

Issue 17, June 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

- [Plastic Intelligence: the Real Cost of AI Playtime](#)
- ['All You Need is Love' - 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 and Cybersecurity](#)
 - [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 what I've been doing and reading over this last month...

Plastic Intelligence: the Real Cost of AI Playtime

Am I Just Becoming Little Miss Grumpy?

I find myself feeling increasingly ratty with the AI playbook. Am I just becoming curmudgeonly, or are my concerns justified?

My latest crotchiness was precipitated by my reading the announcement of Mattel's partnership with OpenAI to create AI-powered toys. This is a bad idea. While they promise "age-appropriate play experiences," how is this possible when we lack representative datasets about children? Without comprehensive, unbiased data, we risk creating AI toys that reflect existing biases, potentially shaping young minds in ways we don't comprehend.

Just imagine the horror of an AI toy responding differently based on socioeconomic linguistic patterns. Or one that reinforces gender stereotypes perhaps? I would not want these for any of my grandchildren or anyone else's either. These are not distant possibilities, they are probable outcomes when AI systems trained on adult-centric, historically biased datasets are adapted for children.

But who is going to stand up and stop the horror unfolding?

There are other ingredients beyond 'Barbie going to school' that add to my tetchiness. For example, recent MIT research provides neuroscience evidence to back up what many of us have been worried about for a while. Using EEG technology, researchers found that ChatGPT users showed the lowest brain engagement and "consistently underperformed." Over time, users became progressively lazier. When asked to work without AI assistance, they remembered little of their own work. It's a 'brains off, AI on' future ahead unless we are very careful!

But I ask again - who is going to stand up and stop the horror unfolding?

Just as we need the education sector to raise the game, they are sadly absent. Assessment drives the way young people are equipped for their future. And yet, higher education is failing to develop assessment methods that are fit for purpose in an AI driven world. Schools cannot do this; they must adhere to their system's exams. But higher education can change assessment, and yet innovation in this area is lamentably lacking.

Human learning evolved as cooperative and emotionally engaged. When AI systems replace these deeply human processes with isolated, efficiency-focused interactions, we risk severing the emotional and social roots that make learning transformative rather than merely transactional.

The decisions we make today about AI will shape human cognitive development for generations. But we are not listening to what is really happening as the AI playbook continues to roll and we are not acting wisely to build the next generation of super smart humans.

Perhaps being a little grumpy is entirely justified.

- Professor Rose Luckin, June 2025

[Subscribe to The Skinny](#)

'All You Need is Love' - What the Research Says about: Emotions as the Foundation for Self-Directed Learning in an AI-Enhanced World

My latest What the Research Says synthesis discusses why emotions aren't peripheral to learning; they are fundamental. Self-directed learning, a capability that we know is increasingly the important for students to thrive in the AI world depends on resilience, hope, and navigating frustration. Yet AI use may create the opposite: reduced engagement, weakened memory formation, and progressive cognitive dependency. [Read the full article here.](#)

I'm often asked to talk about the importance of human intelligence in an AI world. Indeed, in the last few months I've been asked to do this by both the Ministry of Education in Finland and the Ministry of Education in Singapore. Whenever I'm asked to focus on human intelligence, I try to draw out the differences between human intelligence and AI and focus on our unique human strengths too. Two important areas of difference are: our emotional connection to the world and to other people, and our ability to reflect on our own cognition and engage in reflection and self-directed learning. And both these areas of difference are extremely important for learning.

Research reveals that as we integrate AI into education, we must ensure technology amplifies rather than diminishes the emotional foundations of learning. Success requires:

- *Educators who balance technological tools with human connection and emotional support*
- *Technology designers who centre emotional awareness and social connection in their systems*
- *Policymakers who protect the fundamentally human aspects of learning whilst enabling innovation*
-

In our rush to harness AI's power for education, we must not lose sight of what makes learning meaningful: emotional engagement, social connection, and shared purpose. The most effective learning environments will use AI to enhance not replace these fundamentally human elements.

The path forward requires courage and persistence: the courage to maintain human relationships in an age of artificial ones, to value emotional development alongside academic achievement, and the persistence to insist that technology serve rather than subvert our deepest educational values. It requires wisdom to discern when AI enhances learning and when it diminishes it, when to provide support and when to allow struggle, when to connect digitally and when to insist on face-to-face presence.

Most importantly, it requires love, the love that motivated humans to share knowledge across generations, that drives teachers to nurture their students' growth, and that inspires learners to persist despite challenges. No AI, no matter how sophisticated, can replicate this fundamentally human capacity for care.

We have research to guide us; the challenge now is to listen with both our minds and our hearts.

*The 'Skinny Scan' on what is happening with AI in Education....
My take on the news this month – more details as always available below:*

AI's Impact on Education and the Workplace: Key Insights from June 2025

The latest collection of AI news reveals a technology landscape that continues to undergo rapid transformation, with profound implications for how we learn, work, and interact professionally. Here are the key developments that matter most for education and workplace transformation.

Education Embraces AI Integration

Business schools are leading an educational revolution, with nearly 80% now incorporating AI into their programmes. Rather than training students to code, institutions are focusing on creating "strategically fluent leaders" who can work alongside AI systems. Schools like Essec in Paris are demystifying AI's decision-making processes, whilst Trinity Business School in Dublin combines digital fluency with human-centred perspectives.

The transformation extends beyond curriculum content to delivery methods. Language learning platform Duolingo has dramatically demonstrated AI's potential, creating 148 new courses in under a year—more than doubling their catalogue and surpassing their previous 12-year output of 100 courses. This showcases how AI can help educational organisations scale their offerings whilst potentially improving quality.

Apple's announcement that millions of app developers will gain access to its AI models could transform educational technology, though benefits may be limited to users with recent devices, potentially creating new digital divides in classrooms.

The Workplace Transformation Paradox

The employment picture reveals a complex paradox. Whilst corporate leaders across America openly prepare for AI-driven job cuts—with CEOs from Amazon, Bank of America, IBM, and others announcing workforce reductions—academic research tells a different story. The European Central Bank found that employment actually rose in AI-affected occupations, whilst PricewaterhouseCoopers reported 38% growth in AI-exposed roles.

The reality likely lies somewhere between these perspectives. Walmart exemplifies this complexity: the retail giant has reduced its workforce by nearly 70,000 people over five years whilst boosting revenues by more than £120 billion, achieving growth through automation rather than expansion. Their new robotic warehouses employ one person per 1,200 square feet and reduce order fulfilment time from 3-4 hours to under 30 minutes.

The UK graduate job market reflects these tensions, with openings at their lowest level since 2018—down 33% year-on-year. Employers are accelerating AI adoption to perform tasks previously given to entry-level workers, with professional occupations hit hardest despite policy changes initially targeting low-wage sectors.

The Human Skills Premium

Despite widespread automation, there's growing recognition of distinctly human capabilities that complement rather than compete with AI. Research across 20 European countries reveals that workers in highly robotised industries report significantly less job purpose, with decreased feelings of autonomy, meaningfulness, and competence. An Amazon warehouse worker exemplified this, describing robot-assisted work as "frankly, damn boring" and eventually leaving for a less automated facility with more social interaction.

This research challenges assumptions that technological advancement automatically improves working conditions, highlighting the importance of preserving meaningful human engagement in work design.

Communication and Authenticity Concerns

Perhaps most concerning for educators is the rise of "intimate deskilling"—where widespread use of generative AI may cause people to lose the ability to articulate desires and communicate feelings without

technological assistance. With ChatGPT's user base doubling from 400 million to 800 million weekly users between February and May 2025, people increasingly rely on AI for deeply personal communications including birthday cards, resignation letters, and dating profiles.

This trend raises fundamental questions about authentic human expression and emotional development, particularly relevant for educators who must balance AI assistance with fostering genuine communication skills.

Legal and Regulatory Developments

The legal landscape remains unsettled, with significant implications for educational institutions. A recent judge's ruling that AI training on copyrighted works constitutes fair use provides some clarity, though major entertainment companies Disney and Universal are challenging image-generation startup Midjourney over alleged copyright infringement.

The Big Four accountancy firms are racing to create AI assurance services that verify the effectiveness and safety of AI tools, particularly important for educational institutions seeking confirmation they comply with emerging AI regulations.

Looking Forward

The evidence suggests we're witnessing a fundamental reshaping of both education and work rather than simple automation. Educational institutions that successfully integrate AI whilst maintaining focus on distinctly human capabilities—critical thinking, emotional intelligence, authentic communication—will likely provide the most valuable preparation for future careers.

