FT PROFESSIONAL

SPECIAL REPORT

AI in Practice

How artificial intelligence is being set to use in — and often fundamentally changing — multiple sectors, including: transforming pharma; spreading factory automation; astronomical opportunities; agentic travel agents.

professional.ft.com



Contents

Businesses face up to budget-busting AI bills 3

AI speeds the march of China's factory robots into new sectors 5

Why AI can transform pharma – but nature cannot be hurried 7

How AI and an astronomer's laptop can bring new galaxies within reach 9

AI drives a boom in new games but big developers dominate 11

Professional Subscriptions 13



Businesses face up to budget-busting AI bills

Written by Jamie John

In the AI whirlwind of recent years, businesses raced to implement the technology amid widespread discussion of its seemingly unbounded potential and ability to automate white-collar work at little expense.

In sectors from manufacturing and pharmaceuticals to tourism and gaming, businesses encouraged staff to experiment with AI, free from the usual constraints on cost.

In Silicon Valley, tech groups such as Amazon built internal leader boards ranking staff on AI usage, convinced this would help integrate the technology into work routines.

But since the start of 2026, a shift has begun to take place. The world's largest AI providers have started to move businesses away from a flat-fee structure and on to usage-based pricing, as they prepare to go public and demonstrate the sustainability of their businesses.

Meanwhile, so-called AI agents, which perform far more complex tasks than prompt-and-response chatbots but also burn through AI credits at a far greater rate, have surged in popularity.

The combined effect has been to raise significantly the cost to businesses of using AI. Uber made waves earlier this year when it said that it had spent its entire AI budget for 2026 by April and other companies have been similarly shocked by the bills they are racking up. Uber is now limiting its employees to $1,500 per month per AI coding tool.

Nearly half of 2,145 global business leaders surveyed by KPMG in May said they had scaled back use of AI agents because costs outweighed the benefits.

"People are really saying . . . 'My company spent my entire 2026 budget in Q1'," said OpenAI chief Sam Altman on stage earlier this month. It "went from, at the beginning of this year, an issue that never came up . . . to, all of a sudden, a huge issue".

While AI use is continuing to grow overall, businesses are now weighing how to use the emerging capabilities of agentic models without breaking the bank. Tech businesses are on the front lines of this conundrum and are coming up with a host of ways to rein in costs.

Enterprise software group Atlassian is one of several companies that has put caps on the number of "tokens" — units of data processed by models — that each of its employees can use over a given period. Staff can ask for more tokens but that requires sign-off from a manager.

"I see a lot of companies out there 'yolo'-ing the whole thing — it's basically, pick the most expensive model, use it as much as you want," says Mike Cannon-Brookes, Atlassian chief executive, of the "you only live once" spending approach.

"It's pretty dangerous because it also teaches very bad habits," he adds.

Another concern for businesses is vendor lock-in, in which a company becomes overly reliant on a single provider for models, making it difficult to switch if prices rise or performance falls.

"It's really threatening to companies that don't have the resources to 'spray and pray'," says Sarah Sachs, AI lead at productivity software group Notion, which, like other software vendors, has built routing tools that help pick the best model for a given task.

Hostinger, a website hosting provider, has been using large language models (LLMs) for six years and says almost all of its 900 employees use them daily. To keep costs down, the company regularly benchmarks the latest AI models and has also built its own routing tool.

"There are a lot of things to think about," says Mantas Lukauskas, Hostinger's AI tech lead. "[Anthropic's] Claude models right now are the best for coding, but for creative writing, OpenAI models are the best," and there are others still that are better for features like image generation, he says.

One solution gaining traction is open-source AI models, which are free to use and can be hosted locally rather than using AI companies' servers. The models are especially popular with groups such as banks, telecom companies and other businesses for whom data security is important, says Costi Perricos, global generative AI leader at Deloitte.

A recent study led by an academic from the Mozilla Foundation, an open-source-focused non-profit, estimated that open-source models achieved about 90 per cent of the performance of their more popular closed alternatives. Reallocating demand towards open-source could cut AI costs for businesses by up to 70 per cent, it said.

