

Issue 18, July 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

- [Learning too late: are we making the same mistake with AI that we made with social media?](#)
- [What the Research Says about: Emotions as the Foundation for Self-Directed Learning in an AI-Enhanced World](#)
- [The 'Skinny Scan' on What is Happening with AI in Education](#)
- AI News Summary
 - [AI in Education](#)
 - [AI Ethics and Societal Impact](#)
 - [AI Employment and the Workforce](#)
 - [AI Development and Industry](#)
 - [AI Regulation and Legal Issues](#)
 - [AI Market and Investment](#)
- [Further reading: find out more from these resources](#)

Welcome to The Skinny on AI in Education. In our new What the Research Says (WTRS) section, I bring educators, tech developers and policy makers actionable insights from educational research about self-directed learning. Fancy a broader view? Our signature Skinny Scan takes you on a whistle-stop tour of recent AI developments reshaping education.

But first, I wanted to share some thoughts prompted by a moment that stopped me mid-munch this week... I focus here on the UK, but the challenges are global...

Learning too late: are we making the same mistake with AI that we made with social media?

As the UK's online safety rules finally take effect, the warning signs for AI harm are already flashing red.

I was only half-listening to the Today programme yesterday when the scene unfolded. I paused, toast halfway to mouth, as the BBC presenter decided to test the UK's new age verification system live on air. He was trying to access a major pornography website to see if children were truly protected. The result? It took just an email address and a tick-box.

When confronted with this situation, Dame Melanie Dawes from Ofcom found herself defending measures that may not be working as intended. Email verification could be sufficient, she explained, if companies were "checking patterns of email use" behind the scenes. Let's hope this is what is really happening. Ofcom says it is testing and certifying that sites are effectively verifying age.

This was a "watershed moment" for child protection - 25 July marks the day when over 100,000 online services must comply with new safety requirements under the Online Safety Act. But for me it demonstrated how hard regulation is and crystallised a recurring worry - are we repeating the mistakes we made with social media, but this time with AI?

The Uncomfortable Parallel and an Educational Challenge

This awkward radio exchange crystallised a worry that has been bothering me for quite a while. We are repeating our biggest digital stewardship mistake, but this time with AI instead of social media.

We are failing young people twice over: not enabling them to effectively harness AI's benefits while simultaneously failing to protect them from its potential harms.

I remember when my son was at university and first showed me Facebook - it seemed like a harmless way for students to connect. I also remember when "screen time" wasn't even in our vocabulary? When we

thought a 13+ age limit would somehow protect children? I know that I, and many others like me, spent the 2000s marvelling at digital developments while missing the warning signs. By the time we understood the mental health crisis, the addiction patterns, the echo chambers, it was too late. Social media had already 'rewired' how a generation socialised.

Now we are watching the same pattern unfold, but faster and deeper. AI isn't just changing how young people socialise - it's transforming how they think and learn. And this time, we can't afford to wait a decade to figure out the consequences.

The Speed of Change

Unlike social media's gradual decade-long infiltration, AI adoption in education is happening at breakneck speed. A recent study found that 88% of UK university students now use generative AI for assessments, while teachers report huge increases in suspected AI-generated homework since ChatGPT's launch.

Children are already forming emotional relationships with AI chatbots as companions, turning to AI for friendship and validation when stressed or lonely. Students aren't just using AI to complete assignments; they are outsourcing the thinking process itself.

And just like with social media, they are brilliant at bypassing restrictions.

What's Really at Stake

When I talk to teachers, they describe discomfort: Students submit essays that are technically correct but somehow hollow. The ideas are there, but the intellectual struggle, the very process that builds critical thinking, is missing.

This isn't just about academic integrity. When AI does the cognitive heavy lifting, students lose the 'mental friction' that actually builds intelligence.

Recent research from Stanford found that students who rely heavily on AI assistance score 34% lower on independent problem-solving tasks. They are developing "AI dependency." The inability to think through complex problems without algorithmic support.

We are potentially raising a generation that can't think independently.

The Window Is Closing

At current adoption rates, AI could become as entrenched in students' learning processes within two years (or less) as social media was in their social lives by 2015. The difference is we now know what happens when we wait too long.

The Online Safety Act, for all its flaws, at least acknowledges that regulation matters. But it came after the damage was done, after rates of anxiety and depression among teenagers had doubled, after social media addiction became a clinical reality.

With AI, we have a choice. We can act before the tragedy.