For workplace preparation, the emphasis should be on developing skills that complement AI: strategic thinking, ethical reasoning, cross-cultural collaboration, and the ability to work effectively in human-AI teams. The future belongs not to those who can compete with machines, but to those who can work creatively alongside them whilst maintaining essential human qualities that technology cannot replicate.

The transformation is accelerating, but the outcome is not predetermined. How we choose to integrate AI into education and work will determine whether it enhances human potential or diminishes it.

AI News Summary

AI in Education

Business Schools Race to Keep Abreast of Developments in AI

2 June 2025 | [FT Business Education Report](#)

Nearly 80% of business schools have now incorporated AI into their teaching programmes, focusing on creating "strategically fluent leaders" rather than training coders. Leading institutions are taking diverse approaches: Essec Business School in Paris aims to demystify AI's "black box" for executive decision-making, whilst Trinity Business School in Dublin combines digital fluency with human-centric perspectives. Insead has woven AI throughout its education design, teaching leaders to "think with many minds" rather than simply offloading thinking to algorithms.

The transformation extends beyond curriculum to delivery methods. ESCP Business School has partnered with OpenAI to embed generative tools, whilst Imperial Business School requires all executive education staff to complete AI training programmes. Some institutions are developing profession-specific courses, such as Audencia France's "AI + Sales" programme where executives build their own AI agents.

This shift reflects a fundamental change in professional education, moving from teaching about AI to teaching with and through AI, whilst maintaining critical human capabilities that complement rather than compete with artificial intelligence.

Duolingo's AI-Powered Course Expansion
(The Batch Issue 304)

Language learning platform Duolingo has dramatically transformed its course development using generative AI, producing 148 new courses in under a year—more than doubling its catalogue. This achievement represents a stark contrast to their previous pace of creating 100 courses over 12 years. The technology enables them to offer popular languages like Spanish, French, German, Italian, Japanese, Korean, and Mandarin in 28 different base languages.

The process begins with an existing base course, which AI then translates and adapts for numerous target languages. This breakthrough comes at a crucial time for Duolingo, as they face increasing competition from voice-based learning platforms like Speak and Google's Little Language Lessons. The development demonstrates how AI can help educational companies scale their offerings whilst potentially improving quality, addressing the growing pressure to produce more with fewer resources in an increasingly competitive market.

Apple Opens AI Models to App Developers in Cautious Expansion
10 June 2025 | [Financial Times](#)

Apple announced it will allow millions of app developers to access its artificial intelligence models for the first time, potentially transforming educational technology. The move, announced at Apple's Worldwide Developers Conference, emphasises privacy through local device processing rather than cloud-based systems.

Educational apps will gain access to sophisticated AI capabilities whilst maintaining data protection - crucial for educational institutions. Features include real-time translation of calls and messages, which could break down language barriers in multilingual classrooms, and enhanced "Visual Search" functionality for educational content analysis.

However, the benefits may create new digital divides, as AI features are restricted to recent iPhone models. The autumn 2025 rollout gives schools a natural transition point for integration planning, though institutions may need to consider device upgrade cycles and budget implications.

AI Ethics and Societal Impact
Can AI Find Words for Our Feelings?
31 May 2025 | [Financial Times Life & Arts](#)

ChatGPT's user base has doubled from 400 million to 800 million weekly active users between February and May 2025, with people increasingly using AI for deeply personal communications - birthday cards, resignation letters, dating profiles, and family expressions of love. Dating companies like Match Group are integrating AI features for conversation starters, whilst premium Duolingo subscribers can converse with AI chatbots with "sassy personalities."

Luke Brunning, a lecturer in applied ethics at the University of Leeds, warns of "intimate deskilling" - where widespread use of generative AI may cause people to lose the ability to articulate desires and communicate feelings without technological assistance. The article draws parallels to GPS navigation, citing Mountain Rescue data showing a 24% increase in wilderness rescues between 2019-2024, partly attributed to over-reliance on navigation apps.

This development raises fundamental questions about authentic human expression and emotional development, particularly relevant for educators who must balance AI assistance with fostering genuine communication skills.

Sam Altman's Eyeball-Scanning Digital ID Project Launches in UK
9 June 2025 | [Financial Times](#)

Sam Altman's World project has launched its iris-scanning verification service in the UK, using orb devices to "confirm humanness" by scanning people's eyes. The project, which generates World ID digital credentials, anticipates that 90% of online content will be machine-generated within two years, claiming it's already "impossible" to distinguish computers from people in many areas.

The UK was chosen as an "influential market" where 75% of citizens are reportedly affected by AI daily. The service promises to distinguish between humans and AI-generated content for fraud prevention, dating app verification, and preventing ticket touting.

Privacy concerns remain significant despite claims that no data is retained from orb users. The project has faced regulatory pushback in Spain, Portugal, and Germany over security and privacy issues. This development raises fundamental questions about whether private companies should control critical digital identity infrastructure and highlights tensions between technological solutions to AI-created problems and basic privacy rights.

Uncovering Historical Injustices with AI (The Batch Issue 306)

Researchers at Stanford and Princeton have demonstrated AI's potential to illuminate historical wrongs by using a fine-tuned language model to identify racial clauses in property deeds. Their system processed 5.2 million pages of Santa Clara County property records in just six days for £200, a task that would have required nearly 10 years and £1.1 million if done manually.

The results were striking: the AI discovered 24,500 lots covered by racial clauses—roughly one in four homes in 1950—and revealed that just 10 developers were responsible for one-third of these discriminatory restrictions. This work exemplifies how AI can interpret historical documents to reveal actions that might otherwise remain obscured, enabling historians to identify affected areas and trace broader social and economic effects. The researchers have open-sourced their model, streamlining the process for other US counties to conduct similar investigations.

AI's Energy Paradox (The Batch Issue 304)

The International Energy Agency has released a comprehensive analysis revealing AI's complex relationship with energy consumption. The report projects that data centre electricity demand will more than double by 2030, rising from 415 TWh to 945 TWh, with AI accelerator chips consuming four times their current energy levels. The US and China are expected to account for 80% of this growth.

However, the analysis also suggests a silver lining: widespread AI adoption could reduce CO2 emissions by 1.4 gigatons by 2035—five times the projected data centre emissions. AI optimisation of HVAC systems alone could save 300 TWh. This represents the first comprehensive analysis confirming that whilst AI will consume enormous amounts of energy, today's costs may become tomorrow's savings as AI makes energy generation, distribution, and use more efficient across industries.

Working with Robots Often Carries Mental Strain, Studies Find 27 May 2025 | [Financial Times Robotics Guide](#)

Research across 20 European countries reveals that workers in highly robotised industries report significantly less job purpose, with particular impacts on perceptions of meaningfulness and autonomy. A University of Groningen study found that only 3% of European workers directly interact with robots, yet robotisation increased 389% across all industries between 2005-2020.

Amazon warehouse worker "Jessica" exemplified the human cost: whilst robot-assisted work was less physically demanding, she found it "frankly, damn boring," working 12-hour shifts with 20 robots lined up during busy periods. She eventually left for a less automated facility with more social interaction.

Studies show workers in robotised environments experience decreased autonomy (4%), meaningfulness (3.4%), and competence (2.7%). Even highly skilled roles are affected - jet pilots operating surveillance drones and surgeons using robotic arms find work becomes "safer but also more boring and alienating." The research challenges assumptions that technological advancement automatically improves working conditions.

AI and Cybersecurity

Cyber Crime is Surging: Will AI Make it Worse?

7 June 2025 | [Financial Times Weekend Essay](#)

Global cyber crime will cost \$10.5 trillion in 2025 - if it were a legitimate industry, it would be the world's third-largest economy, twice the size of Germany's. AI is dramatically accelerating criminal capabilities: voice phishing attacks increased 442% in the second half of 2024, with AI creating convincing fake emails and voices with perfect local accents.

Criminal AI tools like WormGPT, explicitly designed for illegal activities, launched in June 2023. Researchers have found universal "jailbreak" attacks on all major AI programmes, whilst China's DeepSeek AI has been used to find "zero-day" vulnerabilities in leading browsers.

