

Issue 14, April 2025

Welcome to *The Skinny on AI for Education* newsletter. Discover the latest insights at the intersection of AI and education from Professor Rose Luckin and the EVR Team. From personalised learning to smart classrooms, we decode AI's impact on education. We analyse the news, track developments in AI technology, watch what is happening with regulation and policy and discuss what all of it means for Education. Stay informed, navigate responsibly, and shape the future of learning with The Skinny.

Headlines

- [*The AI race continues with more new models and record-breaking investments, but are we risking “severance”?*](#)
- [*What the Research Says about metacognition and AI in education*](#)
- [*The ‘Skinny Scan’ on What is Happening with AI in Education*](#)
- [*AI News Summary*](#)
 - [*AI in Education*](#)
 - [*AI Employment and the Workforce*](#)
 - [*AI Regulation and Legal Issues*](#)
 - [*AI Development and Industry*](#)
 - [*AI Market and Investment*](#)
- [*Further reading: find out more from these resources*](#)

The AI race continues with more new models and record-breaking investments, but are we risking “severance”?

Welcome to The Skinny on AI in Education. In our [*What the Research Says \(WTRS\)*](#) feature, I bring educators, tech developers and policy makers actionable insights from educational research about metacognition. Fancy a broader view? Our signature **Skinny Scan** takes you on a whistle-stop tour of AI developments reshaping education.

But first I wanted to share some thoughts prompted by what I’ve read this last month...

Is AI ‘severing’ our minds?

I have just finished watching the second season of *Severance* on Apple TV+ and I'm delighted to hear that a third season has been confirmed. For those unfamiliar with the premise, *Severance* is a psychological thriller exploring the concept of extreme work-life separation. The story centres on Mark Scout who, devastated after his wife's death in an accident, becomes a “severed” employee at the mysterious Lumon Industries.

Mark’s decision to become a severed employee involves undergoing a controversial surgical procedure called severance, which divides his memories between his professional and personal lives. The procedure creates two distinct consciousnesses within the same person: an “innie” self who exists only at work with no knowledge of their outside life, and an “outie” self who has no memory of what transpires during work hours. As the narrative unfolds, Mark and his severed colleagues—Irving, Dylan, and Helly—begin to experience disturbing cross-overs, tensions and challenges between

their dual existences, raising profound questions about identity and the nature of consciousness. The series is both intellectually fascinating and deeply unsettling.

We are seeing more reports suggesting that AI doesn't necessarily enhance our cognitive performance—and may actually impair our decision-making, critical thinking, and problem-solving abilities—I worry that we're developing our own version of an "innie" and an "outie": a "dummie" and a "smartie".

Our "dummie innie" self becomes increasingly dependent on AI for efficiency and productivity, gradually disconnecting from deeper engagement with the real world. Meanwhile, our "smartie outie" self remains in touch with reality and, though perhaps slower, maintains a more sophisticated intellectual relationship with the world around us. The concerning parallel with "Severance" emerges when we consider the possibility—hinted at in the conclusion of the second season (spoiler alert)—that our “innie” might eventually dominate what happens to our “outie”—in other words our technology-dependent self might eventually dominate our more deliberate, thoughtful self. This raises questions about the nature of intelligence in the same way that “Severance” raises questions about consciousness.

Are we unwittingly severing parts of our intellectual capabilities, creating a division between convenience and deeper understanding? Let's make sure that we are not...

You can find out more about what prompted my thinking about “Severance” after **What the Research Says**.

[Subscribe to The Skinny](#)

What the Research Says about metacognition and AI in education

In this month's instalment, I focus on the remarkable evolution in the work of a key scholar of metacognition, technology and AI.

In his research over two decades, Roger Azevedo has explored metacognition—the ability to think about one's own thinking—and how technology can enhance these critical processes in education. His work has profound implications for AI developers, educators and policymakers as we navigate an increasingly AI-integrated educational landscape.