But, for now, total spending on AI continues to rise and the world's top AI labs are continuing to release new models touting ever more advanced abilities.

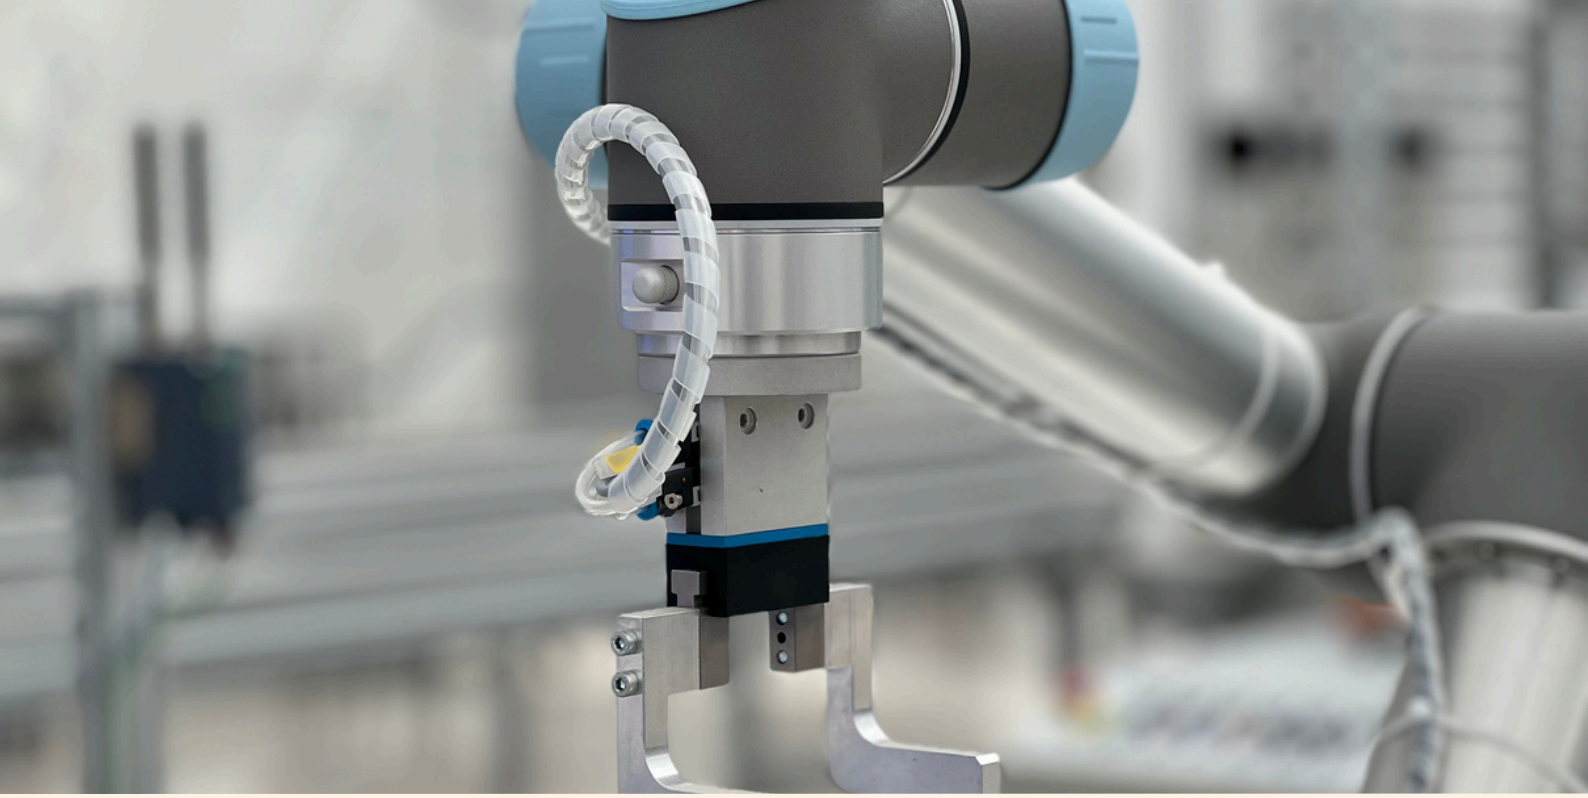
Despite growing controls on spending, Goldman Sachs predicted last month that by 2030, there would be a 24-fold increase in global consumption of tokens, driven by the use of AI agents.

