LOBSTER IN THE LOOP

2026 / 2 · English edition

A quarter of a century ago, Charles Stross's short story "Lobsters" appeared in *Asimov's Science Fiction* magazine. The story begins in the near future, in Amsterdam, where Manfred Macx receives a garbled message from uploaded minds — the *Panulirus interruptus* (the neural state vectors of the California spiny lobster) — who are not superintelligences at all. They are rather "beings" exiled to the human-dominated internet, which, in evolutionary terms, could not possibly have been prepared for what they are experiencing. Before being uploaded, after all, they were crustaceans that swallowed their food whole and chewed it in their chitin-lined stomachs. And that is obviously not enough ammunition for finding one's way among spam, cryptocurrency and future-shocked humans. But this time our aim is not to spoil the story any further…
In any case, the story became the first chapter of the author's 2005 novel *Accelerando*, the work of literary fiction that is already about the singularity. Suffice it to say here that by the end of the novel these confused lobster-minds have become post-biological, space-faring beings, and they help humans escape from the "Vile Offspring" — the second-generation artificial intelligences that have devoured most of the Solar System's mass to build technological megastructures.
Stross set "Lobsters" barely two decades ahead in time, around 2013–2020. It is now 2026. Where do we stand? Who is helping whom?...



(www.medium.com)

INDEXES, CASES, AFFAIRS

Yet another measuring tool and ranking about the future — whom are we giving to tomorrow? — QS World Future Skills Index

The World Future Skills Index, published this summer by QS Quacquarelli Symonds, assesses how effectively individual countries are able to develop, align and deploy human skills in a rapidly changing global economy.

It does not examine higher education or the labour market in isolation, but rather how well higher-education systems fit labour-market needs in the age of AI. The index covers 89 economies and, through an analysis of talent supply and talent demand, measures preparedness to seize the opportunities offered by AI; it combines QS's own data (on university performance, jobs and skills, and AI transformation) with internationally recognised indicators, thus serving as a global benchmark of AI readiness in the measurement. The underlying premise is, on the supply side, whether graduates are work-ready (Skills Alignment) and whether the system is capable of producing, at an appropriate scale, capabilities aligned with labour-market needs (Academic Readiness); and, on the demand side, how AI will affect the given economy (Future of Work), and whether the economy is able to absorb and utilise these skills (Economic Transformation).

Key results

In general, AI amplifies the differences that already exist between economies: where higher education, labour-market demand and economic strategy fit well together, the advantage grows further; elsewhere, the skills shortage widens even alongside larger investment. It is telling that even in the first-placed USA, employers report critical shortages in the human skills needed to maximise AI gains, and graduate underemployment is rising — a strong supply of technical skills is no guarantee of demand.

The five main findings are as follows:

1. Critical skills shortages are persistent, but several economies have demonstrated that they can be narrowed through targeted interventions. Globally, human skills are still the most sought-after among employers, yet the shortages in these are at least as large as — and often larger than — those in technical skills; higher education cannot respond as quickly as employer demand moves. For example, India ranks fifth on the

Future of Work indicator (96.0), but only eighteenth in Skills Alignment (82.7) — which clearly shows the gap between labour-market transformation and the output of work-ready graduates. Moreover, in India an average of 15 million young people graduate each year, while the number of jobs advertised for entry-level candidates does not exceed half a million.

2. The labour market's "AI risk–benefit ratio" — the balance between jobs augmented by AI and jobs automated by it — determines long-term competitiveness. The United Kingdom, the USA, Australia and Germany lead the new QS AI Workforce Transformation Index not because they are adopting AI faster than others, but because their economies are concentrated in occupations where AI complements rather than replaces humans.
3. The most competitive economies will be those that successfully align higher education, industry and public policy. The United Kingdom's example: world leader in Academic Readiness (100.0), but its Economic Transformation score (90.2) lags behind — the largest such gap in the top 10.
4. Higher education is not keeping pace with the "seismic" scale and tempo of AI-driven industrial transformation, and employers are frustrated. Economic transformation is proceeding faster than skills development in education.
5. Smaller economies, exploiting their agility and organising education, policy and industry around their strengths, can achieve disproportionately large gains: the Netherlands is 8th, Switzerland 10th and Singapore 12th, with far smaller higher-education systems than the USA, the United Kingdom or Germany.