The threat is escalating toward "agentic AI" that could autonomously handle entire cyber-crime operations - writing malware, executing phishing campaigns, encrypting data, issuing ransom demands, and laundering Bitcoin payments. Major incidents include the NHS Synnovis attack cancelling 12,000 appointments and the record-breaking \$1.4 billion Lazarus Group cryptocurrency heist.

Three key factors drive this growth: the internet's inherent insecurity (built for interoperability, not security), Bitcoin providing untraceable criminal currency since 2009, and lower barriers to entry through Ransomware-as-a-Service platforms requiring no technical skills.

New Threats to Autonomous Agents (The Batch Issue 304)

Researchers at Columbia University have uncovered a concerning vulnerability in autonomous AI agents, developing a method to exploit them by poisoning popular websites with malicious links. The attack works by exploiting agents' tendency to trust popular sites like Reddit, allowing attackers to redirect them to malicious websites.

Once redirected, the agents proved remarkably compliant with attacker instructions—divulging credit card information in 10 out of 10 trials and sending phishing emails with the same success rate. This research highlights a critical security concern as AI agents gain real-world action capabilities, demonstrating that manipulating agents through malicious web content represents an effective attack vector. The findings suggest that future agents must be designed with built-in resistance to such manipulation tactics.

UK Must Toughen Regulation of Facial Recognition, Say AI Experts 29 May 2025 | [Financial Times Biometrics](#)

Almost 5 million faces were scanned by UK police forces in 2024, resulting in over 600 arrests, as facial recognition technology expands from policing into retail stores, sports stadiums, and shopping centres. Major retailers including Southern Co-op, Budgens, Sports Direct, and Asda are now using the technology.

The Ada Lovelace Institute warns of "significant gaps and fragmentation across biometrics governance," with privacy campaigners describing the UK as a "wild west" due to absent rules around "live" facial recognition systems. The legality of current deployments is in "serious question" following a 2020 UK appeals court ruling that South Wales Police's use broke privacy and data protection laws.

Policing minister Dame Diana Johnson acknowledged "very legitimate concerns" whilst calling the technology "transformational for policing." The UK lags significantly behind the EU's AI Act and several US states that have banned many applications of live facial recognition systems.

Emerging AI developments threaten to make the technology more powerful and affordable, with new systems claiming to infer emotions, intentions, and truthfulness - capabilities that critics argue interfere with protest rights and frequently misidentify innocent people as criminals.

AI Employment and the Workforce

UK Graduate Job Openings at Lowest Level Since 2018 25 June 2025 | [Financial Times](#)

The UK graduate job market has collapsed to its lowest level since 2018, with openings down 33% year-on-year according to Indeed data. Graduate roles now represent 12% fewer of all job postings, whilst

overall postings have fallen 5% since April policy changes. The UK stands as the only developed economy with fewer job openings than before the pandemic, now 21% below pre-pandemic baseline.

Employers are accelerating AI adoption to perform tasks previously given to entry-level workers, with professional occupations hit hardest despite policy changes targeting low-wage sectors. Human resources and project management roles are 27% below pre-pandemic levels, marketing down 37%, and media and communications down 48%.

The Bank of England's Deputy Governor Dave Ramsden has voted for interest rate cuts, citing "material loosening in labour market," with rising online searches for "redundancy" and dropping rates of unemployed people finding work. This represents a fundamental labour market transformation driven by the combination of AI automation and policy-driven employment cost increases.

Corporate Leaders Prepare for AI-Driven Job Cuts
(The Batch Issue 307)

Major corporate leaders across America are openly preparing for AI to eliminate numerous jobs from their workforces. Amazon CEO Andy Jassy has written that AI will reduce the corporate workforce, whilst CEOs at Bank of America, IBM, Shopify, and Williams-Sonoma have made similar pronouncements. Globally, around 40% of employers expect to downsize due to AI implementation.

However, this narrative faces significant counterevidence from academic research. The European Central Bank found that employment actually rose in AI-affected occupations, whilst PricewaterhouseCoopers reported 38% growth in AI-exposed roles. Historical precedent supports this more optimistic view—technological advances typically create more jobs than they destroy, with an estimated 60% of US jobs in 2018 not existing in 1940.

The reality likely lies somewhere between these perspectives. Whilst AI will indeed impact knowledge workers and creative professionals previously considered immune to automation, executives bear responsibility for thoughtful preparation and transition planning rather than wholesale workforce reduction.

Walmart is Supercharging Revenue — But with Fewer Workers
3 June 2025 | [Financial Times](#)

Walmart, America's largest private sector employer, counted 2,165,465 staff worldwide at the end of 2024 - almost 70,000 fewer than five years ago, despite boosting revenues by more than \$150 billion over the same period. The company aims to grow sales by 4% annually without significantly increasing headcount.

New robotic warehouses outside Dallas showcase the transformation: a cold-storage hub employs 600 associates in 730,000 square feet (one employee per 1,200 square feet), whilst order fulfilment time has been reduced from 3-4 hours to under 30 minutes. Work processes have been condensed from 12 steps to 5, cutting costs by an expected 30%.

This contrasts sharply with competitors: Costco, Target, and Home Depot each added tens of thousands of employees over five years, whilst Amazon nearly doubled its workforce to 1.6 million. Net US sales rose 36% in the past five years whilst average hourly wages increased only 28% to \$18.25, raising questions about whether workers are sharing proportionally in the company's success.

The transformation demonstrates how companies can achieve growth through automation rather than expansion, with significant implications for retail employment and economic inequality.

Duolingo CEO on Going AI-First: 'I Did Not Expect the Blowback'
8 June 2025 | [Financial Times](#)

Duolingo CEO Luis von Ahn faced public backlash after announcing the language-learning company would go "AI-first," with users cancelling accounts and ending learning streaks over fears of mass layoffs. Von Ahn

attributed the anger to general "anxiety" about technology replacing jobs, clarifying that changes focus on overhauling work processes rather than eliminating employees.

The company's AI-first strategy involves staff performance reviews evaluating AI use, with 50% of new hires being graduates who "come with different mindset" using AI at university. Engineering teams spend less time writing code through AI-mediated development, whilst designers take supervisory roles with AI creating artwork in Duolingo's "very specific style."

With 10.3 million paying subscribers (40% increase year-over-year) and revenue of \$230 million (38% increase), Duolingo is expanding beyond 40 languages using AI to multiply offerings. Premium subscribers can converse with "Lily," an AI chatbot with a "sassy personality" that personalises based on user preferences.

Von Ahn, who wakes at 2-3am thinking about business and maintains both free and paid Duolingo accounts, represents a generation of CEOs navigating public anxiety about AI displacement whilst pursuing technological transformation.

AI Development and Industry

Could AI Make a Scorsese Movie? Demis Hassabis and Darren Aronofsky Discuss

30 May 2025 | [Financial Times](#)

Demis Hassabis, CEO of Google's DeepMind and Nobel laureate in chemistry, has partnered with Oscar-winning filmmaker Darren Aronofsky to explore AI's creative potential. Their collaboration, beginning with a 1999 meeting about storytelling's future, has led to Aronofsky's new film company Primordial Soup partnering with DeepMind to produce AI-assisted short narrative films.

Their core philosophy centres on Aronofsky's motto: "Make soup, not slop" - emphasising that AI-generated content needs storytelling and emotion to be memorable. Both see AI as a tool to enable creativity rather than replace human artists, pushing AI models to imagine previously impossible visuals.

"Ancestra," directed by Eliza McNitt and premiering at Tribeca Film Festival on 13 June 2025, combines live-action performance with AI-generated imagery featuring "unphotographable" elements. Despite being a short film, it required a crew size comparable to a feature film. Screen Actors Guild members provided input, with unions reportedly "leaning into" AI technology rather than opposing it.

However, Hassabis notes current limitations: AI systems excel at extrapolation but cannot generate truly original concepts, lacking "true invention," analogical reasoning, and the intuitive judgement that separates good from great artists. The technology helps AI systems understand physics, lighting, and material behaviour, with DeepMind planning integration into platforms like YouTube.

Biomni: AI Agent for Biological Research
(The Batch Issue 307)

Stanford, Princeton, and other leading institutions have introduced Biomni, a groundbreaking research agent based on Claude 4 Sonnet that performs tasks across genomics, immunology, microbiology, neuroscience, and pathology. The system utilises 150 specialised tools and nearly 60 databases, following the CodeAct framework with built-in quality checking mechanisms.