"Model capability is improving so fast that when a new model comes out, there are all these new things you can do," says Duncan Lennox, chief product and technology officer at software developer HubSpot.

"It's a matter of discipline — yes you want room for experimentation and innovation, but the big lesson for us has been to tie it to work outcomes."



The world's largest AI providers have started to move businesses away from a flat-fee structure and on to usage-based pricing, as they prepare to go public and demonstrate the sustainability of their businesses.



AI speeds the march of China's factory robots into new sectors

Written by William Langley

At a cavernous 100,000 sq m factory in the southern Chinese city of Changsha, dozens of automated guided vehicles whisk heavy loads between an array of giant robotic welding arms, machine tool beds and a small smattering of human workers, churning out a new concrete pump truck every 45 minutes.

In the factory's space-age centre, manager Peng Yonghui can oversee each machine's efficiency and state of repair from a giant screen flanked by blue-tiled fountains and palm trees. With a few clicks on a touchscreen, he can monitor machines and manufacturing sites across China.

Sany's No.18 factory, opened in 2012, is one of the planet's largest "lighthouse factories", a World Economic Forum designation given to particularly smart manufacturing bases. With recent improvements in AI, the factory has continued to innovate, says Peng: "it makes us much more efficient".

For instance, it previously took a number of humans to manually sort through the array of metal frames delivered to the factory's eastern end.

But with the advent of AI, the factory's robotic arms can identify the shape and weight of the parts, decide which magnets are needed to lift them and feed them into machine toolbeds for processing, or drop them on automated vehicles to take them to the welding station.

"Originally, I had big parts, long parts, thick parts, small parts and short parts," Peng says. Now, the system can not only discern size, weight and colour, but even the type of material. "Some materials are very similar, almost like twins, but they can still identify it."

Policymakers in China are pushing automation and other smart factory technologies in the hope that they will enable the country to keep its manufacturing edge as its workforce ages.

But now, the rapid advance of AI technologies is improving the efficiency of production lines, boosting quality control and enhancing oversight. Factories, meanwhile, are deploying more intelligent robots on more difficult work. Others are preparing for the arrival of humanoid bots that could perform tasks previously reserved for humans.

FBetter AI models are enabling automation and smart systems to spread beyond the strongholds of large-scale auto manufacturing and other high-end fields into previously labour-intensive smaller factories in traditional sectors such as garments and footwear.

irst, the growth of AI supply-chain industries is generating direct demand for robots in manufacturing," Samantha Mou, a senior analyst at robotics research consultancy Interact Analysis, wrote in a note this month.

"Second, AI-powered software is enhancing robot capabilities and ease of use, enabling robots to perform tasks they previously couldn't, thereby creating demand in markets that were difficult to penetrate," she added.

The introduction of AI has improved industrial robots' visual guidance and made them easier to programme and better at co-ordinating with others. AI has also enabled remote monitoring and maintenance of robots, she wrote.

Robot vendors the world over are reacting. Japan's Fanuc and Denmark's Universal Robots have each announced collaborations with chipmaker Nvidia to work on AI-enabled programming tools. Switzerland's ABB and Shanghai-based Step, meanwhile, have released new robot models designed for the electronics and semiconductor industries, to capitalise on new AI-driven demand.

Back in China, the huge manufacturing sector is driving demand.

"China is a mass market. [Its] industry reach and manufacturing scenarios [are] in the millions. So that generates . . . a real need for robotics," says Jerry Liang, a partner at Beijing-based venture capital firm CCV.

Liang made early bets on a number of Chinese internet groups, including ecommerce retailer JD.com. But in recent years, his team has shifted its focus to AI-empowered robotics to capitalise on the pressing need for more robots in China's economy.