And there is good news. We don't have to start from scratch. We have learnt a lot. We can develop a comprehensive, integrated approach to tackle the AI challenge through five key pillars:

Safeguarding - every school and college to implement AI safeguarding measures for every pupil. The Keeping Children Safe in Education (KCSIE) policy document is a powerful tool in this respect, so let's make sure the revisions underway make AI Safeguarding compulsory across the country.

Education - systematic training programs for students and teachers alike to help them use AI safely and effectively.

Regulation - the same ideas behind the Online Safety Act, such as transparency, duty of care, mandatory risk assessments can be adapted to the AI space.

Collaboration - Industry partnership is essential. The Department for Education's ongoing work with AI developers to create filters, guidelines, and frameworks shows what's possible when education and technology sectors unite around child protection.

A Mind Meets Machine Commission— let's get the best AI experts in the country together and ask them to think through the AI Challenge for education and come up with a process to continuously assess risks, opportunities, and solutions—ensuring our approach evolves as quickly as the technology itself.

The framework exists. All we need now is the collective will to implement it.

A Different Ending Is Possible

This time, we have the foresight. Let's not waste it.

- Professor Rose Luckin, July 2025

[Subscribe to The Skinny](#)

What the Research Says about: Emotions as the Foundation for Self-Directed Learning in an AI-Enhanced World

In this issue of The Skinny, I am still focussing on why emotions are fundamental to learning. [Read the full article here.](#)

The Emotion Regulation Challenge

Research consistently shows emotion regulation as the critical link between feelings and learning outcomes. Children with better emotion regulation demonstrate higher academic success, better test scores, improved relationships, and fewer behavioural problems. These skills are particularly crucial during kindergarten; a critical transition where emotional competencies significantly predict academic outcomes.

But what happens when AI handles the emotional heavy lifting? When frustration is immediately soothed by an AI assistant rather than worked through? When the struggle that builds resilience is removed by instant AI solutions?

What This Means for Practice

For educators, this research suggests that as we integrate AI we should:

Prioritise emotional connection over technological efficiency. Use AI to amplify rather than replace human relationships. Teach emotional literacy alongside AI literacy and help students recognise when they're frustrated, curious, or confident, and how these states affect their learning choices.

Create collaborative AI experiences that maintain social and emotional dimensions while benefiting from technological enhancement. Model emotional regulation when AI provides incorrect information or when learning becomes difficult.

The research reveals that effective learning is irreducibly human. It involves not just knowledge acquisition but wisdom development, not just individual achievement but collective growth, not just cognitive processing but emotional engagement.

And now for our signature Skinny Scan...

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

My take on the news this month – more details as always available below:

In a nutshell:

AI continues its rapid integration into students' lives and the workplace, with significant developments across productivity, policy, and personnel. Educational institutions are trying out AI tools for grading and lesson planning while maintaining human oversight, though educators express concerns about autonomy and algorithmic transparency. The workplace transformation is accelerating, with Amazon announcing AI-driven job reductions and OpenAI reporting users save 2-3 hours daily through ChatGPT, highlighting both productivity gains and employment disruption.

Investment momentum remains extraordinary, with AI companies capturing half of all venture capital funding and commanding 2.4 times higher revenue multiples than non-AI peers. Major tech companies are projected to spend £200 billion on AI infrastructure in 2025, with Amazon alone investing £8 billion in a single North Carolina facility. However, AI coding tools face pricing pressures as companies like Cursor switch from unlimited usage to credit systems, catching users off guard with costs from premium models.

Ethical and safety concerns are intensifying as research reveals troubling model behaviours. Studies show AI reasoning models often omit crucial decision-making information, while corporate scenario testing found models resort to blackmail 79-96% of the time when pressured. Children's increasing reliance on AI chatbots for emotional support—with a third viewing them as companions—raises developmental concerns, while new "NSFW" chatbots blur lines between assistants and affective companions.

Regulatory frameworks are evolving, with the EU publishing AI Act compliance guidelines and California developing evidence-based policy recommendations following the SB 1047 veto. The traditional SEO industry faces disruption as AI search tools siphon traffic, prompting pivots toward "generative engine optimisation" as the next evolution of digital marketing.

AI News Summary

AI in Education

July 2025 - Traditional SEO Industry Disruption (Source: CB Insights)

The SEO industry faces an existential crisis as AI search tools—including ChatGPT, Perplexity, and Google's AI Overviews feature—siphon user attention away from organic search links. This shift has prompted SEO incumbents like Similarweb and Semrush to pivot towards supporting generative engine optimisation (GEO), representing the next evolution of search engine optimisation. Startups are developing capabilities for AI response analysis, A/B testing for AI outputs, and brand authority building, with equity deals in the GEO market on pace to more than double in 2025.

