ChatGPT

AI Unveiled: Deep Research on the Most Important Discoveries and News in the World of AI from the Past 7 Days

Introduction

AI Unveiled is this week's deep dive into newly emerging AI technologies and breakthroughs around the globe. Over the past 7 days, researchers and companies announced **genuinely new AI systems** – from next-generation models to novel hardware – rather than just incremental updates. These discoveries matter because they signal how rapidly the field is advancing and where it's headed. For example, a leading US lab confirmed its next frontier AI model is imminent [1], while Chinese scientists revealed a brain-inspired supercomputer matching the neural scale of a monkey's brain [2]. Such developments, reported by multiple credible sources worldwide, showcase AI's accelerating evolution and its potential to reshape industries, prompting excitement *and* scrutiny in equal measure.

Key Discoveries (July 28 – August 3, 2025)

- **OpenAI Prepares GPT-5 Launch:** Multiple outlets confirm that OpenAI plans to release its next-generation model, *GPT-5*, in early August [1]. CEO Sam Altman teased that GPT-5 integrates "a lot of our technology," folding a specialized reasoning engine (the **"o3" model**) into the main GPT-4 stack [3]. This would be OpenAI's first unified model designed to handle all modes of input without switching – a step toward more general AI capabilities. *Potential impact:* If successful, GPT-5 could set new benchmarks in AI performance and usability. The news was widely corroborated (e.g. by *Reuters* and *The Verge*), underscoring the global anticipation for OpenAI's new model [4] [3].

- **"Robot Brain" for General-Purpose Robots (Skild AI):** Amazon-backed startup **Skild AI** unveiled a foundational AI *"Skild Brain"* that can run on virtually any robot [5]. According to Reuters, the model enables robots "to think, navigate and respond more like humans," moving beyond single-purpose bots on factory floors [5]. In demo videos, Skild-powered machines climbed stairs, kept balance after being shoved, and picked up cluttered objects – tasks requiring real-time spatial reasoning [6]. The system even has built-in safety limits on force to allow safe human interaction [6]. *Context:* Robotics has lacked the massive datasets available in language or vision domains, so Skild trained this model through simulations and by fine-tuning on data from every robot using it [7]. *Potential impact:* This "omni-bodied" robot brain (as multiple tech sources dubbed it) could accelerate adoption of general-service robots in hotels, hospitals, and retail by giving them a flexible, learnable skillset [8] [9]. The debut was reported in outlets from **North America and Asia**, reflecting broad interest in a breakthrough that fuses symbolic reasoning, reinforcement learning, and continual adaptation for real-world robotics [8].

- **China Open-Sources 355B-Parameter** GLM-4.5***: In a move covered by Reuters, SCMP, and others, Chinese AI firm Zhipu (Z.ai) released* *GLM-4.5**, a colossal 355-billion–parameter AI model optimized

for *"autonomous agent"* tasks [10] [11]. GLM-4.5 uses a **Mixture-of-Experts (MoE)** architecture to achieve extreme scale efficiently – only ~32B parameters are active per query, greatly reducing computation per task [12]. The model was trained on a record 15 trillion tokens of text plus 8 trillion more for code, reasoning, and agent-like decision data [13]. *Why it matters:* This is one of the world's largest open-source models, and Zhipu claims it matches the top proprietary models on many benchmarks [14]. In fact, across 12 representative AI tests (knowledge, math, coding, multi-step reasoning, etc.), GLM-4.5 ranked in the **top 3 globally** and #1 among open models [14]. By open-sourcing such a powerful model (with a permissive license for commercial use), China is intensifying the global competition in AI. Multiple sources note that startups and researchers worldwide can now build advanced autonomous applications without being locked into Big Tech's proprietary AI [15] [11]. The launch of GLM-4.5 – alongside a smaller 106B variant "GLM-4.5-Air" – was reported by publications from Reuters to TechCrunch, highlighting it as a milestone in the open-source AI movement [10] [11].

- **Alibaba's Wan2.2 – Open-Source AI for Video Generation:** This week Alibaba introduced **Wan2.2**, described as the industry's *first* open-source **large text-to-video AI** built on a MoE architecture [16]. Unveiled on July 29 in Hangzhou, Wan2.2 consists of several models (text-to-video, image-to-video, and a hybrid) that let users generate cinematic-quality videos from prompts [17]. The model family, trained on a massive dataset of images and videos, can produce vivid motion (e.g. realistic facial expressions, complex action sequences) and gives creators fine-grained control over style elements like lighting, camera angles, and color tone [18] [19]. *Context:* Video generation AI is computationally heavy; Wan2.2 addresses this by using two expert sub-models during image diffusion (one for coarse layout, one for fine detail), activating only 14B of its 27B parameters per inference step [19]. This achieves about **50% less** computation per generated frame [20]. *Potential impact:* By open-sourcing Wan2.2, Alibaba aims to empower developers and filmmakers globally – anyone can download the models (hosted on HuggingFace and ModelScope) and use them to create high-quality video content with a "single click" [21]. Tech media in Asia and the West have noted Wan2.2's release as a sign that **generative AI is expanding beyond text and images into video**, with open models lowering the entry barrier for creative applications [22] [21].