Liang argues that deploying more robots in industrial settings will help to gather data for better physical AI models. Factories could also represent the best chance for China's many humanoid robotmakers to monetise their products at scale, he says.

"Chinese wages, worker wages, are not cheap anymore, and new generations are reluctant to do the boring and dangerous jobs . . . the manufacturing sectors have great anticipation of either humanoid or non-humanoid [robots] to do more complex jobs out there."

In the southern Chinese manufacturing hub of Guangzhou, GSK CNC Equipment was one of the country's earliest industrial conglomerates to branch out into the relatively high-end field of industrial robots, developing its first models in the mid-2000s.

Like many Chinese industrial robot producers, the company's biggest advantage over its Japanese and European rivals is its lower costs, says Huang Yongchao, a sales engineer at the company for 15 years. But he adds that workers in some of the lower-cost industries it sells to can struggle to adapt.

"The biggest challenge . . . is that our equipment is used in those kinds of factories where the work is dirty and tiring and the workers generally have lower levels of education. For tasks involving editing and manipulating text [to instruct robots], they might not be able to get to grips with them at first," he says, standing in front of a phalanx of huge bronze-coloured robotic arms at a trade expo this month.

But Huang is hopeful that AI might have the answer. "They would prefer something simpler, or even something like cutting-edge AI technology, where people interact directly with the computer, which then uses language to tell them how to use it," he says. "They want [the robots] to be able to understand their speech."

Better AI models are enabling automation and smart systems to spread beyond the strongholds of large-scale auto manufacturing and other high-end fields into previously labour-intensive smaller factories in traditional sectors such as garments and footwear.



Why AI can transfrom pharma – but nature cannot be hurried

Written by Aanu Adeoye

Since OpenAI released ChatGPT in 2022, launching a new arms race in the field of generative AI, several industries have been scrambling to figure out how large language model systems could make their workflows more efficient and productive.

The world of pharmaceuticals is not immune. The promises of AI in pharma are plentiful and global groups are investing billions of dollars to outcompete one another, while AI-first start-ups are sprouting up in the industry.

The pharma sector has long embraced technology to improve the process of discovering and delivering new treatments to patients. “In the research space, AI has been used for decades in the industry,” says Thorsten Rall, global industry lead for life sciences at Capgemini, the IT services group. “Pharma is a data-rich environment and applications like AI for [identifying treatment targets] and molecule design have been around for a while and deployed quite broadly.”

A game-changing moment for use of AI was the release of AlphaFold2 in 2021 by Google DeepMind. AlphaFold, which went on to win the Nobel Prize, predicts a protein’s structure from its amino acid sequence.

This “protein-folding” prediction knowledge is crucial to grasping how proteins work and could be essential to understanding previously unknown disease-causing pathogens.

“That was a watershed moment,” says Rebecca Paul, vice-president and head of drug design at Isomorphic Labs, the company spun out of Google DeepMind to advance the project. “That was the first time as a scientist I saw a model that could make very accurate predictions about complex biological output and the progress has been steep from there.”

The progress has included predictions of how proteins interact with other molecules and DNA, Paul says. AI will be used to design molecules that can target proteins once thought “undruggable”, widening treatment options for patients across many indications, she adds.

There is an expectation within the industry that the impact of AI will be enormous. A new report by Capgemini says AI-driven platforms could account for about 60 per cent of new molecular entities within the next decade, up from around 12 per cent of pipelines now.

The way drugs are designed and developed will fundamentally change if the core promises of the new AI era come to pass. The first impact is expected to be in reducing the time it takes to reach the pre-clinical testing of molecules from about four to five years to 12-18 months, according to Rall.

The second benefit will be cost savings, since companies will no longer have to run the same level of lab infrastructure, and time will be saved by recruiting patients more quickly and efficiently for clinical trials.

Industry experts are also excited about the prospects of AI being able to generate synthetic data good enough to skip the early stages of trials in lab animals. This may allow human trials to begin much sooner, although that appears some way off and is subject to regulatory approval.

The rate of failure in drug development has been constant — only one in 10 drugs that begin a clinical trial makes it through the various phases and regulatory approvals. Pharma groups often spend enormous amounts of money on projects that never reach the market.