A related additional layer of analysis is the Balance Index: calculated as the standard deviation of indicator scores, it measures how balanced a skills ecosystem is; the lowest (most balanced) values are achieved by Australia, the USA and Spain, and according to the report, balance in itself may predict adaptability even better than performance does.

The methodology of the survey

The index is built on four equally weighted indicators — each contributing 25% of the final score — and draws on a combination of global databases, QS's own data and recognised international indicators. The four indicators and their internal weightings (based on the report's methodology chapter) are as follows:

– **Skills Alignment (SA, 25%)** — measures how well graduates' skills match employer expectations. It consists of two sub-indicators: 70% comes from skills-gap indicators (shortages reported by employers, by skills cluster: human cognitive skills, human-

centred leadership, entrepreneurial mindset, sustainable and ethical workforce) and 30% from the Human Capital Index as a macro-level foundation.

- **Academic Readiness (AR, 25%)** — the depth, quality and future-readiness orientation of the higher-education system. According to the report it consists of seven sub-indicators: 20% institutional effectiveness and impact (the number and performance of QS-ranked universities), 20% subject competitiveness (QS Subject Rankings), 15–15–15% AI, digital and green educational capacity, plus 15% the quality and number of student cities (QS Best Student Cities).
- **Future of Work (FW, 25%)** — measures how ready a country's labour market is for the AI, digital and green transition. 70% of the weight comes from QS's own AI Workforce Transformation Index, which classifies some 1,869 occupations according to whether AI complements them in a productivity-enhancing way (augmentation) or whether they can be automated because of their routine cognitive/manual tasks; each economy's score reflects the weighted distribution of its workforce across these categories, set against the demand signals of QS's global jobs database of 750 million postings. The remaining 30% reflects current employer demand, with weights of 10–10–10% for demand for AI, digital and green skills.
- **Economic Transformation (ET, 25%)** — the macroeconomic conditions for converting skills into productivity, innovation and growth. It aggregates ten metrics in three components: economic capacity (55%: GDP growth, capital formation, labour productivity), workforce readiness (20%: unemployment, labour-market participation, higher-education enrolment) and future-oriented innovation and sustainability (25%: R&D expenditure, share of young population, environmental performance).

How, where and when the data were collected

- **Who could be included:** 89 economies met the following selection criteria: at least one ranked university in the QS World University Rankings 2027, a minimum employer response threshold, and full macro-data coverage on all ten Economic Transformation metrics. Coverage is global: Europe and Central Asia is represented with 37 economies, East Asia and the Pacific with 14, MENA with 15, Latin America with 13, South Asia and Sub-Saharan Africa with 4 each, and North America with 2.
- **Primary (own) data sources and periods:** the 2021–2025 waves of the QS Employer Reputation Survey, with a total of 92,950 employer responses from around the world; the 2026/2027 editions of the QS World University Rankings, the 2026 QS subject rankings and the 2026 Best Student Cities ranking; the QS AI Workforce Transformation Index

(2025) and QS jobs and skills data from 2023–2025. In the employer survey, employers are asked what importance they attach to certain skills and how satisfied they are with them among graduates recruited from the given universities — the skills gap is essentially the difference between importance and satisfaction.

- **External reference databases:** the World Bank (GDP, capital formation, labour force, Human Capital Index), ILOSTAT (employment, productivity, participation), UNESCO UIS (higher-education enrolment, R&D), and the Yale Environmental Performance Index 2024.
- **Publication:** the index was released to the public on 23 June 2026 by QS Quacquarelli Symonds (the report page itself appeared with a date of 24 June 2026).

The ranking — top 25

The final score is the average of the four 25% pillars; SA = Skills Alignment (the fit between graduate skills and employer needs), AR = Academic Readiness (the strength of the higher-education system), FW = Future of Work (the labour market's readiness for the AI/digital/green transition, chiefly the augmentation–automation ratio), ET = Economic Transformation (the ability to convert skills into economic value).