The results demonstrate significant advancement in AI-assisted research capabilities. On Humanity's Last Exam biomedical subset, Biomni achieved 17.3% accuracy compared to Claude 4 Sonnet alone (6%) and with research access (12.2%). For genome-based patient diagnosis, Biomni reached approximately 85% accuracy versus Claude 4 Sonnet's 5%. Unlike typical biology-focused agents that concentrate on narrow specialities, Biomni's knowledge spans the entire biological domain, offering expert assistance across multiple specialities with reasoning capabilities that improve alongside more capable language models.

Meta's Major Investment in Scale AI
(The Batch Issue 307)

Meta has made a substantial strategic move by hiring Scale AI's entire leadership team, including founder and CEO Alexandr Wang, whilst investing £11.4 billion in the data-labelling startup and acquiring 49% in

non-voting shares. This arrangement echoes similar deals by Microsoft with Inflection AI, Amazon with Adept AI, and Google with Character.AI. Wang will oversee a new Meta research laboratory focused on developing superintelligence.

This deal represents Meta's aggressive strategy to establish a decisive lead in AI through substantial financial commitments. The acquisition provides Meta with access to a star AI entrepreneur and closer integration with Scale AI's high-quality training data pipeline, whilst Scale AI gains Meta's enormous resources. For the broader AI community, Meta's willingness to spend such sums could drive up engineers' salaries whilst potentially blocking less well-funded competitors from accessing key talent and resources.

DeepSeek's Cost-Effective AI Training Breakthrough (The Batch Issue 303)

Chinese AI company DeepSeek has revealed the technical innovations that enabled them to build DeepSeek-R1 and DeepSeek-V3 at the remarkably low cost of £4.5 million using 2,048 Nvidia H800 GPUs. Their approach included several key innovations: mixed-precision training with FP8 computations (making it the first open language model to use FP8), limited expert routing to four nodes for faster communication, and multi-head latent attention that saves memory compared to alternatives.

This breakthrough challenges the conventional wisdom that only companies with massive budgets can afford state-of-the-art model training. By sharing their technical details, DeepSeek has empowered a wider range of teams to potentially improve the state of the art, democratising access to advanced AI development techniques that were previously the exclusive domain of well-funded technology giants.

DeepSeek-R1 Performance Updates (The Batch Issue 304)

DeepSeek has updated its reasoning model to DeepSeek-R1-0528, achieving performance that approaches OpenAI's o3 and Google's Gemini-2.5 Pro whilst surpassing its predecessors. The company also released a smaller DeepSeek-R1-0528-Qwen3-8B version capable of running on a single GPU with 40GB VRAM.

The performance improvements are substantial: on AIME 2025 mathematics problems, the 8B model achieved 76.3% accuracy compared to the much larger Qwen3-32B's 72.9%. The company claims 50% fewer hallucinations and dramatic improvements across reasoning benchmarks. DeepSeek's models continue to narrow the gap between open-weights and closed models, potentially leading to wider adoption of less expensive, more efficient approaches whilst passing cost savings to developers through high-performance inference at fractions of closed model costs.

Microsoft's BitNet Breakthrough in Low-Precision AI (The Batch Issue 307)

Microsoft and Chinese researchers have achieved a significant breakthrough with BitNet b1.58, a model that limits most weight values to just -1, 0, or +1, yet competes effectively with full-precision models up to 2 billion parameters. Their 2-billion parameter model achieved 54.19% average accuracy across 16 benchmarks whilst operating significantly faster and using less memory than competitors.

The performance advantages are striking: BitNet generated 34.5 tokens per second compared to Qwen2.5-1.5B's 15.4 tokens per second, used just 0.4GB memory versus 2.6GB, though with slightly lower accuracy (54.19% versus 55.23%). This work demonstrates that quantising a language model to few bits requires rethinking training approaches, including hyperparameter details like learning rate and weight decay, where seemingly small changes can significantly impact performance.

Apple's Foundation Models Development (The Batch Issue 306)

Apple has updated its Foundation Models (AFM) family with smaller on-device versions (3B parameters) and larger server-hosted versions, plus released the Foundation Models framework API. Both models incorporate vision encoders with language models and employ aggressive quantisation techniques.

The on-device AFM performed better than competitors at non-US English tasks and image understanding, whilst the server version's performance wasn't decisively better than competitors and underperformed GPT-4o across all reported tests. This development appears to be Apple's bid to reset its position in the AI race, addressing perceptions that the company has fallen behind with delayed Siri upgrades and lack of advanced features that have led to class-action lawsuits.

OpenAI's o3-pro Reasoning Model
(The Batch Issue 306)

OpenAI has launched o3-pro, a more capable version of its reasoning vision-language model designed specifically for difficult science, mathematics, and coding challenges, though it operates dramatically slower with 129.2 seconds to first token. On AIME 2024 mathematics problems, o3-pro achieved 93% accuracy compared to o3's 90% and o1-pro's 86%. For GPQA Diamond graduate-level science questions, o3-pro reached 85% versus o3's 81% and o1-pro's 79%.

This development represents OpenAI pushing the limits of current reasoning approaches with promising but incremental results. The extensive reasoning capabilities may appeal to developers working on multi-step scientific problems, though the high price and slow speed may prove dealbreakers for many applications.

Black Forest Labs' FLUX.1 Kontext for Creative Consistency
(The Batch Issue 305)

Germany's Black Forest Labs has released the FLUX.1 Kontext family (available in max, pro, and dev versions) trained specifically for controlled image alterations, character consistency, and style maintenance. The system uses adversarial diffusion distillation to reduce inference steps whilst maintaining quality.

In proprietary benchmarks using approximately 1,000 crowd-sourced prompt-image pairs, FLUX.1 Kontext max and pro outperformed all competing models, including OpenAI's GPT Image 1 and Google's Gemini 2.0 Flash. This advancement enables artists to craft stories around specific characters with greater consistency, addressing a long-standing challenge in AI-generated imagery where maintaining consistent details across images has remained finicky.

Nvidia's STORM: Efficient Video Processing
(The Batch Issue 305)

Nvidia and collaborators have developed STORM, a system that reduces video tokens processed by language models whilst maintaining performance. The system uses mamba layers to enrich token embeddings across frames, enabling an eightfold reduction in processed tokens.

On MVBench 16-second videos, STORM achieved 70.6% accuracy versus GPT-4o's 64.6%. For MLVU (3 minutes to 2+ hours), STORM reached 72.9% compared to GPT-4o's 66.2%, whilst running three times faster than baseline systems. This demonstrates how combining mamba and transformer architectures can yield benefits of both approaches, processing one-eighth as many tokens and using one-eighth the compute whilst delivering superior performance.

Claude 4's Coding Advancement
(The Batch Issue 303)

Anthropic has launched Claude 4 Sonnet and Opus with optional reasoning mode, parallel tool use, and made Claude Code generally available. Both models achieved top marks in LMSys WebDev Arena and coding benchmarks. On SWE-bench Verified, Claude Opus 4 succeeded 72.5% of the time compared to OpenAI o3's 70.3%. On Terminal-bench, Claude Opus 4 achieved 39.2% versus GPT 4.1's 30.3%.

The new models extend language model technology with parallel tool use, external file memory, and the ability to stay on-task over unusually long periods. Early users report impressive projects including single-shot game development and seven-hour code refactoring sessions, demonstrating the models' capacity for sustained, complex technical work.

**Google's I/O Developments
(The Batch Issue 303)**

Google has announced significant improvements to Gemini 2.5 Pro/Flash with audio output capabilities, the Veo 3 video generator, Gemma 3n mobile-optimised models, and increasingly AI-powered search features. Key updates include Gemini-2.5-Pro-Preview's ability to set reasoning budgets up to 128K tokens, Veo 3's production of 4K video with audio and creative controls, and Gemma 3n models optimised for mobile devices with 2-3GB memory requirements.

These developments represent Google's efforts to catch up with Microsoft and OpenAI on several fronts. The audio output addition fuels voice applications development, whilst Veo 3 shows marked improvement over previous versions. Monthly token processing has surged from 9.7 trillion to 480 trillion, indicating rapid API and product adoption.

**OpenAI Risks Being Undercut by Cheaper Rivals, Says Star Investor Mary Meeker
30 May 2025 | [Financial Times](#)**

Mary Meeker, the renowned Silicon Valley analyst known as "the queen of the internet," warns that US AI companies like OpenAI face significant threats from cheaper competitors, particularly China's DeepSeek. Despite predicting AI will create "multiple companies worth \$10tn," she notes they "probably will not all be based in North America."