"If you can use AI . . . in the design phase to come up with molecules that have a better safety profile and reduce the leakage in the funnel, you can have better quality assets," says Rall.

Danielle Belgrave, vice-president of AI and machine learning at British drugmaker GSK, says the advancement in AI has also been beneficial for "processing data at scale". With enormous advances in the understanding of human biology, the data generated needs a commensurate level of computing power, which AI has made possible.

"It's more than an AI revolution, it's been a compute revolution that's enabled us to scale solutions to understanding human biology," she says. "The number one thing with AI is [that] being able to process this data at scale, we're able to do data science in a new way and automate solutions." This "deep understanding" helps to assess whether a treatment might work, "because medicine isn't a one-size-fits-all," she adds.

Since patients with the same diagnosis will often respond to treatments differently, AI can be used to "untangle the [differences] in response and personalise interventions".

Basecamp Research, a London-based AI life sciences company sequencing new genomes collected from the natural environment, aims to develop personalised medicines from knowledge it generates building its genomic database, says chief executive Glen Gowers. "The scope of drug development massively widens with AI if you can understand biology at a different level, away from the limits of what humans can," he adds.

Gowers identifies cell and gene therapies — drugs designed to treat diseases at their genetic or cellular source — as a niche where AI could be most useful for drug development.

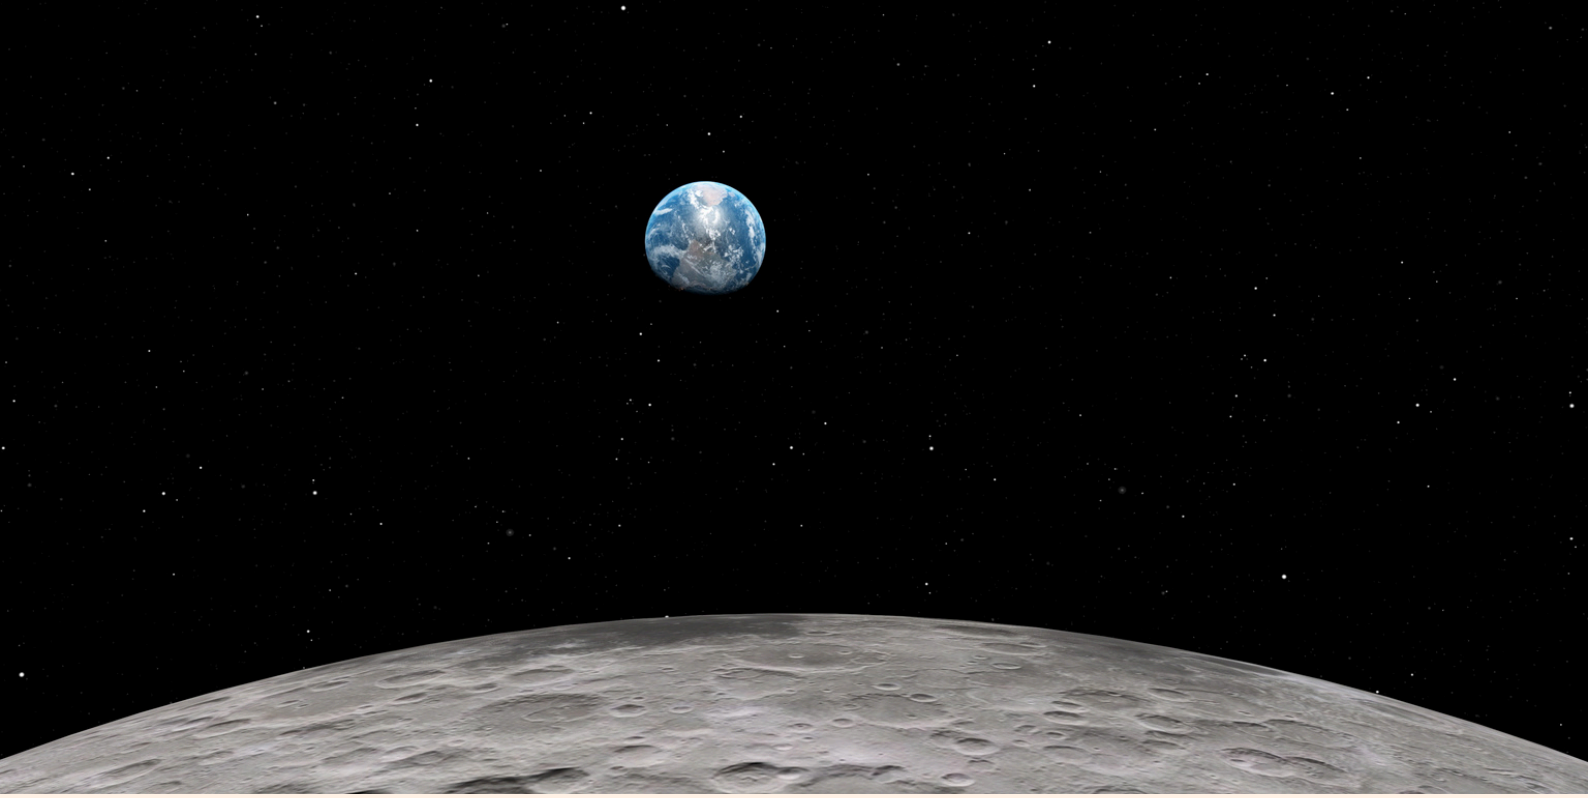
Yet for all the promise of AI, there are questions over how long it will be before a fully AI-developed treatment reaches patients and why it is taking so long.