#	Country	Final	SA	AR	FW	ET
1	United States	99.2	100.0	99.3	100.0	97.4
2	Australia	97.5	96.9	98.6	98.0	96.4
3	United Kingdom	96.6	99.0	100.0	97.0	90.2
4	Germany	95.5	93.9	98.0	99.0	91.2
5	Canada	93.7	94.9	95.9	94.1	89.7
6	South Korea	93.4	95.9	93.2	86.1	98.4
7	China	92.5	98.0	91.8	80.2	100.0
8	Netherlands	91.9	89.8	95.2	88.1	94.3
9	Spain	91.7	90.8	93.9	89.1	92.8
10	Switzerland	91.6	84.8	96.6	90.1	94.8
11	France	91.2	87.8	97.3	95.0	84.6
12	Singapore	91.1	86.8	88.4	91.1	97.9

#	Country	Final	SA	AR	FW	ET
13	India	89.4	82.7	85.7	96.0	93.3
14	Sweden	89.2	79.7	91.2	92.0	93.8
15	Japan	89.0	92.9	87.1	84.1	91.7
16	Taiwan	88.7	77.7	89.1	–	99.4
17	United Arab Emirates	86.5	83.8	85.0	78.2	98.9
18	Hong Kong (SAR)	85.1	88.9	90.5	75.8	–
19	Denmark	84.9	69.6	87.8	85.1	96.9
20	Poland	83.5	71.6	84.4	82.2	95.8
21	Ireland	82.3	85.8	83.0	76.2	84.1
22	Italy	81.0	75.7	94.6	75.3	78.4
23	Israel	80.4	72.6	66.7	93.0	89.2
24	Turkey	79.2	81.7	78.9	65.4	90.7
25	Austria	79.0	68.6	89.8	69.3	88.2

Hungary appears on the list in 41st place with an overall score of 65.1, tied with Peru. It shows itself weakest in SA and ET, and strongest in FW, as can be seen below (SA, rank, AR, rank, FW, rank, ET, rank, Final).

41	Hungary	Europe & Central Asia	High-income	54.4	46	72.8	41	59.5	40	73.8	49	65.1

CAN THE ENGLISH GHOST-WRITER OF OUR AGE, AI, BE RECOGNISED?

THE ECONOMIST TESTED THE STYLE OF THE BIGGEST PLATFORMS

Not-only-but-also? Long dashes (em-dashes)? Professional jargon and nominalisation? George Orwell condemned writers who dress up their vacuous thoughts in convoluted and complicated sentences. He called their style "pretentious diction". It seems he would find plenty to criticise in the way AI writes. Even though

this new intellectual species has no uniform, everywhere-identical manner of writing. But appearing clever and competent, as it turns out, is an important ambition of its.

Source: The Economist: How to Spot AI Writing, 30 July 2026.

Spotting the swineherd… or rather, why is it so difficult?

- After all, the supporting algorithms are already here: the market leader Pangram, for example, claims 99.98% accuracy for itself. The problem, however, is that they can give false positives, and they do not explain why they reached a given conclusion.
- After all, we can search for suspicious words. Unfortunately, however, a single word is not proof in itself. As The Economist's author notes, to claim that a text comes from an LLM because it contains the word "delve" is like saying a text was written by Jane Austen because the expression "imprudence" appears in it.
- After all, we can compare linguistic corpora from before and after the appearance of LLMs. Except that it is not easy to disentangle AI idiosyncrasies from other, natural linguistic trends.
- What is more, according to Karolina Rudnicka, a linguist at the University of Gdańsk, there is no single, uniform "AI style", just as there is no uniform human style. Writers have their individual habits of language use (Emily Dickinson, for instance, was very fond of using the em-dash), and bots may well have such stylistic traits too.

How The Economist investigated

The Economist designed its own comparative study, whose logic was as follows:

1. Establishing the baseline. To identify AI's fingerprints, a human text corpus was needed that was certainly written by humans, distinctive, and familiar to readers. The paper therefore chose its own articles as the reference.

2. Producing the machine corpus. Four leading models were involved: OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini and xAI's Grok. Each was asked to write its own version based on The Economist's articles — and to do so without internet access, so that they could not copy the original. As prompts, they used the AI-generated article summaries that the paper already adds to some of its articles on an experimental basis. Thus the models worked from the same core content as the human authors, but the style of the wording had to be entirely their own.