What the Research Says: Azevedo's Key Insights and Their Implications for AI in Education

Azevedo's work is particularly notable for tracing the evolution from simple technological tools to sophisticated AI systems that can detect, model and foster metacognitive processes in real-time. While his early research focused primarily on individual learners using basic hypermedia, his recent work embraces multimodal data integration, adaptive AI systems, and serious games as powerful contexts for metacognitive development. The core principles from Azevedo's work stand the test of time, most notably in three ways:

- **The Critical Role of Self-Regulated Learning.** Research consistently demonstrates that teaching students how to learn is as important as what they learn.
- **The Power of Multimodal Assessment and Adaptive Support.** Azevedo's shift from relying on simple think-aloud protocols to integrating multiple data streams (eye tracking, log files, facial expressions, physiological sensors) creates a more comprehensive understanding of learners' metacognitive processes.
- **The Evolution Towards Human-AI Partnership.** Perhaps most significantly, Azevedo's recent work emphasises a co-evolutionary relationship between human learners and AI systems, where both adapt to each other. This represents a profound shift from seeing AI as merely a tool to viewing it as a partner in the learning process.

When it comes to AI in Education: while AI presents exciting possibilities for enhancing metacognition, Azevedo's research suggests that **truly effective systems must go beyond simple content delivery to fostering specific metacognitive processes like judgments of learning, feelings of knowing, content evaluation, and monitoring progress towards goals.**

Key practical takeaways from Azevedo's research

For Educators:

In general teaching, practice benefits can come, for example, from explicitly teaching metacognitive strategies, dedicating class time to show students how to plan, monitor and evaluate their own learning. Targeting specific aspects of metacognition can also be effective, helping students accurately assess what they know, recognise when they have relevant knowledge, critically evaluate information quality, and track their progress.

For AI Developers:

Create AI systems that adjust support based on real-time metacognitive data, providing more challenging content for learners who demonstrate strong metacognitive skills and offering simplified tasks for those who are struggling. Transparent algorithms are also important for helping educators and students understand how AI systems interpret data and make recommendations.

For Educational Leaders and Policymakers:

Invest in training teachers to understand and foster metacognitive development, not just content delivery and develop assessment frameworks that value metacognitive development alongside content knowledge.

The challenge ahead lies not merely in technological innovation but in leveraging AI to create learning environments that foster self-regulated learning while maintaining the irreplaceable role of skilled teachers in guiding metacognitive development.

[For a fuller version of this piece click [here](#).

The 'Skinny Scan' on what is happening with AI in Education....

As promised, here is more about what aspect of my Skinny Scan prompted my thinking about ‘Severance’

Is AI 'Severing' Our Minds?

At the start of this issue of The Skinny I talked about Apple TV+'s Severance and said that my thinking had been prompted by what I had been reading as I've been doing The Skinny Scan. Here are a few of the items that prompted my Severance analogy.

The Decline in Human Cognitive Performance

[An interesting piece by John Burn-Murdoch](#) describes how the cognitive performance in reasoning and problem-solving of both teenagers and adults has declined from a peak in 2012. The OECD's [Programme for International Student Assessment \(PISA\)](#) and the [Programme for the International Assessment of Adult Competencies \(PIAAC\)](#) provides comprehensive data supporting this trend. Particularly concerning is the rise in basic numeracy deficiencies, with 25% of adults in high-income countries now struggling with basic mathematical reasoning.

This timing is significant, as it coincides with our shifting relationship with information—moving from active, self-directed information seeking to passive consumption of infinite content feeds and notifications.

The 'Atrophying' Mind and the “Attention” Economy

[Tej Parikh's examination of "brain capital"](#) explores how technology strains our brain health, capacity, and skills. He identifies concerning trends in brain health drawing on [data from the World Health Organisation](#). Meanwhile, our attention is increasingly fragmented by digital distractions, with daily screen time across devices increasing from 9 hours in 2012 to 11 hours in 2019, according to BOND Internet Trends data referenced in the article.