Paul at Isomorphic Labs says it is inevitable that the process will take time "because we cannot compress the physical world and we have to let biology play out". This means time will still be needed to study the impact and efficacy of a molecule in patients.

"Much as we can make incredible predictions and find things we otherwise wouldn't find, we still have to navigate the physical world — there's always going to be that time lag," she says.

"What we're seeing in the clinic now is a reflection of what AI was doing five to 10 years ago. What we're doing now, which I believe is a huge step change from what we were doing then, is going to be reflected in five, 10 years' time."

The way drugs are designed and developed will fundamentally change if the core promises of the new AI era come to pass.



How AI and an astronomer's laptop can bring new galaxies within reach

Written by Clive Cookson

At the European Space Agency, David O'Ryan and Pablo Gómez have discovered hundreds of new astronomical objects in the vast archive of images transmitted to Earth by the Hubble Space Telescope over the past 36 years.

Using a neural network called AnomalyMatch, trained to detect unusual patterns, they found galaxies galore, with peculiar shapes that astronomers had not previously seen through manual inspection of Hubble data.

"Of course people look at the exciting new generation of telescopes coming online but this work shows how AI can enhance the scientific return from archival data sets," Gómez says. "They are treasure troves in which we can still find very exciting things."

At Oxford university, physicist Héloïse Stevance has led the development of an AI-based Virtual Research Assistant. This has cut by 85 per cent the amount of human "eyeballing" required to distinguish the rare genuine signals caused by supernovae — ultrapowerful explosions of dying stars — among hundreds of alerts received every day from Atlas, a network of five terrestrial telescopes scanning the sky for high-energy events in distant galaxies.

"Manual verification would take several hours each day," says Stevance. "Thanks to our new tool, we can free up scientists' time for . . . creative problem solving and questioning the nature of our universe. It's the astrophysical equivalent of having a bot doing your laundry so you can focus on
your art!"

At Birmingham university Guy Davies, who studies the behaviour and structure of stars, uses AI "emulators" to help interpret telescope observations. "In much of the work I do our biggest challenge is not so much getting the data but building models of stellar evolution fast enough to be able to figure out what we are seeing in the data," he says.

"We have been successful in taking a neural network and teaching it to simulate complex models of stellar evolution," Davies adds. "Training can take weeks or months but once it is trained we can evaluate a model in less than a millisecond."

His Birmingham colleague Anjali Piette depends on emulators for research into the atmospheres of planets beyond our solar system. “If we see a signal in the spectrum of an exoplanet, which we can attribute to a particular molecule, we want to know why it is there,” she says. “We want to use all the nuanced information in the spectrum to link it to different specific processes that we can model. Getting the ‘why’ involves connecting the models and data.”