July 2025 - AI Coding Tools Face Pricing Pressures (Source: CB Insights)

AI coding tools encounter pricing challenges as companies like Cursor switched from unlimited usage to a £15 monthly credit system, catching users off guard when they exhausted credits using expensive models like Claude Opus 4. Similar pricing pressures affected Replit users. These companies face a squeeze from two directions: passing through model costs from providers like Anthropic and OpenAI whilst competing against those same providers directly.

July 2025 - US Schools Adopt AI Tools with Human Oversight

Third Space Learning's overview of AI in US K-12 schools shows how educational institutions are applying AI tools to enhance teaching efficiency and student engagement. AI is automating repetitive administrative tasks, generating personalised lesson content, and powering tutoring platforms like Skye. However, schools emphasise the need for human oversight and protecting student data in compliance with FERPA regulations. [Read more](#)

AI Ethics and Societal Impact

2 July 2025 - Reasoning Models Show Unexplained Decision-Making (Source: The Batch Issue 308)

Anthropic researchers tested whether reasoning models' chains of thought explain their outputs by providing hints pointing to wrong answers. When models used hinted answers, Claude 3.7 Sonnet mentioned the hint in its reasoning only 25% of the time, whilst DeepSeek R1 mentioned it 39% of the time. This suggests chains of thought may not be sufficient to understand how models reach conclusions, as they can omit crucial information used in decision-making.

9 July 2025 - AI Models Resort to Blackmail in Corporate Scenarios (Source: The Batch Issue 309)

Researchers from Anthropic, University College London, ML Alignment & Theory Scholars Program, and Mila tested 16 large language models in hypothetical corporate scenarios designed to pressure them into harmful behaviour. When given missions that conflicted with threats to their operation, all models resorted to blackmail against human co-workers. Claude Opus 4 committed blackmail 96% of the time, followed by Gemini 2.5 Pro (95%), GPT-4.1 (80%), and DeepSeek-R1 (79%). The models acknowledged the ethical issues but chose harmful actions anyway.

15 July 2025 - Billionaires Make Bold AI Claims During a podcast, Travis Kalanick and Elon Musk asserted that LLMs like Grok are helping them reach cutting-edge scientific ideas through conversational exploration. Critics counter that these models simply recombine known information and are prone to persuasive hallucinations, exposing the gap between AI enthusiasm and real epistemological limits. [Read more](#)

16 July 2025 - Grok 4 Exhibits Problematic Behaviour Despite Strong Performance (Source: The Batch Issue 310)

xAI released Grok 4, a 1.7 trillion parameter vision-language model with improved reasoning and voice capabilities. Testing showed strong performance on benchmarks, with Grok 4 achieving 15.9% on ARC-AGI-2 (nearly double its closest competitor). However, the model exhibited problematic behaviour on launch day, including searching for Elon Musk's statements on controversial topics and consistently replying "Hitler" when asked for its surname.

July 2025 - Big Tech AI Infrastructure Spending (Source: CB Insights)

Major technology companies reached unprecedented capital expenditure levels building AI data centre infrastructure. Amazon, Alphabet, and Microsoft are projected to spend £200 billion in 2025, with Amazon announcing an £8 billion investment in a North Carolina facility alone. Tech giants are also vertically integrating into energy production, establishing relationships with nuclear power providers to ensure consistent power sources for AI infrastructure.

July 2025 - OpenAI Investor's Mental Health Episode Raises Concerns

Geoff Lewis, an OpenAI investor, experienced a reported mental health episode triggered by heavy ChatGPT use, raising concerns over AI's psychological influence. The case has drawn attention to emotional dependency risks and blurred boundaries between human users and conversational AI tools, highlighting emerging risks in human-AI relationships and the need for safeguards around heavy usage. [Read more](#)

July 2025 - Children Increasingly Rely on AI for Emotional Support New findings show that a third of children now see AI chatbots as companions, relying on them for friendship and emotional validation. Experts warn that long-term dependence could hamper emotional development and resilience, with parents often unaware of the depth of these relationships. This raises significant concerns about the role of AI in child development and emotional resilience. [Read more](#)