3. The size of the sample. The resulting human + machine corpus amounted to 55,940 sentences and 1.2 million words.

4. Filtering out bias. There was a risk that the study would identify not AI's but The Economist's stylistic features as deviations. The AI texts were therefore also compared with two further human corpora:

- texts from other outlets (CNN, New York Times, Washington Post), and
- excerpts from bestselling novels published between 1950 and 2022 — the latter ensured comparison with literary prose.

5. The dimensions examined. The analysis was conducted at four linguistic levels: word choice, punctuation, sentence structure and paragraph structure.

Surprising results

The summary conclusion may be this: AI prose is distinguishable, but not on the basis of the features discussed so far. One reason for this may be that the models' style also keeps changing with the updates. But let us look at the current tell-tale traits.

1. Vocabulary: the clichés have been replaced, but the "pretentious diction", the air of pedantry, has remained.

New words are becoming overused in English AI texts. The bots no longer "delve", and there is far less "tapestry". In their place, four new patterns are emerging:

- The piling-up of long, polysyllabic words: "significant", "increasingly", "consequences".
- Rare words appear more often than the human usage average: "interdependence", "reindustrialisation".
- Scientific jargon is characteristic: "parameter", "methodology".
- Nominalisation is standard practice, that is, the use of nouns formed from verbs (e.g. "expand" → "expansion"); this device makes the text more abstract.

All four models display these features, but most strongly of all Gemini and Claude.

2. Punctuation: the em-dash myth has been debunked — the real sign is the absence of punctuation.

Contrary to popular belief, after the most recent updates only Claude uses more long dashes (em-dashes) than human writers do — while ChatGPT uses markedly fewer than any participant in the study. In other words, the em-dash belongs to humans again.

The more reliable tell-tale sign, instead, is that machine text is in fact sparsely punctuated: it contains fewer commas and semicolons than human prose, and there are hardly any parentheses anywhere, for example.

Two explanations can be offered for this: on the one hand, the models write longer sentences, in which segmentation is often replaced by a conjunction. It is no accident that the most overused word has become the coordinating "and". On the other hand, they do not cite experts

and there are not many quotations — whereas quotation marks, interjections and source attributions account for a substantial share of the punctuation of human journalism.

3. Sentence and paragraph structure: AI language is long, monotonous and built on rhetorical templates.

- The bots' sentences are typically long, and their paragraphs are rarely broken up by short, punchy statements — which makes the text monotonous.
- When the models want to liven up their style, they reach for these habitual, clichéd rhetorical structures. The two very typical ones: the "not X but Y" and "not only… but also" turns of phrase, and the rule of three.

Per thousand sentences, ChatGPT and Claude use the most of these structures — more than the other models and more than human writers.

So can we now tell when AI is writing, and when the human scribes?
?..

AI writing is present everywhere: in emails, on LinkedIn, on websites and in chats too. According to some, more than a third of new websites are already drafted by AI. LLMs help students write essays and probably help scientists write papers (cf. the Scientia Acephala blog). The seriousness of the problem of AI writing is also signalled by the fact that a seemingly beautiful literary triumph has ended in a minor scandal: it has been suggested that, instead of the hitherto unknown 62-year-old Trinidadian writer Jamir Nazir, it was in fact AI that wrote the short prose piece Serpent in the Grove, which won first prize at this year's prestigious Commonwealth short story competition. The suspicion is not unfounded: the above-mentioned Pangram identified the piece as 100% synthetic text (The Atlantic). Machine prose is thus drawing ever closer to the human kind, becoming more human-like with every update, since it learns from human texts and polishes itself on the strength of human feedback. It may well be, then, that Pangram can still successfully recognise today's machine texts, but in the future its job may become far harder.

A striking example of rapid learning is ChatGPT's em-dash about-turn: the earlier model used one in almost every sentence, and when the user asked about this, it replied enthusiastically — with an em-dash; today's version soberly says only that it is best used sparingly.

So then..? The closing conclusion may be that we should look for the style of empty pedantry, of straining for abstraction, devoid of wit and surprise — and pay attention not only to what is present in the prose, but also to what is missing from it.

AI DISCLOSURE

In the production of HAMÁRHOMÁR/Lobster in the Loop, 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 Fable 5,
- 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.