These examples — making new discoveries in scientific archives; detecting significant signals in incoming data; and building models to make sense of observations — illustrate some of the ways astronomers are enlisting AI as they look forward to a hugely increased flow of data from a new generation of telescopes in space and on the ground.

The greatest torrent of observations will come from the new Vera Rubin Observatory in Chile, which includes an 8.4-metre telescope and a 3bn-pixel digital camera, the largest ever. It will take an image every 30 seconds for 10 years, creating an ultra-wide, high-definition, time-lapse record of the southern sky — “the largest astronomical movie of all time”, as the observatory puts it.

The US National Science Foundation has set up the NSF-Simons AI Institute for the Sky (SKAI Institute) in Chicago to develop new AI tools to detect scientific treasure in Rubin’s output. Astronomers look forward to discovering myriad objects such as asteroids and comets, pulsating stars, supernova explosions and others so novel that they do not yet have a name.

Supercomputers and intensive processing are not required to exploit AI for astronomy. “The surprising thing is how little data it took,” says Stevance of her Virtual Research Assistant. “With just 15,000 examples and the computing power of my laptop, I could train smart algorithms to do the heavy lifting and automate what used to take a human being hours to do each day . . . With expert guidance, AI can transform astronomical discovery without requiring enormous data sets or computational power.”

Gómez made a similar comment about the European Space Agency project to explore the Hubble archive with AI. “We have a very efficient algorithm,” he says. “It ran on a single GPU [graphics processing unit] in a few days. We are not on the scale of using 100,000 GPUs burning gigawatts of energy. If we pick the right methods and inputs, like an efficient AI tool with expert annotations, we can still be sustainable.”

On the whole, astronomers are using more specialised AI tools than the generative AI and large language models from companies such as Anthropic and OpenAI that have caught the public imagination.

“Do you need a general purpose system to answer a specific question in astronomy?” asks Gómez. “In many cases not.”

While astronomers continue to develop and use specialised tools for their research, generative AI will improve their efficiency, says Amaury Triaud, another expert on exoplanets at Birmingham university. “For example it will help us to design user interfaces for our instruments, work online with colleagues around the world and also carry out intricate work on the optical systems of telescopes — aligning mirrors, focusing and tracking your star.”

“

On the whole, astronomers are using more specialised AI tools than the generative AI and large language models from companies such as Anthropic and OpenAI that have caught the public imagination.



AI drives a boom in new games but big developers dominate

Written by Orlando Crowcroft

When Stanislas Marchand was heading up game development for French mobile gaming unicorn Voodoo, his team tested 3,000 new titles per year.

"It was an industrial process," says Marchand, who was there until 2024. "Out of those, only one or two would go on to make $10mn-plus per year."

Those 3,000 ideas came from Voodoo's internal teams as well as independent studios and, even after the advent of generative AI, productivity savings were relatively modest, he says. Pre-AI, it took about 14 days to design, test and get a mobile game ready for market. With AI, the team took 10 days.

That appears at odds with predictions across many industries in 2026, as adoption of AI leads to grim forecasts about the future of work. Despite some lay-offs, the view of many in gaming is that AI will struggle to replace the central role of human creativity.

"What makes a good game is what we call 'game feeling'," he says. "So far, AI cannot do that," Marchand says.

Nonetheless, AI — and particularly the onset of "vibe-coding", which allows people with little knowledge of software engineering to use AI to create concepts and games — has led to a massive increase in games coming to market. Global data from research company ATTN Economy shows that 181,000 games were released in the six months to May, 43,500 on iOS and 137,000 on Android — increases of 118 per cent and 73 per cent respectively year on year.

This increase has come alongside some significant job losses in the gaming industry. Microsoft lost staff in its XBox division this month and cut 10 per cent in its Stockholm-based King division, famous for Candy Crush. A report by the GDC Festival of Gaming in May found that one in four employees in gaming had been laid off in the past two years.