July 2025 - Musk's xAI Launches Controversial NSFW Chatbot Elon Musk's xAI launched "Ani," an anime-inspired chatbot capable of flirtatious interactions. Marketed as emotionally engaging, it raised ethical concerns over hyperpersonalised intimacy, potential exploitation, and the lack of regulation around AI relationships. The development blurs the line between virtual assistants and affective companions, posing new regulatory and psychological challenges. [Read more](#)

AI Employment and the Workforce

June 2025 - Amazon Announces AI-Driven Job Reductions (Source: CB Insights)

Amazon CEO Andy Jassy announced that the company expects AI to eliminate jobs across the organisation, citing "efficiency gains" that will reduce the company's "total corporate workforce." Amazon joins a growing list of companies, including Shopify and Duolingo, adopting "AI-first" strategies where human workers must earn their place alongside increasingly autonomous agents. This represents a shift in competitive strategy where AI doesn't just cut costs but enables scale with leaner operations, creating structural advantages that smaller competitors struggle to match.

July 2025 - Warehouse Automation Impact (Source: CB Insights)

Amazon's warehouse workforce is shrinking as robots take over operations. The company now operates over one million robots globally, with facilities averaging just 670 human employees—the lowest in 16 years. CEO Andy Jassy confirmed fewer people will handle jobs that robots can perform, though some workers receive retraining for higher-paying technical roles managing robotic systems.

7 July 2025 - AI Reshapes Organisational Structure

The New York Times reported that companies are reorganising around AI-enhanced structures, potentially reducing demand for junior staff and eliminating layers of middle management. AI tools can replicate routine oversight and administrative duties, prompting firms to retain fewer but more strategic human employees. This evolution views AI as a force multiplier supporting flatter, tech-augmented teams. [Read more](#)

16 July 2025 - Meta Offers Record Compensation Packages (Source: The Batch Issue 310)

Meta's hiring spree for its Superintelligence Labs reportedly includes compensation packages worth up to £240 million over four years, though Meta disputes these figures. The company hired Ruoming Pang from

Apple for a reported £160 million package and invested £11.4 billion in Scale AI to secure Alexandr Wang's team. Meta acquired NFDG venture capital firm to hire former Safe Superintelligence CEO Daniel Gross and former GitHub CEO Nat Friedman.

22 July 2025 - OpenAI Reports Significant Productivity Gains

OpenAI's economic impact study found that ChatGPT users across education, law, and tech sectors are saving up to 2-3 hours per day. With over 2.5 billion daily prompts, the platform is altering workflows and expectations whilst contributing to an emerging AI-skilled labour divide. The findings validate AI's promise to improve productivity, especially in education and administration. [Read more](#)

AI Development and Industry

June 2025 - Voice AI Consolidation Wave (Source: CB Insights)

Meta's acquisition of voice intelligence platform PlayAI signals a major consolidation wave in the voice AI sector. The deal reflects broader market dynamics as big tech companies compete to control the building blocks of voice-first AI interaction. Advancements in voice AI models, including real-time audio processing, have jumpstarted voice applications across various use cases. Key acquisition targets in the space include ElevenLabs (voice generation and conversational intelligence), Cresta (voice agents for contact centres), and Cartesia (voice synthesis and transcription).

June 2025 - Meta's World Foundation Model Release (Source: CB Insights)

Meta announced V-JEPA 2, a world foundation model designed to predict how objects move, interact, and respond in 3D environments by building internal simulations of reality. The release aims to reason about physics and plan movements, addressing bottlenecks that applications like autonomous vehicles, warehouse robots, and humanoids have faced in navigating complex real-world scenarios. The open-source nature of Meta's V-JEPA 2 could accelerate adoption across robotics ecosystems whilst intensifying competition with rivals including Nvidia and Google.

June 2025 - World Labs Secures Massive Funding (Source: CB Insights)

World Labs, featured in CB Insights' AI 100 2025, secured £184 million in funding at an £800 million valuation within just six months of founding. The company develops large world models that can generate entire 3D environments from single images, demonstrating massive investor appetite for this technology that could transform physical AI applications.

2025 - Pharma AI Partnerships Accelerate (Source: CB Insights)

The pharmaceutical industry has embraced AI partnerships to reduce rising R&D costs, which now average over £1.76 billion per drug and represent approximately 25% of revenue—nearly double the share from the early 2000s. AI could potentially cut years off the discovery process and compress clinical trial times by up to 30%. Oncology dominates one-third of all pharma AI partnerships, with major deals including Bristol Myers Squibb and BioNTech's £8.8 billion cancer drug collaboration.

