelling and rendering can now be performed by AI.

Global context (other AI curriculum reforms)

- India has built AI modules directly into the national school curriculum.
- The United Arab Emirates has launched its National AI Education Initiative.

- Kazakhstan is pursuing a new education strategy, calling the adoption of AI a matter of "national survival."
- Spain has updated its curriculum to emphasise AI literacy as part of the "Digital Spain" strategy.

The British thread

- UK Education Secretary Bridget Phillipson said the department is examining the possibility of a new data-science and AI qualification, which would sit alongside the standard GCSE and A-level exams.
- The largest national curriculum review in a decade will be published next year, with rollout beginning in September 2028.
- Phillipson noted that the curriculum had not been updated in more than a decade, and stressed the need to prepare young people for today's challenges.

...AND YET THEY'RE BOOING IT

At this year's University of Arizona commencement, students booed Eric Schmidt, Google's former CEO, when he spoke about the rise of artificial intelligence and likened today's AI breakthrough to the ascent of computers four decades ago. The speaker responded to the jeering ("I hear you") and called the students' fears "rational" — but urged them to adapt, and to be the ones who shape the technology.

The incident is not unique: commencement speakers who mention AI are drawing this kind of criticism more and more often. Real-estate executive Gloria Caulfield got the same reception from young people at the University of Central Florida, as did Scott Borchetta, head of Big Machine Records, at Middle Tennessee State University.

The phenomenon is reinforced by **Pew Research's "Americans and AI 2026"** survey and analysis, published on 17 June 2026. The research was carried out on a sample of 5,119 people, surveyed between 17 and 23 February 2026. Its headline finding fits the commencement boos: AI use is growing, but perceptions — **even among the young** — are negative.

What the survey found

- Roughly half of American adults now use an AI chatbot, up from a third in 2024; about a quarter of them use one daily.
- The share of ChatGPT users is 44% (34% last year), more than double the 2023 level. It is followed by Gemini (24%), Copilot (17%), Meta AI (14%), then Grok (8%), Claude (6%) and Character.ai (3%).
- Six adults in ten read the AI overviews that appear at the top of search-engine results.

The mood — and the verdict — are not positive

- **More people predict a negative social impact:** 40% say AI will be bad for society over the next 20 years, while far fewer expect a positive effect. Views of their own lives also tilt negative: 31% expect a negative personal impact, and about a quarter a positive one.
- **The young are the most sceptical:** 48% of under-30s say AI is affecting society negatively (versus "only" 37–39% among those aged 30 and over), and 37% say it affects them personally as well.
- **Too fast:** roughly two thirds of respondents (63%) say AI is developing too quickly; a mere 2% say too slowly.
- **Data anxiety:** about seven in ten (71%) believe AI does not adequately protect their personal data.
- **No effective regulation:** 67% of Americans have little or no confidence that the government will regulate AI effectively (up from 62% in 2024).

AI DISCLOSURE

In the production of HAMÁRHOMÁR, artificial-intelligence tools play a strictly supplementary role, confined to defined tasks.

Areas of use: data gathering and the cross-checking of sources; text extraction and summarisation.

Platforms used:

- ChatGPT Pro,
- Claude Sonnet Pro,
- Gemini 3 Pro,
- Genspark Pro,
- Consensus Pro,
- Perplexity,
- DeepL.

The key elements of the editorial process do not rely on artificial-intelligence tools. These include:

- the selection of topics and the generation of ideas,
- the verification and interpretation of sources, texts and data,
- linguistic shaping and adaptation into Hungarian,
- and the final editing and review.