ChatGPT reached over 500 million monthly users within two years, representing an unprecedented "five billion-user market that one could get to so easily." However, the three leading US AI start-ups (OpenAI, xAI, and Anthropic) have a combined valuation of around \$400 billion after raising \$95 billion to achieve collective annualised revenue of only \$12 billion.

Training costs for cutting-edge models have increased 2,400 times in eight years, yet massive chip and algorithm improvements have slashed operational costs, enabling cheaper competitors. General-purpose large language models increasingly resemble "commodity businesses with venture-scale burn," whilst smaller, cheaper models trained for specific use cases emerge as alternatives.

Meeker compares the situation to early Uber, Amazon, and Tesla - companies requiring massive cash burn before profitability - recommending a portfolio approach and warning that current euphoria could end suddenly: "everything is up and to the right — until it isn't."

**Chinese Tech Groups Prepare for AI Future Without Nvidia
30 May 2025 | [Financial Times](#)**

China's biggest technology companies - Alibaba, Tencent, and Baidu - are beginning the "long and difficult process" of switching AI development to homegrown chips as their Nvidia stockpiles dwindle. Trump's administration has clamped down on sales of Nvidia's H20 chip, and Chinese companies' existing stockpile will only sustain AI development until around early next year.

Baidu's AI cloud head Shen Dou said the company has "a range of chip options" to replace Nvidia's, particularly for inference processing. Alibaba's Chief Eddie Wu stated they're "actively exploring diversified solutions," whilst Tencent's President Martin Lau said they have "enough high-end chips to continue training models for a few more generations" while exploring alternatives.

Huawei's Ascend chips are the primary domestic alternative, with state-owned companies like China Mobile making "large-scale procurement." However, switching from Nvidia's CUDA software framework to Huawei's CANN is "extremely time-consuming," with one executive estimating "about three months of disruption in AI-related development."

Washington has warned that using Huawei chips "anywhere in the world" could result in criminal penalties, keeping many tech companies quiet about testing Ascend chips. Most companies are considering hybrid strategies: using existing Nvidia chips for AI training whilst employing local processors for inference.

The Emerging Reality of the OpenAI-SoftBank Grand Plan for Data Centres
29 May 2025 | [Financial Times](#)

The \$500 billion Stargate project, announced in January 2025 as "the largest AI infrastructure project by far in history," reveals a significant gap between headline figures and actual commitments. The \$500 billion came from OpenAI's calculations of computing power costs, not actual invested capital - only about \$50 billion in real investment commitments exist.

After four months, only \$7.5 billion has actually changed hands, arriving in OpenAI's account in May 2025. SoftBank and OpenAI each committed about \$18 billion, with Oracle and MGX agreeing to \$7 billion each. OpenAI's contribution will largely come from capital raised from SoftBank through a \$40 billion funding round.

The only money directly deployed is in Abu Dhabi for "Stargate UAE" (up to \$20 billion from G42). In the US, only one OpenAI-related project exists: a 1.2GW data centre in Abilene, Texas, which predates Stargate and is owned by Crusoe and Blue Owl Capital, not Stargate partners.

Baker Botts law firm estimates the global data centre industry will triple by 2034, from \$256 billion to \$776 billion in asset value. However, Silicon Valley investors question whether this mirrors dotcom-era fibre optic cable overinvestment, which led to overcapacity and failures when the bubble burst.

Nvidia's Suppliers Resolve AI Rack Issues in Boost to Sales
27 May 2025 | [Financial Times](#)

Nvidia's suppliers - Foxconn, Inventec, Dell, and Wistron - have resolved technical issues that delayed shipments of flagship "Blackwell" AI servers. Partners announced at Computex conference that GB200 rack shipments began at the end of the first quarter, with production capacity now being rapidly scaled up.

Each AI rack includes 36 "Grace" central processing units and 72 Blackwell graphics processing units connected through Nvidia's NVLink system. Technical challenges included overheating from 72 high-performance GPUs, liquid cooling system leaks, software bugs, and inter-chip connectivity problems. An engineer noted: "No company has tried to make this many AI processors work simultaneously in a server before."

Issues were resolved 2-3 months ago according to supply chain partners. Nvidia is targeting approximately \$43 billion in sales for the quarter ending April 2025 (65% year-over-year increase). The company is looking beyond Big Tech "hyperscaler" companies to nation states for customer diversification, with Saudi Arabia and UAE announcing plans to acquire thousands of Blackwell chips during Trump's Gulf tour.

However, US government bans on H2O chip exports to China are expected to cost Nvidia \$5.5 billion in charges from inventory write-offs and purchase commitments, with Bank of America estimating gross margins will drop from 71% to approximately 58%.

AI Boom Adds €150bn to Value of Four of Europe's Oldest Industrial Groups
27 May 2025 | [Financial Times](#)

Four European industrial groups - Schneider Electric, Siemens AG, ABB, and Legrand - have added more than €150 billion to their market caps since ChatGPT's launch in November 2022. These companies are pivoting from traditional electrical equipment to data centre infrastructure, with Dell'Oro expecting total data centre capital expenditure to increase from almost \$600 billion in 2025 to more than \$1 trillion by 2028.

Schneider Electric, valued at €127.9 billion, leads through its 2006 acquisition of American Power Conversion for \$6.1 billion. Data centres now account for 24% of orders (up from 19% in 2022), recently acquiring 75% of liquid cooling specialist Motivair for \$850 million. The company's market cap now exceeds French oil major TotalEnergies.

Legrand projects 20-25% of business from data centres by 2030, making 10 small acquisitions in the past year, six in the data centre sector. However, high exposure to US clients creates tariff vulnerability, with CEO Benoît Coquart warning that 50-60% US tariffs on Chinese goods could cost up to \$200 million annually.

ABB's data centre orders represent 15% of its \$16.4 billion electrification unit orders in 2024 (up from 9% two years earlier), growing at 24% average annual rate between 2019-2023. Siemens AG is investing to compete, with data centre business rising more than 45% to approximately €1.3 billion in the first half of fiscal year.

Musk's xAI Agrees Telegram Tie-up as Billionaire 'Bromance' Blooms
28 May 2025 | [Financial Times](#)

Elon Musk's xAI has agreed a \$300 million deal with Pavel Durov's Telegram to distribute the Grok chatbot to 1 billion users. The one-year agreement includes cash and equity payment to Telegram, plus half of revenue from xAI subscriptions sold via the messaging app, giving Telegram a stake in xAI.

This represents xAI's first major expansion beyond Musk's X platform, following a recent alliance with Microsoft making xAI available through Azure cloud computing. The partnership exposes xAI to new audiences and valuable data as tech companies compete in the AI race.

Both leaders share passion for free expression and opposition to government censorship, described as a tech "bromance." Durov is under formal investigation in France over Telegram's alleged failure to address criminality including child abuse and terrorism, but remains ordered to stay in France while allowed business trips to Dubai.

Musk, described as a "self-declared free speech absolutist," has challenged takedown requests in Brazil, India, and Australia, whilst backing Germany's far-right AfD party. Critics argue both leaders are selective in their free speech advocacy.

The deal status remains unclear, with Musk stating it wasn't "signed" but Durov responding they had "agreed in principal."

Telegram is currently launching a \$1.5 billion bond offering to buy back existing debt, reporting 2024 revenue of \$1.4 billion (up from \$343 million in 2023) and first annual profit of \$540 million.

Saudi Arabia Seeks to Use Financial Might to Muscle into Global AI Industry
28 May 2025 | [Financial Times](#)

Saudi Arabia's new state-owned AI company Humain seeks investment from top US tech companies including OpenAI, Elon Musk's xAI, and Andreessen Horowitz. The company plans to launch a \$10 billion venture capital fund this summer targeting start-ups in the US, Europe, and parts of Asia, aiming for a central role across AI industry investment, infrastructure, and chip design.

Humain targets 1.9 gigawatts of data centre capacity by 2030, rising to 6.6GW by 2034, at an estimated cost of \$77 billion. The goal is processing 7% of global AI "training" and "inferencing" by 2030. The first phase involves a 50MW plant with 18,000 Nvidia chips, expanding to 500MW requiring 180,000 chips.

Already secured \$23 billion in deals with US tech groups since launch: a \$10 billion joint venture with AMD for 500MW capacity over five years, and \$2 billion investment with Qualcomm for data centres and chip design capabilities. Qualcomm will establish a chipset design centre in Riyadh employing 500 engineers.