This fragmentation appears to affect our critical thinking abilities. Studies from the [Brain Health Atlas](#) show concerning trends in healthy years lost due to poor mental wellbeing across different age groups.

Students are using GenAI, but not as we would like

The [HEPI survey of UK undergraduates](#) revealed that 92% are using generative AI this year compared with 66% last year, while 88% have used it in assessments. But, when Cambridge University tested students by having them improve AI-generated content, the results were disappointing—undergraduates made only cosmetic improvements, failing to restructure arguments or apply critical thinking.

Reading this reminded me of recent research by Lee et al. (2025) that identified concerns about how GenAI affects critical thinking. Their survey of 319 knowledge workers found that "GenAI tools reduce the perceived effort of critical thinking while also encouraging over-reliance on AI, with confidence in the tool often diminishing independent problem-solving." Workers with higher confidence in AI were less likely to engage in critical thinking, creating a dangerous feedback loop where diminished skills lead to increased reliance.

The researchers identified a critical shift as "workers shift from task execution to AI oversight, they trade hands-on engagement for the challenge of verifying and editing AI outputs, revealing both the efficiency gains and the risks of diminished critical reflection". This transition to "task stewardship" rather than execution could have long-term implications for skill development.

The 'Severed' Mind

Much like the "innies" in *Severance* who exist only within the confines of Lumon Industries with no knowledge of their outside lives, we may be developing a similar split in our cognitive capabilities. But unlike in *Severance*, where the procedure is deliberate, our cognitive bifurcation is happening gradually and largely unnoticed.

Preventing Complete Severance

With research indicating that those with higher self-confidence—i.e. those who knew they could perform tasks without AI—applied more critical thinking when using these tools, we clearly need to focus on building these foundational capabilities in our students and ourselves.

This aligns with recent empirical evidence from a comprehensive meta-analysis of AI and humans working together. Vaccaro et al. (2024) found that human-AI teams achieve synergy only under specific conditions. Their analysis in *Nature Human Behaviour* of 106 experimental studies revealed that on average, human-AI combinations performed worse than the best performer alone (either humans or AI). They also found that Task Type significantly affects outcomes with decision-making tasks showing performance losses when combining humans and AI, whereas content creation tasks showed performance gains. This suggests that educational strategies should focus on teaching students to use AI as a creative collaborator rather than a decision-making replacement.

Interestingly, when humans outperformed AI in a task, the human-AI combination showed positive effects. However, when AI outperformed humans, the combination typically underperformed compared to AI alone. This suggests that strong human capability is a prerequisite for effective collaboration with AI.

Unlike the fictional employees of Lumon Industries, we still have the choice to maintain integration between our technology-assisted selves and our more deliberate, thoughtful selves. Effective integration will require what Vaccaro et al. (2024) describe as "innovative processes" that appropriately allocate subtasks between humans and AI based on their respective strengths. The question remains whether we'll exercise this choice before our minds become truly "severed".

And now for our signature Skinny Scan...

In a nutshell: AI adoption is rising across sectors, with 88% of UK university students now using generative AI for assessments, and 48% of white-collar workers reporting daily AI usage. However, there's a significant training gap, with only 25% of workers receiving any AI training.

Regulatory and ethical concerns are intensifying around copyright issues, with creative industries warning that UK government plans to weaken copyright laws for AI training pose an "existential threat." Meanwhile, voice-based AI is emerging as a critical frontier, with companies like Meta enhancing voice capabilities to enable more natural interactions.

AI advancement has continued at a rapid pace with multiple new model releases, including OpenAI's GPT-4.5, Anthropic's Claude 3.7 Sonnet, and Elon Musk's Grok 3. Chinese company DeepSeek has emerged as a significant competitor with models that offer comparable performance at lower costs, prompting Western companies to explore "distillation" techniques to create more efficient AI systems.