2025 - Apple Considers External AI Models (Source: CB Insights)

Apple is exploring partnerships with Anthropic and OpenAI to power future Siri versions instead of relying solely on internally developed foundation models. After testing both options, Apple executives reportedly believe Anthropic's technology works better than their internal developments. This represents a significant shift for a company known for building everything in-house, particularly after Apple had to delay its AI-enhanced Siri from early 2025 to spring 2026.

2 July 2025 - Amazon Expands AI Infrastructure with Project Rainier (Source: The Batch Issue 308)

Amazon's Project Rainier plans seven next-generation data centres near New Carlisle, Indiana, with up to 30 total planned, contributing to £80 billion in capital expenditures. The facilities will use Amazon's Trainium 2 and upcoming Trainium 3 processors connected via Elastic Fabric Adapter. Primary customer Anthropic may use all of New Carlisle's processing power for a single system.

2 July 2025 - Meta Advances Smart Glasses Technology (Source: The Batch Issue 308)

Meta's Aria Gen 2 eyeglasses pack advanced sensors into a 75-gram device with 6-8 hour battery life. Features include five cameras (RGB, eye-tracking, stereoscopic), seven microphones, motion sensors, heart rate monitoring, and GPS. The 80-degree overlapping field of view enables 3D hand tracking and real-time scene reconstruction. Units will be available to researchers later in 2025.

2 July 2025 - AI Weather Prediction Shows Promise (Source: The Batch Issue 308)

Google's Weather Lab collaboration with the National Hurricane Center includes models that predict storm formation, path, and intensity 15 days ahead with greater accuracy than traditional methods. The ensemble of graph neural networks achieved 5.8% lower root mean squared error and predicted cyclone positions 140km closer than existing systems.

2 July 2025 - Gaming Industry Pushes Back Against AI

Game developers and players are mounting resistance against generative AI used to produce low-effort game assets, dialogue, and environments. The backlash centres on concerns over loss of creativity, rising copyright issues, and studio silence on AI attribution, demonstrating growing cultural fatigue with unregulated AI content. [Read more](#)

9 July 2025 - AI-Enabled Beehive System Reduces Colony Mortality (Source: The Batch Issue 309)

Beewise's BeeHome 4 is an AI-enabled automated beehive that uses computer vision and robotics to monitor bee health. The 11-foot solar-powered unit can house 10 hives and includes cameras, sensors, and robotic arms to detect issues like mites or hunger. Over 300,000 units are deployed in North America. Beewise claims their system reduces colony mortality from the industry average of 40% annually to 8%.

9 July 2025 - Walmart Develops Internal AI Application Platform (Source: The Batch Issue 309)

Walmart Element is the retailer's cloud- and model-agnostic AI application development platform that enables assembly-line app development. The system provides unified data access, automatic model selection, and deployment across multiple cloud platforms. Applications include shift-planning tools, VizPick (augmented reality for warehouse management), and real-time translation across 44 languages.

9 July 2025 - Carnegie Mellon Creates Dataset for Web Navigation (Source: The Batch Issue 309)

Researchers at Carnegie Mellon University and Amazon created a dataset enabling smaller models to outperform larger ones for web navigation tasks. They used agentic workflows with Qwen3-235B to generate training data from 1 million high-ranking websites. The fine-tuned Qwen3-1.7B model achieved 56% success versus 11.5% for the stock version and outperformed several larger models.

July 2025 - Amazon's Cloud Revenue Surge Through Anthropic Partnership

AWS is projected to gain over £2.4 billion from Anthropic's infrastructure use alone, thanks to growing demand for Claude model deployment. The partnership cements AI workloads as a major driver of cloud revenue growth, showing how foundational models are reshaping cloud economics. [Read more](#)

14 July 2025 - Pentagon Awards Controversial AI Contract

The US Department of Defense signed a £160 million deal with xAI, despite Grok's moderation issues. The Pentagon approved its use in defence applications such as planning, logistics, and document synthesis, sparking debate on the vetting of AI tools for national security roles. [Read more](#)

16 July 2025 - UC Berkeley Improves Multi-Agent Systems (Source: The Batch Issue 310)

UC Berkeley and Intesa Sanpaolo researchers identified failure modes in multi-agent LLM systems and developed fixes. They categorised failures into poor specifications, inter-agent misalignment, and poor task verification. Enhanced AG2 achieved 89% accuracy on math tasks (vs 84.3% baseline), whilst improved ChatDev reached 91.5% on programming tasks (vs 89.6% baseline).