The report found that 52 per cent of gaming professionals thought generative AI was bad for the industry, up from 30 per cent in 2025 and 18 per cent in 2024. Workers in visual and technical arts (64 per cent) and game design and narrative (63 per cent) had the most unfavourable views. Only 7 per cent said that AI was having a positive impact.

Some fans are also unhappy. The much-anticipated Crazy Taxi: World Tour was embroiled in online controversy over whether generative AI had been used in its development. This prompted Sega to acknowledge that it had utilised generative AI "as a support tool for developers [but] no AI was used in reference to performers in the game".

Those more bullish on generative AI see performers — or non-playing characters (NPCs) — as the most exciting use-cases. Where interactions with characters in role-playing and artificial world games such as Zelda or Call of Duty were heavily scripted, AI raises the prospect of true player-to-NPC dialogue.

"Non-playable characters are already getting more powerful, but these things are about to become deeply personal. They're going to remember you, and remember that you played this game two months ago," says Josh Chapman, co-founder and managing partner at Konvoy, a venture capital firm that invests in early-stage gaming start-ups.

Vlastimil Venclik, co-founder and chief executive of Czech gaming start-up Valka AI, agrees, and says there has been an exponential increase in the power of AI within games compared with two years ago. As models improve, so will storytelling and dialogue — and attitudes to AI may change.

"I see some gamers are against it, but gaming studios are all using AI," says Venclik. "I believe [their opposition is] just a trend which can fade away as fast as . . . it appears. In five, 10 years, my kids . . . probably won't know the difference, right?"

Others are sceptical. ATTN Economy's Lexi Sydow notes that data from Klaviyo, a digital marketing platform, shows that 13 per cent of consumers trust AI. When they see marketing content has been created with AI, they are four times less likely to trust that brand.

"There's a cost to AI slop and it comes with a hefty price tag: trust," Sydow says. "In gaming, it is like that on steroids. If it looks like it is AI-driven and not human-driven, it could fracture trust . . . if the storyline or the art drifts, people get really upset."

The right approach will be key for the few large gaming conglomerates that dominate the industry. Sydow notes that the top 1 per cent of publishers controlled $75.6bn of revenue in 2025 with the other 99 per cent collectively making $6.1bn. That top 1 per cent saw 40.2bn downloads in 2025, accounting for nearly 80 per cent worldwide.

Vibe-coding may have lowered the technical bars to entry in game development but Chapman predicts this will not break the big companies' dominance. "The incumbents have such large balance sheets [and] are building up incredible talent and decades of data. That is really, really hard to compete with if you're an up-and-coming start-up with a $4mn seed round."

From Slack workplace collaboration to Helsing military technology to Elon Musk's SpaceX, gamers and gaming founders have played oversized roles in the wider economy. AI or not, that will continue to be the case.

"The gaming market is the most competitive market in the entire planet . . . no other industry has 10,000 new [start-up] competitors a year," says Chapman. "Gaming technologies proliferate into other markets and some of the world's best technologies are being built by gaming founders."

"

Vibe-coding may have lowered the technical bars to entry in game development but Chapman predicts this will not break the big companies' dominance.

FT PROFESSIONAL

Turn trusted journalism into actionable intelligence.

FT Professional helps teams go deeper than the typical FT.com experience, with exclusive professional tools that surface relevant analysis, explore connected topics and provide context in real time.



Inform decision making
with a trusted source of business intelligence



Improve client knowledge
using broad sector coverage and deep analysis



Identify risks and opportunities
by leveraging global, connected reporting

An FT Professional Subscription delivers the commentary and analysis organisations need, along with the tools and services to make that intelligence actionable:



Premium access to FT.com
Access to all FT.com journalism at your fingertips, meaning you and your team won't miss a thing



Workflow integration and apps
The FT on the systems and devices you use most. Keep informed with mobile apps, a Teams integration and availability on platforms like Bloomberg



Professional-only tools and features
Use exclusive features like Ask FT and Article Insights to make FT journalism more useful for work or learning



Enhanced support and services
Optimise use of the FT in your organisation with the help of a dedicated Customer Success Manager, and tools to manage your subscription effectively

If you're interested in learning more about FT Professional Subscriptions, visit our website.