Investment in AI continues to soar, with Anthropic raising \$3.5 billion at a \$61.5 billion valuation, while OpenAI is reportedly seeking \$40 billion. Tech giants are massively increasing their infrastructure spending, with companies collectively investing hundreds of billions in AI data centres, computing power, and talent acquisition.

- Professor Rose Luckin, April 2025

AI News Summary

AI in Education

Higher education globally continues to face considerable challenges when it comes to the adoption of AI. Institutions must balance embracing AI's potential benefits with thoughtful consideration of its impact on teaching quality, equity, content delivery models, and student wellbeing. Clear institutional strategies and values-based approaches will be essential as the sector navigates this transformative period in education.

The Rise of AI in UK Higher Education

According to a recent Higher Education Policy Institute survey (February 2025), there has been a remarkable surge in AI usage among UK university students. The study reveals that 88% of students now use generative AI tools like ChatGPT for their assessments, a significant increase from 53% last year. Science students appear more likely to embrace these tools than their humanities counterparts, with primary motivations being time-saving and work quality improvement.

The survey also highlights shifting attitudes, with 25% of students now considering it acceptable to include AI-generated text after editing, up from 17% previously. Despite improvements in staff preparedness, many students report receiving "mixed messages" about acceptable AI use. Digital divides persist, with men and students from wealthier backgrounds more likely to be frequent users.

Educational Technology Disruption

A February 2025 report details how educational technology company Chegg has taken legal action against Google's parent company Alphabet. Chegg claims that Google's AI Overviews search tool has significantly reduced its traffic and revenues by keeping

users on Google's site rather than directing them to Chegg's services. This has resulted in a 25% drop in Q4 revenues year-on-year and a 21% decline in service subscribers.

This case illustrates how AI search tools are disrupting traditional traffic patterns and revenue models for educational content providers, forcing institutions to reconsider their digital strategies. Google has defended itself by stating that AI Overviews actually helps discover a greater diversity of sites, highlighting the tension between technology providers and educational content platforms.

Resource Constraints in UK Universities

A March 2025 report highlights concerning trends in UK university assessments. With potentially 10,000 academics facing job losses this year, remaining staff are reporting increased workloads and reduced capacity to provide detailed feedback.

Universities are increasingly moving towards "standardisation, more formulaic and perhaps shorter assessments" with "maybe more automation."

The UK higher education sector has traditionally prided itself on offering tailored feedback and small-group teaching, but resource constraints are pushing many institutions away from this model. There are growing concerns that university administrators implementing cuts may not fully understand these impacts on educational quality and student experience.

Strategic Approaches to AI in Higher Education

George Siemens, a noted educational technology expert, criticises universities for adopting AI without clear vision. According to an EDUCAUSE survey cited by Siemens, only 22% of institutions report having an institutional AI strategy, with the University of Florida initiative standing out as "the most intentional" approach.

Siemens poses fundamental questions for educational leaders: "What should be taught/learned?" and "How should it be taught?" in the AI era. He suggests universities can legitimately choose not to use AI based on their values, emphasising that values-based decision-making is essential.

Educational Models and AI Impact

Siemens highlights promising developments in educational models, including a K-12 approach where "kids crush academics in 2 hours, build life skills through workshops, and thrive beyond the classroom." He also notes reports of an "AI tutor" that has helped "rocket student test scores to top 2% in the country," suggesting significant potential for AI-assisted learning.

However, research cited by Siemens also suggests concerning correlations between AI usage and wellbeing, noting that "higher daily usage—across all modalities and conversation types—correlated with higher loneliness, dependence, and problematic use, and lower socialisation." This raises important questions about AI's broader social impact in educational settings.

AI Employment and the Workforce

The Rise of AI in UK Consulting

According to a February 2025 report in the Financial Times, the UK consulting industry is bouncing back after two years of cutbacks, with artificial intelligence driving the recovery. A recent survey found that 67% of respondents expect AI to be the biggest growth area over the next three years. Lisa Fernihough of KPMG UK confirms that "hiring is back on," particularly for skills in data, cloud services, and technology. The UK consulting market, which contracted by about 3.4% in 2024, is predicted to grow by approximately 5% in 2025. Three-quarters of consulting firms are planning "significant" investments in AI, averaging £1.9 million over the next two years, allowing consultants to focus on strategic thinking while AI handles routine tasks.