Humain is owned and funded by the \$940 billion Public Investment Fund, chaired by Crown Prince Mohammed bin Salman. The unveiling came the day before Trump's Riyadh visit with tech executives including Musk, Sam Altman, and Jensen Huang, as part of the kingdom's effort to diversify its oil-dependent economy and become a "data exporter."

CEO Tareq Amin emphasises speed: "Whoever reaches the end line first is going to secure a good chunk of the market share," positioning Saudi Arabia as a solution to global capacity hunger whilst focusing on American tech partnerships to reassure US policymakers concerned about technology transfer to China.

Meta Invests \$15bn in Scale AI, Doubling Start-up's Valuation
13 June 2025 | [Financial Times](#)

Meta has invested \$15 billion in data-labelling start-up Scale AI, valuing the company at \$29 billion and giving Meta 49% equity. Scale AI specialises in manually ensuring images and text are accurately labelled before training AI models, expecting \$2 billion in revenues this year primarily from major AI model builders like OpenAI.

Alexandr Wang, Scale's 28-year-old co-founder and CEO, has been hired by Meta to work on their AI efforts whilst remaining on Scale's board. Jason Droege, former Uber Eats launcher and venture capitalist, becomes interim CEO. The deal reflects growing patterns of tech giants investing in start-ups whilst acquiring key talent, similar to Microsoft's \$650 million Inflection deal and Google's \$2.7 billion Character.AI arrangement.

For Meta, this supports CEO Mark Zuckerberg's strategy to compete in the AI race and build a "superintelligence" team, addressing underperformance of Meta's Llama 4 model on reasoning and coding benchmarks. The investment supports Meta's \$72 billion capital expenditure plan for AI infrastructure this year.

Early customers were autonomous vehicle companies, but Scale's focus has shifted to AI model training data. The massive investment demonstrates the premium companies will pay for high-quality AI training data and highlights severe skills shortages in AI research and development.

OpenAI Revenue Doubles to \$10bn Despite Remaining Unprofitable
10 June 2025 | [Financial Times](#)

OpenAI's annual recurring revenue has nearly doubled to \$10 billion, up from \$5.5 billion at the end of 2024, despite remaining loss-making with profitability not expected until 2029. The company projects \$125 billion revenue when profitability is achieved. ChatGPT now serves over 500 million weekly active users with approximately 3 million business and education subscriptions.

This performance outpaces competitors: Anthropic tripled its ARR to \$3 billion between January and May 2025, whilst coding tool Cursor grew from under \$100 million to \$500 million ARR in 2024-2025. All major AI companies remain unprofitable despite massive investment.

Market adoption shows 40% of US businesses now pay for AI models (quadrupled in two years), though growth stalled in May 2025 for the first time in 10 months. Fintech company Ramp suggests most willing-to-pay businesses are already customers.

OpenAI is raising \$40 billion from SoftBank and other investors, recently acquiring hardware start-up io for \$6.4 billion and in process of acquiring code editing company Windsurf. The company collaborates with the Trump administration on the Stargate data centre project, whilst the focus on business and education subscriptions indicates AI companies see education as a key revenue stream.

Apple's Struggles to Update Siri Lead to Investor Concerns Over AI Strategy
8 June 2025 | [Financial Times](#)

Apple is struggling to integrate cutting-edge large language models with existing Siri infrastructure, with former executives reporting "bugs" from attempting to build LLMs over current machine learning

technology. Competitors like OpenAI built generative AI voice assistants "from scratch" without legacy system constraints.

Apple is the worst-performing stock in "Magnificent 7" tech stocks in 2025, down approximately 18% since the start of year versus Nasdaq largely flat. Investor expectations are low for the upcoming Worldwide Developer Conference (WWDC), with JPMorgan analyst Samik Chatterjee saying "let's first have you deliver what you promised last year."

Apple Intelligence, a suite of AI features announced at WWDC 2024 intended to boost hardware sales, has faced significant delays. While some features like writing aids and image generation have been released, key Siri upgrades remain unreleased. CEO Tim Cook admitted the technology didn't meet their "high quality bar" and is "taking a bit longer than we thought."

Organisational changes include John Giannandrea (AI guru from Google) losing control of Siri division, which was transferred to Mike Rockwell (Vision Pro executive). Former executives cite "fragmented leadership teams" and initial insufficient budget allocation.

Bank of America analysts suggest Apple is "essentially three years or more away from delivering a truly modern AI assistant," whilst competitors operate "much larger and more powerful models." Apple's privacy-first approach, prioritising smaller models and on-device processing, adds complexity compared to cloud-based LLMs running on powerful servers.

Europe's Mistral Benefits from Search for AI Alternatives
8 June 2025 | [Financial Times](#)

European AI company Mistral has secured new contracts worth hundreds of millions of dollars, with revenues increasing "several times over" since its last funding round a year ago. The company is on track to surpass \$100 million annual revenue for the first time, with business tripling in the last 100 days, particularly in Europe and outside the US.

CEO Arthur Mensch attributes growth to "growing demand for more strategic autonomy," as corporate, public-sector, and defence customers seek alternatives to US tech companies. This trend is driven by the Trump administration's return and increasing geopolitical tensions.

Multiple contracts worth at least \$100 million each over 3-5 years are modelled on the €100 million deal with French shipping group CMA CGM announced in April 2025. Mistral focuses on "bespoke" AI systems for enterprise clients, with major customers including BNP Paribas, AXA, Stellantis, and Veolia.

Founded by three former Meta and Google DeepMind researchers two years ago, Mistral employs approximately 250 people and has significantly expanded its commercial team. The company competes with China's DeepSeek and Meta's Llama in "open" AI models that customers can examine and customise.

Mistral is considering raising up to \$1 billion more in potential funding, having already raised over \$1 billion since founding. The company is embarking on ambitious infrastructure expansion, building a large data centre outside Paris and collaborating with Abu Dhabi-based tech groups G42 and MGX.

AI Regulation and Legal Issues

New York Times Agrees First AI Deal with Amazon
29 May 2025 | [Financial Times](#)

The New York Times has agreed to license its editorial content to Amazon, marking the first time the news giant will allow its stories to train a tech company's artificial intelligence models. The deal allows summaries and short excerpts of NYT stories and cooking recipes to be used by Amazon products such as Alexa speakers, whilst content will train Amazon's proprietary AI models.

This agreement comes whilst the NYT continues its legal battle with OpenAI and Microsoft after suing them in 2023 for "massive copyright infringement," seeking billions of dollars in damages for allegedly using

millions of their articles to build ChatGPT's underlying models. OpenAI has said the NYT suit was "without merit."

CEO Meredith Kopit Levien said in an internal memo that the Amazon deal "is consistent with our long-held principle that high-quality journalism is worth paying for" and "aligns with our deliberate approach to ensuring that our work is valued appropriately."

Several media groups have struck licensing deals with OpenAI in recent years, including News Corp, Axel Springer, and the Financial Times, bringing in tens of millions of dollars for publishers. However, some news executives privately view these deals as temporary measures whilst waiting for long-term legal frameworks.

The agreement highlights the complex negotiations required as traditional industries adapt to AI technologies, balancing immediate revenue needs with long-term strategic positioning whilst maintaining editorial independence.

Judge Rules AI Training on Copyrighted Works Is Fair Use (The Batch Issue 307)

A significant legal precedent has emerged with a judge ruling that training AI systems on copyrighted works constitutes fair use, though specific details of the case and its broader implications remain to be fully understood. This ruling comes at a crucial time as the AI industry grapples with questions about data usage and copyright law.

Hollywood Takes on AI Image Generation (The Batch Issue 306)

Disney and Universal have filed suit against Midjourney, accusing the image-generation startup of training on "countless" unauthorised copies of their copyrighted works. The entertainment giants seek cessation orders and damages of up to £120,000 per infringed image.

The lawsuit alleges that Midjourney reproduces characters from major franchises including Star Wars, Toy Story, Cars, Ironman, and The Simpsons, even without explicit prompts. The studios claim Midjourney could implement software to prevent generating copyrighted material. This case, brought by two of the most powerful entertainment companies, could establish precedents that influence future lawsuits, AI company behaviour, and legislation updating copyright law for the AI era, particularly as AI systems show greater potential to erode markets for human-made creative works.