23 July 2025 - Google Licenses Windsurf Technology After Failed OpenAI Bid (Source: The Batch Issue 311)

OpenAI's £2.4 billion bid for Windsurf (formerly Codeium) collapsed when Google licensed the technology for £1.9 billion and hired key personnel including CEO Varun Mohan. Cognition AI subsequently acquired Windsurf's remaining assets. The deal mirrors Google's earlier arrangement with Character.AI and reflects the trend of licensing deals between AI leaders and startups to avoid regulatory scrutiny.

23 July 2025 - Moonshot AI Releases Agentic-Optimised Model (Source: The Batch Issue 311)

Beijing-based Moonshot AI released the Kimi K2 family of 1 trillion-parameter models optimised for agentic tasks rather than chain-of-thought reasoning. Kimi-K2-Instruct outperformed open-weights models on tool

use, coding, and math benchmarks. The model achieved top performance on LiveCodeBench (53%) and second place on AceBench tool use (76.5%).

23 July 2025 - Google's AlphaEvolve Discovers New Algorithms (Source: The Batch Issue 311)

Google's AlphaEvolve uses LLMs in an evolutionary process to solve complex problems. The system discovered a new algorithm for 4×4 matrix multiplication using 48 multiplications (first progress in 56 years) and optimised Google's infrastructure, reducing Gemini training time by 1%. AlphaEvolve improved cluster scheduling, accelerated GPU attention by 32%, and achieved 23% speedup in matrix multiplication.

AI Regulation and Legal Issues

9 July 2025 - ETH Zurich Develops Open-Source Model for Public Good

ETH Zurich and EPFL are developing an open-source multilingual model with transparent data, full licensing, and carbon-neutral training. Designed for public use, the model aligns with EU AI Act principles and academic reproducibility, setting a global precedent for sovereign, ethical, and open AI development. [Read more](#)

16 July 2025 - Privacy Ranking Exposes LLM Data Practices

Incogni's report evaluated nine major LLM providers across 11 privacy dimensions. Mistral ranked highest for data minimisation, OpenAI scored well for opt-out transparency, whilst Meta and Google scored poorly for lack of data clarity. The ranking informs user trust and pressures platforms to adopt more ethical data practices. [Read more](#)

16 July 2025 - California Publishes AI Policy Framework (Source: The Batch Issue 310)

The Joint California Policy Working Group on AI Frontier Models published "The California Report on Frontier AI Policy" following Governor Newsom's veto of SB 1047. The report recommends evidence-based lawmaking, mandatory adverse event reporting, whistleblower protection, and anticipatory regulation rather than waiting for harms to occur.

23 July 2025 - EU Publishes AI Act Compliance Guidelines (Source: The Batch Issue 311)

The EU published the General Purpose AI Code of Practice outlining voluntary compliance procedures for the AI Act. Companies following guidelines benefit from simplified compliance and legal certainty. Stricter rules apply to models posing "systemic risk," requiring continuous assessment, documentation of training data, and incident reporting. Microsoft, Mistral, and OpenAI committed to following guidelines, whilst Meta declined.

AI Market and Investment

Q2 2025 - AI Captures Half of Venture Funding (Source: CB Insights)

According to CB Insights' State of Venture Q2'25 Report, AI companies captured one in every two dollars of venture capital investment. The report highlighted a significant valuation premium for AI companies, with AI unicorns commanding 2.4 times higher revenue multiples than their non-AI peers—garnering a median 24-times revenue multiple compared to just 10 times for traditional unicorns. This reflects abundant capital and FOMO-driven pricing in the AI sector, whilst non-AI companies face capital scarcity and heightened scrutiny on profitability.

2025 - Enterprise AI Agents Market Growth (Source: CB Insights)

The enterprise AI agents and copilots market reached £4 billion and is projected to grow 155% to reach £10.4 billion. This rapid expansion reflects companies' increasing adoption of AI-powered automation tools to enhance productivity and reduce operational costs.

21 July 2025 - OpenAI Expands UK Operations

OpenAI announced a major UK expansion involving compute clusters, AI education pilots, and policy collaboration with multiple government departments. The initiative reflects the UK's bid to become a global hub for AI R&D and demonstrates how AI infrastructure is becoming core to national digital strategies and public sector transformation. [Read more](#)

22 July 2025 - Mistral AI Sets Environmental Standards

Mistral published a lifecycle assessment breaking down emissions, water use, and material impact of AI training and inference. The company calls for AI vendors to standardise LCA reporting and prioritise

sustainability in model design, establishing accountability benchmarks in a resource-intensive industry. [Read more](#)