Government AI Initiatives

The UK government is developing several AI tools to improve public services. As reported by the Financial Times in February 2025, ministers are creating AI tools to help jobseekers write CVs and covering letters. This initiative aims to free up Jobcentre staff to focus on complex cases and support the Labour government's goal of increasing employment rates from 75% to 80%. The plan could help jobseekers tailor CVs, identify gaps in experience, and practice for interviews, though ironically, the Department for Work and Pensions' own guidelines state that AI-generated application materials are "unacceptable" for departmental job applicants.

In a separate Financial Times article from March 2025, Sir Keir Starmer announced that digitisation of government services could achieve up to £45 billion in annual savings across the UK public sector. The government is developing "Humphrey," an AI tools package to perform administrative tasks such as transcribing meetings and summarising policies. A beta version of the gov.uk app is launching this summer to offer citizens a single point of online access for all state interactions.

Worker Adoption Challenges

Despite the promise of AI to free up time for high-value work, a Financial Times article from March 2025 reveals that workers aren't fully embracing its potential and employers aren't providing adequate support. Demographic factors significantly influence AI adoption: younger workers (18-29) are twice as likely to use AI chatbots frequently compared to those aged 50-64, while education level also impacts usage. Only 25% of white-collar workers have completed any AI training, despite 48% reporting daily AI use. Regular AI users do recognise productivity benefits, with nearly 30% using time saved to check work accuracy and engage in more creative tasks.

Automation in Action: Amazon's Robotics Revolution

The Financial Times reported in March 2025 that Amazon is significantly expanding its robotics operations, having deployed over 750,000 mobile devices and tens of thousands of robotic arms. A new warehouse in Shreveport, Louisiana features 10 times as many robotic pieces of equipment as previous versions, reducing order fulfilment costs by 25%. Morgan Stanley analysts estimate these robotics investments will generate around \$10 billion in annual savings by 2030. While Amazon maintains that humans still have roles for complex tasks, labour organisations warn that robotics-

enabled warehouses increase work pace and cause injuries, with Amazon warehouses recording 30% higher injury rates than the industry average.

Looking Ahead

These developments paint a complex picture of AI's impact on the UK workforce. While consulting firms and government agencies are investing heavily in AI capabilities, there remains a significant gap in worker training and adoption. The Amazon case study demonstrates both the potential efficiency gains and the challenges of integrating automation into existing workplaces. As AI continues to transform various sectors, businesses and policymakers will need to address training gaps and ensure that technological advancement benefits both organisations and their employees.

AI Regulation and Legal Issues

Creative Industries Face Copyright Challenges

According to a Financial Times report from 1 March 2025, the UK's creative sector is raising serious concerns about proposed changes to copyright laws. Eric Fellner, co-chair of Working Title Films and producer of the Bridget Jones films, has warned that government plans to implement an "opt-out" system for AI training data represent an "existential threat" to British creative industries.

The proposal would require artists, authors and companies to actively exclude their work from being used to train AI systems, rather than requiring AI developers to seek permission first. This approach has drawn criticism from numerous British creative figures, including musicians Kate Bush, Damon Albarn and composer Hans Zimmer.

Fellner emphasises the potential economic impact, stating: "If they're going to give away the IP to the tech companies, that is going to make a huge dent in our future ability to operate and our future ability to generate revenue here in the UK."

This comes at a time when the UK film industry is thriving, with production spending reaching £5.6 billion in 2024 – nearly a third higher than 2023 figures and exceeding pre-pandemic levels.

AI Could Speed Up Criminal Justice System

The FT reported on 20 March 2025 that a government-backed review led by Jonathan Fisher KC recommends updating the UK's criminal case disclosure rules to incorporate artificial intelligence. The review found that current rules have failed to keep pace with technological advancements and modern online crime.