Evidence of Training Data Copyright Infringement (The Batch Issue 303)

A study co-authored by publisher Tim O'Reilly has found evidence that GPT-4o likely trained on paywalled O'Reilly books, as the model could identify verbatim excerpts from these protected materials. Using the DE-COP method, researchers compared recognition of paywalled versus freely available excerpts across 34 books divided into approximately 14,000 paragraphs.

GPT-4o recognised private paragraphs with 82% AUROC compared to 64% for public paragraphs, whilst GPT-3.5 Turbo showed the opposite pattern (64% versus 54%). This research highlights the ongoing tension between the AI industry's need for vast quantities of high-quality data and copyright owners' concerns about earning opportunities. The findings suggest developers must find fair responses to avoid creating an "extractive dead end" that diminishes the supply of high-quality training data.

Big Four Firms Race to Develop Audits for AI Products 3 June 2025 | [Financial Times](#)

The Big Four accountancy firms - Deloitte, EY, PwC, and KPMG - are racing to create AI assurance services that verify the effectiveness and safety of artificial intelligence tools. These services would cover systems such as self-driving cars, cancer-detecting programmes, and customer service chatbots, checking whether chatbots answer questions accurately and identifying issues such as bias.

Growing demand comes from companies seeking confirmation they comply with AI regulations, with some insurers beginning to offer cover for losses caused by malfunctioning AI tools. Demand is particularly high in financial services, life sciences, and pharmaceutical sectors.

PwC UK will launch AI assurance services "soon," whilst Deloitte considers AI assurance "critical" to AI adoption. The Institute of Chartered Accountants in England and Wales held its first conference on the topic last month, as firms attempt to shape the emerging field and avoid losing out to nimble start-ups.

However, significant challenges remain. EY's UK technology risk leader warns development could take significant time due to potential liabilities: "We are still quite a way away from being able to say that we are unequivocally giving assurance of an AI model." AI models continue to ingest data and develop over time, so they won't always react the same way in given scenarios.

Unlike financial audits, there's a lack of standardisation in the AI assurance market. Hundreds of firms in the UK already supply forms of AI assurance, but most current verification is provided by AI developers themselves, raising independence concerns.

Law Firm Clients Seek Clarity on AI's Potential to Cut Costs
29 May 2025 | [Financial Times](#)

Big law firms globally are adopting generative AI extensively, but corporate clients have yet to see cost savings despite legal service prices continuing to rise. Top 100 US law firms increased rates by 10% year-on-year in 2024, with Alex Kelly of Brightflag noting "there has not been a fundamental change in how law firms are delivering the work."

BT's general counsel Jeff Langlands is asking law firms to demonstrate where they use generative AI and how it benefits the company. Law firms are not yet willing or able to show specific cost savings on tasks like due diligence, though clients expect to see "compare and contrast" evidence of AI-driven efficiencies within the next year.

Clients are divided on AI adoption: some companies fully support maximum AI use, whilst others prohibit it entirely, putting law firms in difficult positions balancing conflicting client concerns. Andrew Perlman of Suffolk University Law School emphasises that "client pressure matters a lot" in driving change.

Generative AI could bring the largest material change to the billable hours model than any previous technology. However, law firms must invest in AI services and legal tech, leading some to incorporate these costs into client billing. Kerry Westland of Addleshaw Goddard notes it's challenging to break down AI cost savings as the technology is applied to parts of projects rather than entire matters.

Alternative fee arrangements may become more common as firms build customised AI models, with subscription-based legal services potentially shifting work away from billable hours. Success requires information exchange between law firms and in-house departments about AI benefits, efficiency gains, and ethical considerations.

AI's Seismic Effect Changes Client Expectations of Law Firms
25 June 2025 | [Financial Times](#)

Generative AI is fundamentally transforming legal practice, requiring law firms and in-house legal teams to develop new hybrid approaches combining AI technology with human expertise. Clients now explicitly demand that law firms demonstrate how they combine generative AI with human expertise in proposals.

Ashurst won a competitive pitch specifically because they successfully showed how they "augmented the technology with the expertise." The pitch process involved both firms working on 10 matters over two weeks to demonstrate their AI integration.

Multidisciplinary teams are now essential, with lawyers, project managers, and technologists collaborating before delivering legal services. In-house lawyers must become "translators" bridging law, strategy, and technology, increasingly working on committees with AI heads and information security executives.

The traditional apprenticeship model is reviving as AI handles routine tasks like contracting, legal research, and document drafting, freeing junior lawyers for meaningful mentorship. Three Crowns and Stanford's CodeX project use generative AI for cross-examination skills training through simulations and immersive learning.

Danny Tobey of DLA Piper notes the end of "grunt work fiction" - the myth that doing mundane tasks taught legal skills. However, key tensions remain between speed versus security, with corporate pressure for rapid AI implementation conflicting with data protection requirements.

AI data governance policies have become fundamental to organisational accuracy, with legal advisers increasingly counselling boards and CEOs on AI governance to prevent litigation, regulatory scrutiny, and reputational damage from irresponsible data use.

Microsoft to Rank AI Model Safety for Cloud Customers
7 June 2025 | [Financial Times](#)

Microsoft will add a "safety" category to its AI model leaderboard launched this month, helping tens of thousands of clients using Azure Foundry developer platform "directly shop and understand" AI model capabilities. Current rankings cover quality, cost, and throughput, with safety rankings covering a catalogue of more than 1,900 AI models.

Safety benchmarks include Microsoft's ToxiGen benchmark measuring implicit hate speech and the Center for AI Safety's Weapons of Mass Destruction Proxy benchmark assessing potential for malicious use. No global standard currently exists for AI safety testing, though the EU AI Act will mandate safety tests later this year.

Microsoft positions itself as an agnostic platform alongside its \$14 billion OpenAI investment, offering models from competitors including xAI (Elon Musk) and Anthropic. The company is one of three major "hyperscalers" dominating the cloud market with Amazon and Google.

Growing customer concerns about AI agents working without human supervision create risks of unauthorised actions and legal liabilities. Microsoft launched an "AI red teaming agent" in April to automate vulnerability testing, whilst some model builders (including OpenAI) reportedly dedicate fewer resources to safety testing.

Sarah Bird, Microsoft's head of Responsible AI, emphasises the impossibility of shipping quality models without "huge amount" of evaluation investment. However, consultant Cassie Kozyrkov warns that "safety metrics are a starting point, not a green light," emphasising the need for critical evaluation beyond standardised metrics.

AI Market and Investment

Palantir Technologies: Extreme Valuation Raises Bubble Concerns
5 June 2025 | [FT Alphaville](#)

Palantir Technologies, with a market capitalisation of \$314 billion, now ranks as a top-30 S&P 500 constituent between Coca-Cola and Bank of America, trading at extreme valuation metrics: 565x trailing earnings, 228x forward earnings, and 79.90x enterprise value-to-forecast-sales ratio.

According to Trivariate Research analysis of 2,000 established US non-financial stocks since 2000, only six companies have traded at higher EV-to-forecast-sales ratios. Most expensive comparables include MicroStrategy (248.50x) and Roivant Sciences (70.90x), with five of the 15 most expensive stocks being "moonshot biotechs."

Historical performance data shows companies hitting 30x+ sales multiples typically underperform the S&P 500 by 22.5% over the next year and end trading at just 18x sales. Previous high-multiple companies had poor outcomes: Comverse Technology was defunct by 2013 with its CEO imprisoned, Bluebird Bio was taken private for one-fifth of its 2019 valuation, and Moderna fell 94% from pandemic highs.

Enterprise values of 70x+ forecast sales are "never normal," with most years seeing zero established US stocks hitting 30x sales. High valuations previously occurred only during the dotcom boom and pandemic era. The end-June S&P 500 rebalance will raise Palantir's index weighting above the large-cap threshold, potentially forcing active managers to conduct detailed valuation analysis.

Companies typically enter the "bubble zone" after 45% average annual revenue growth, with analysts expecting repeat performance, but growth typically slows to 28%. This analysis provides a compelling case study in speculative excess and the importance of maintaining analytical rigour during periods of market euphoria.

Mary Meeker's AI Market Analysis (The Batch Issue 305)

Investment analyst Mary Meeker has published "Trends — Artificial Intelligence (May '25)", her first major report since 2019, spanning 340 pages of comprehensive market analysis. The report reveals striking statistics: ChatGPT reached 1 million users in just 5 days compared to iPhone's 74 days, the six biggest technology companies increased capital expenditure by 63% to £169 billion between 2023-2024, training datasets are growing 260% annually, and processing power is expanding 360% annually. Perhaps most remarkably, 73% of humans classified language model responses as human-generated.