The proposed framework would enable law enforcement agencies to use AI to process documents more efficiently, potentially reducing the time cases take to reach trial. This could address a major bottleneck in the UK justice system, where disclosure failures have derailed numerous high-profile criminal trials in recent years.

The scale of the challenge is significant – the average case handled by the Serious Fraud Office involves approximately 5 million documents. However, the review stopped short of recommending a US-style "keys to the warehouse" approach that would give defence

lawyers full access to all prosecution materials, noting this would require substantial increases in state spending.

US Export Controls Stimulate Chinese Chip Industry

On 5 March 2025, the FT reported that US export controls designed to restrict China's access to advanced chips have inadvertently accelerated the development of domestic Chinese alternatives. Huawei, working with Chinese chipmaker SMIC, has made significant progress in chipmaking technology, improving the yield of its latest AI chips to about 40% – double the rate from a year ago.

This improvement has made Huawei's production line profitable for the first time, with Chinese tech giants like Baidu and ByteDance increasingly adopting Huawei's AI chips for deep-learning workloads. The article suggests that without US export bans, Huawei would have continued relying on Taiwan Semiconductor Manufacturing Company (TSMC), with less urgency to innovate domestically.

Despite this progress, challenges remain for Chinese chipmakers, including Nvidia's deeply established software ecosystem and limited access to advanced manufacturing equipment.

Taiwan Leverages Chip Industry for Security Guarantees

A 4 March 2025 FT article highlights how Taiwan is using its semiconductor industry as leverage to secure stronger US security commitments. Taiwan Semiconductor Manufacturing Co (TSMC), which produces 90% of the world's most advanced chips, has pledged to increase its investment in Arizona from \$65 billion to \$165 billion, including a research and development centre.

This comes at a time of increased anxiety in Taiwan following former President Trump's approach to Ukraine, with concerns that US support might waver in the face of Chinese pressure. Taiwan's defence minister Wellington Koo acknowledged the need to balance values with national interests, while former lawmaker Lo Chih-cheng stated bluntly: "We need to put our bargaining chips on the table."

President Lai Ching-te has committed to increasing defence spending to 3% of GDP and boosting investment and procurement from the US. However, some in Taiwan worry that transferring advanced technology to the US might weaken their "silicon shield" – the theory that American dependence on Taiwanese chips helps ensure continued US defence commitments.

Legal Battle Over OpenAI's Corporate Structure Continues

The FT reported on 5 March 2025 that a US federal court has denied Elon Musk's attempt to immediately block OpenAI's conversion from a non-profit to a for-profit entity. Judge Yvonne Gonzalez Rogers dismissed all four arguments underpinning Musk's injunction request, writing that they "failed to meet their burden of proof for the extraordinary relief requested."

While rejecting the immediate injunction, the court has expedited the broader trial to autumn 2025, citing "the public interest at stake and potential for harm if a conversion contrary to law occurred."

Musk, who co-founded OpenAI in 2015 and donated approximately \$45 million before leaving the board in 2018, has accused CEO Sam Altman of "perfidy and deceit [of] Shakespearean proportions." The judge dismissed Musk's claim about anti-competitive behaviour, citing Altman's declaration that he had not told investors backing OpenAI meant they couldn't invest in rivals.

AI Development and Industry

The AI industry continues to evolve at a remarkable pace, with developments in model capabilities, infrastructure, and applications occurring simultaneously. As AI becomes more integrated into our daily lives and critical systems, balancing innovation with accessibility, cost-effectiveness, and regulatory considerations remains a central challenge.

The Democratisation of AI

According to Andrew Ng's "The Batch," AI development is becoming increasingly accessible. Ng pushes back against the notion that people shouldn't learn programming due to AI automation, arguing instead that "as coding becomes easier through technological advancements, more people should learn to code, not fewer." This democratisation is evident in developments like Alibaba's QwQ-32B, a relatively small language model that can run on consumer hardware while rivalling much larger models in reasoning capabilities.

Voice-Based AI Systems Gaining Momentum

Voice interaction is emerging as a crucial frontier in AI development. The Batch reports that "foundation models that directly process audio input and generate audio output are driving growth in voice applications." However, challenges remain: voice models are harder to control, have fewer guardrails, and users are highly sensitive to latency.

This trend is reflected in recent industry moves. According to a Financial Times article (7 March 2025), Mark Zuckerberg is enhancing Meta's voice capabilities as part of their upcoming Llama 4 model. The company is focusing on creating more natural two-way dialogue between users and its voice model, allowing for interruptions rather than rigid question-and-answer formats.

The Race for New AI Models

The AI landscape continues to evolve rapidly with new model releases. The Financial Times (27 February 2025) reports that OpenAI has launched GPT-4.5, boasting a significantly lower hallucination rate of 37% compared to nearly 60% in its predecessor. This launch came amid intense competition, with Anthropic revealing Claude 3.7 Sonnet and Elon Musk's xAI launching Grok 3 in the same timeframe.

The Infrastructure Challenge

As models grow more complex, infrastructure becomes a critical bottleneck. According to the Financial Times (26 February 2025), Nvidia's profits and revenues soared last

quarter as technology companies rushed to build AI infrastructure, with sales increasing 78% year over year to \$39.3 billion.

OpenAI CEO Sam Altman highlighted these challenges, revealing that the company is "out of GPUs," though more are expected soon. To address infrastructure needs, OpenAI has signed a near-\$12 billion, five-year contract with CoreWeave to supply computing power for training and running its AI models (Financial Times, 10 March 2025).

Innovative Applications of AI

AI is being applied in increasingly diverse fields:

- **Weather Forecasting:** The Financial Times (20 March 2025) reports on Aardvark, an AI-powered weather prediction model that can run on desktop computers rather than supercomputers, aiming to "democratise" advanced weather forecasting for countries with fewer resources.
- **Healthcare:** The UK Biobank has launched a major project with 14 pharmaceutical companies to use AI to analyze proteins for better disease treatment (Financial Times, 5 March 2025).
- **Robotics:** Google DeepMind has launched two new AI models called Gemini Robotics and Gemini Robotics-ER, described as a milestone in making general-purpose robots more practical (Financial Times, 12 March 2025).

Cost Challenges and Solutions

According to the Financial Times (2 March 2025), leading AI firms including OpenAI, Microsoft and Meta are turning to "distillation" to create AI models that are cheaper for consumers and businesses to adopt. This technique involves taking a large "teacher" model and using it to generate data that trains a smaller "student" model, transferring knowledge and capabilities at a fraction of the price.

Regional Developments

China's AI sector continues to advance rapidly. The Financial Times (27 February 2025) reports that DeepSeek's advances have sparked a nationwide push in China to deploy its large language models across diverse sectors as Beijing aims to consolidate its gains in generative AI.

Meanwhile, Jack Ma's Alibaba has undergone a strategic transformation focused on AI, with its share price soaring 66% since the start of 2025 (Financial Times, 17 March 2025). The company is now positioned to capitalize on China's AI boom, with its Qwen large language model considered a market leader in China and chosen by Apple to run AI functions on iPhones in China.

Regulatory Challenges

As AI becomes more integrated into critical infrastructure, regulatory concerns are emerging. Microsoft has accused the UK's Competition and Markets Authority of "looking backwards" by ignoring how AI is reshaping the technology industry in its investigation into the UK cloud computing market (Financial Times, 29 February 2025).

AI Market and Investment

The AI Investment Boom Shows No Signs of Slowing

The artificial intelligence sector continues to attract unprecedented levels of investment, with the Financial Times reporting that US startups raised over £30 billion this quarter alone, with an additional £50 billion of fundraising in process. This represents the biggest venture capital splurge in three years (FT, 9 March 2025).