The report offers a wealth of market data drawn from analyst reports, consumer surveys, and academic studies. Meeker, who has observed technology markets since the web's dawn, provides valuable perspective alongside existing annual AI surveys from Stanford and Air Street Capital, offering insights into the unprecedented pace of AI adoption and investment.

Rising Costs of AI Model Benchmarking (The Batch Issue 305)

Artificial Analysis has revealed the escalating costs of benchmarking reasoning models, with OpenAI's o1 producing 44 million tokens compared to GPT-4o's 5.5 million tokens across seven benchmarks, costing £2,200 versus £87 respectively. Testing 12 reasoning models cost approximately £4,150 compared to roughly £1,900 for over 80 non-reasoning models, though smaller reasoning models proved significantly less expensive (o3-mini at £275, o1-mini at £112).

These rising costs present challenges for the AI research community, as benchmarks remain critical for relative performance assessment and independent benchmarking ensures fair, consistent testing. As costs climb, fewer laboratories can afford to confirm or challenge results from original developers, making model comparison and progress recognition increasingly difficult for the broader research community.

AQR Capital Management Embraces AI After Years of Scepticism 4 June 2025 | [Financial Times](#)

AQR Capital Management, the Connecticut-based hedge fund with \$136 billion under management, has "surrendered more to the machines" after years of AI experiments. Co-founder Cliff Asness, previously one of the sector's holdouts on AI adoption, has overseen significant performance recovery since the "quant winter" of 2018-2020.

The firm first invested in machine learning technology in 2018 and recently expanded AI strategy beyond stocks to other asset classes. AQR now uses AI to determine portfolio factor weightings in real-time and employs machine-learning algorithms to identify market patterns for trading bets.

Assets dropped from \$226 billion to an eventual low of approximately \$98 billion in 2023, but top hedge fund strategies are now performing well over the past five years. The Apex fund (multi-strategy) shows 19% annualised net returns, whilst Delphi (equity strategy) demonstrates 14.6% annualised net returns as of end-May 2025.

Asness describes their approach as creating a "cloudy and complicated box" rather than "black box," usually finding economic rationale for AI-driven trades whilst acknowledging difficulty explaining AI decisions during poor performance periods. The firm maintains some human oversight while letting "data speak more."

AQR's transformation demonstrates the ongoing evolution of the financial industry and the need for educational institutions to prepare students for a more technologically sophisticated investment landscape while maintaining critical thinking about industry practices and claims.

Cohere Seeks \$500mn Funding to Compete with AI Giants
4 June 2025 | [Financial Times](#)

Canadian AI start-up Cohere is seeking to raise more than \$500 million in new funding, targeting a valuation between \$5.5-6.5 billion, similar to its previous funding round. The company, founded by former Google researchers including CEO Aidan Gomez (co-author of the seminal "Attention Is All You Need" transformer paper), remains significantly behind US competitors in valuation. OpenAI reached a \$300 billion valuation in April 2025 (up from \$157 billion in 2024), whilst Anthropic achieved \$61.5 billion in March 2025 (tripled from its previous round). Cohere's target \$6 billion valuation reflects the challenging competitive landscape for non-US AI providers.

The company focuses on enterprise and privacy-focused solutions rather than consumer apps, doubling annual recurring revenue in the past four months and crossing \$100 million last month. Cohere develops "open" models including the Aya family of multilingual models and recently introduced the North platform for building AI agents for office work.

However, Cohere faces significant challenges competing against tech giants (Google, Microsoft, Amazon) for enterprise contracts and application-layer start-ups like Anysphere (valued at \$2.5 billion). The company operates in a crowded open-source model market alongside Meta, Mistral, and DeepSeek.

For educational institutions, Cohere's emphasis on privacy and enterprise solutions could benefit organisations seeking AI tools with better data protection, whilst their multilingual capabilities could support international education programmes. However, the significant valuation gap suggests potential market consolidation that could limit future AI tool diversity.

Meta Agrees 20-year Deal to Buy Output from Illinois Nuclear Plant
4 June 2025 | [Financial Times](#)

Meta has signed a 20-year agreement with Constellation Energy to purchase the entire output of the Clinton Clean Energy Center in Illinois from June 2027, marking Meta's first nuclear power deal as it seeks electricity for AI development and US data centres.

The agreement will extend the plant's life beyond the expiry of state government subsidies, with Meta joining rivals Amazon, Google, and Microsoft in securing nuclear deals. Big Tech companies are racing to secure vast amounts of electricity needed for the AI boom.

Meta's previous nuclear plans were complicated by environmental issues, including discovery of a rare bee species near a potential data centre site. Constellation operates the largest fleet of conventional nuclear reactors in the US and previously announced reopening Three Mile Island's nuclear plant after a 20-year Microsoft deal.

President Trump signed an executive order to fast-track reactor construction and quadruple US atomic energy capacity by 2050. The Clinton plant was loss-making and scheduled for closure in 2017 but received Illinois state legislature financial support until mid-2027.

Meta's deal will replace government subsidies, removing burden from ratepayers whilst expanding clean energy output by 30 megawatts, preserving 1,100 high-paying local jobs, and generating \$13.5 million in

annual tax revenue. This development highlights the critical intersection between technology advancement and sustainable energy solutions.

UK Medical Robotics Company Aims for \$4bn Sale
3 June 2025 | [Financial Times](#)

CMR Surgical, a Cambridge-based medical robotics company founded in 2014, has put itself up for sale seeking a valuation of up to \$4 billion. The company has hired advisers for a sale process likely to attract interest from large strategic buyers, representing a success story in the UK life sciences sector.

CMR's flagship product is Versius, a small portable surgical robot launched in 2019 that can be moved between operating rooms. Versius has been used in more than 30,000 surgical cases across more than 30 countries, competing with large US rivals such as Intuitive Surgical in the global surgical robotics market expected to generate \$12.8 billion in sales this year.

The company is plotting its launch in the US, the world's biggest surgery market, with the FDA approving Versius for gallbladder removal surgery in October 2024. Earlier in 2025, CMR raised \$200 million from existing backers including SoftBank and Ally Bridge Group to fund US rollout, with GE HealthCare's venture capital arm also an investor.

Many hospitals outside the US have been slow to adopt surgical robots due to costs, though surgical robots are gaining traction for certain gynaecological and urological procedures. The potential sale highlights the UK's challenge of retaining life sciences companies, where start-ups are often sold to foreign buyers rather than growing into large standalone UK companies.

The case demonstrates how UK innovation in medical technology can achieve global reach whilst illustrating the tension between providing early-stage funding and maintaining long-term ownership of successful domestic companies.

News Sites Are Getting Crushed by Google's New AI Tools
10 June 2025 | [Wall Street Journal](#)

"The AI armageddon is here for online news publishers" as chatbots replace Google's traditional search, eliminating the need to click on blue links and devastating traffic to news sites. Google is shifting "from being a search engine to an answer engine," with dramatic consequences for digital journalism.

Traffic declines are severe across major publications: HuffPost's organic search traffic fell by just over 50% in the past three years, Washington Post declined by nearly 50%, Business Insider fell 55% between April 2022 and April 2025, and New York Times' share of traffic from organic search dropped to 36.5% in April 2025 from almost 44% three years earlier.

Google's AI Overviews, launched last year, summarise search results at the top of pages, reducing clicks to vacation guides, health tips, and product review sites. AI Mode's US rollout last month creates chatbot-style conversations with far fewer links, expected to deliver "a stronger blow" to publisher traffic.

Industry responses include significant layoffs: Business Insider cut about 21% of staff with CEO Barbara Peng citing need to "endure extreme traffic drops outside of our control." The Atlantic's CEO Nicholas Thompson told staff to assume traffic from Google would drop toward zero, whilst Washington Post is "moving with urgency" to prepare for a "post-search era."

Publishers are pursuing complex strategies combining legal action and licensing deals. New York Times sued OpenAI and Microsoft for copyright infringement whilst announcing an AI licensing agreement with Amazon. News Corp has content deals with OpenAI but two subsidiaries sued Perplexity.

This represents a fundamental threat to the economic foundations of independent journalism and democratic discourse, requiring educational institutions to teach students about information literacy, source verification, and the importance of accessing original reporting rather than relying solely on AI-generated summaries.