The investment landscape has become notably concentrated, with elite companies receiving the lion's share of funding. Recent notable funding rounds include:

- Anthropic raising £3.5 billion, tripling its valuation to £61.5 billion after launching its Claude 3.7 Sonnet model (FT, 3 March 2025)
- OpenAI reportedly in talks to raise £40 billion at a £260 billion valuation
- Alibaba pledging "aggressive" investment in AI over the next three years, spending more on cloud and AI infrastructure than in the past decade (FT, 20 February 2025)

Industry experts suggest this investment cycle differs from the 2021 peak, with Hemant Taneja of General Catalyst noting these investments are reasonable because AI is "a transformative force" rather than merely speculative.

The Battle of the AI Titans

The competition between major AI companies is intensifying. OpenAI is fighting off an unsolicited £97.4 billion takeover bid from Elon Musk while considering granting special voting rights to its non-profit board to maintain control as it converts to a for-profit structure (FT, 18 February 2025).

Meanwhile, Mira Murati, OpenAI's former chief technology officer, has launched a rival startup called Thinking Machines Lab, which aims to make "AI systems more widely understood, customisable and generally capable." The company has already attracted several senior former OpenAI employees (FT, 18 February 2025).

AI's Growing Energy Demands Spark Nuclear Renaissance

The massive computational requirements of AI systems are driving a resurgence in nuclear power investment. Developers of small modular nuclear reactors (SMRs) have raised at least £1.5 billion in funding over the past year, with investment surge linked to power supply deals with major technology companies seeking energy sources for AI data centres (FT, 19 February 2025).

This trend is further evidenced by eight large energy-intensive companies—including Amazon, Google, and Meta—signing a pledge supporting the goal of tripling nuclear capacity by 2050. Since November 2023, eight new nuclear reactors have been connected to grids worldwide, and construction has begun on 12 more (FT, 12 March 2025).

AI's Expanding Practical Applications

Beyond the investment headlines, AI is demonstrating increasingly practical applications:

- Google has developed an AI "co-scientist" tool to help scientists accelerate biomedical research. Early tests with experts from Stanford University, Imperial College London, and Houston Methodist hospital showed promising results, with the tool reaching in days conclusions that took researchers years to discover (FT, 19 February 2025).
- European startups are focusing on the AI application layer, reimagining workflows across industries. As Taavet Hinrikus, co-founder of Wise, notes: "Every application will be rebuilt with AI and there Europe has a fantastic opportunity to compete" (FT, 13 March 2025).
- Autonomous driving startup Wayve is accelerating its international expansion after raising more than £1 billion from investors. The company claims its system operates on a single AI model that can quickly adapt to different driving environments (FT, 3 March 2025).

The Scale of Investment Reaches New Heights

The sheer scale of AI investments has reached unprecedented levels, with Apple announcing a £500 billion investment in the US over the next four years, matching OpenAI's pledge of the same amount for its "Stargate" data centre project (FT, 25 February 2025).

Tech sector capital expenditure has increased dramatically, from just £23 billion a decade ago to collective spending of £320 billion this fiscal year by Amazon, Meta, Alphabet and Microsoft—primarily on AI initiatives.

Complex Investment Structures Emerge

As the AI sector grows, so does the complexity of its financial backing. Wealthy Chinese investors are quietly funnelling tens of millions of pounds into Elon Musk's private companies (xAI, Neuralink and SpaceX) through special-purpose vehicles designed to shield their identities (FT, 9 March 2025).

Meanwhile, Mark Walter, billionaire CEO of Guggenheim Partners, and Thomas Tull, former owner of Legendary Entertainment, have formed a £40 billion holding company called TWG Global to make large bets on artificial intelligence (FT, 3 March 2025).

Looking Ahead

The AI sector continues to evolve at breathtaking speed, with enormous capital flows supporting increasingly practical applications across industries. While concerns about sustainability, governance, and competition remain, the transformative potential of AI continues to drive unprecedented investment and innovation.