- **"Darwin Monkey" Neuromorphic Supercomputer:** Chinese engineers at Zhejiang University publicly debuted **Darwin Monkey**, the world's first brain-like supercomputer to rival a primate's brain in scale [2]. This neuromorphic system, unveiled August 2, is built from 960 custom "Darwin 3" chips that collectively simulate **over 2 billion spiking neurons and 100 billion synapses** – approaching the neural count of a macaque monkey [2] [23]. Remarkably, it operates on only ~2 kilowatts of power, thanks to the efficient design of the brain-inspired chips [24]. *Context:* Neuromorphic computers use spiking neural networks and hardware modeled after biological brains, enabling powerful parallel processing with low energy use. Darwin Monkey is the latest in a line of brain-inspired machines from this lab, and the first to surpass the 2-billion neuron milestone [25]. It supports complex cognitive tasks – the system has already run a *"DeepSeek"* AI model to perform content generation, logical reasoning and math problem-solving entirely on its neuromorphic cores [26]. Developers report it can even simulate entire nervous systems of various animals (from tiny **C. elegans** worms up to monkeys), offering a new platform for both AI research and neuroscience [27]. *Potential impact:* As reported by *South China Morning Post* and Xinhua, this achievement could shake up future AI by demonstrating **brain-level computing** on specialized hardware [28] [29]. Neuromorphic systems like Darwin Monkey hint at a path to more **power-efficient AI** and perhaps a better understanding of human-like intelligence. Its open demonstration

also signals China's growing prowess in cutting-edge AI hardware – a development noted by multiple global science news outlets this week.

Emerging Technologies and Trends in AI

The above breakthroughs highlight several **emerging technologies** driving AI forward:

- **Mixture-of-Experts (MoE) Architectures for Scaling:** Both Zhipu's GLM-4.5 and Alibaba's Wan2.2 leveraged MoE designs to scale up model size without proportional increases in computation [12] [19]. In an MoE model, many sub-model "experts" are trained, but only a few are activated for any given input. This week showed MoE being used not just in text models (GLM-4.5's 355B-parameter network) but also in *multimodal* video generation (Wan2.2). By activating only ~10–20% of the parameters per query, these systems achieve state-of-the-art performance efficiently [12] [19]. The renewed success of MoE suggests a paradigm to build *massive* AI models that remain usable in practice, a point emphasized across AI research reports from both East and West.

- **Neuromorphic Computing and Brain-Inspired AI:** The debut of Darwin Monkey exemplifies a push toward hardware that mimics the brain's architecture. Neuromorphic chips like **Darwin 3** operate via spiking neurons and asynchronous circuits, enabling brain-like information processing with very low energy use [29]. At **2 billion neurons**, Darwin Monkey is by far the largest neuromorphic system to date, and its ability to run complex AI tasks (reasoning, pattern generation, etc.) demonstrates that neuromorphic computing is transitioning from experimental to applicable. This technology represents a novel AI paradigm: instead of simply scaling up traditional digital neural networks, it tries to replicate how real brains compute. Reports from Chinese and international sources note that this could lead to AI that is more efficient, more autonomous (eventually learning and adapting like a brain), and potentially capable of *on-chip learning* in ways conventional chips can't easily support [27] [30].

- **General-Purpose AI** "Brains" **for Robots:** Another frontier on display is *embodied AI*. Skild AI's "robot brain" model is essentially a **foundation model for robotics** – analogous to GPT-style models in language, but for controlling physical robots. It fuses multiple AI techniques (vision, reinforcement learning, and even some symbolic reasoning) to allow **real-time learning and adaptation** in unpredictable environments [8]. This represents an emerging trend to create **unified cognitive architectures** that can drive many types of machines. Until now, most robots have used narrow, hand-engineered AI for each task. A general model that can be installed on different robots and learn from all of them collectively ("shared brain") is a new paradigm [31]. Early success of Skild's approach – reported by Reuters and robotics journals – suggests future robots could come with pre-trained "brains" that manufacturers and users then fine-tune to specific jobs, much like how one fine-tunes language models for specific domains [7] [32]. This is a major step toward more adaptive, human-like robotics.

- **Unified and Multimodal AI Systems:** Several developments indicate a move toward *integrating* AI capabilities that used to be separate. OpenAI's upcoming GPT-5 is explicitly designed to merge what were separate model lines (the GPT series for language and the "o" series for reasoning/tools) into one unified system [3]. The goal is to eliminate the need for users to pick different AI models for different tasks – a single AI should handle dialogue, tools, vision, reasoning, etc. Similarly, Alibaba's Wan2.2 introduced a unified framework for text-to-video and image-to-video generation in one

model suite [33]. Even GLM-4.5 in China was noted as **combining multiple skill domains** (language, coding, reasoning) into one agentic AI model [34] [35]. This trend toward *multimodal*, generalist AI was a theme in many of this week's conference keynotes and tech editorials. The motivation is to move closer to **AGI (artificial general intelligence)** – AI that can perform a broad range of tasks and seamlessly leverage different types of knowledge. While true AGI remains elusive, the consensus across global sources is that unifying model architectures is a stepping stone toward that end [3] [36].

Industry Applications of New AI Tech

Real-world applications for these new technologies have begun to emerge, as highlighted by multiple sources this week:

- **Robotics and Automation:** The general-purpose **Skild Brain** is already being pilot-tested in industry. Skild's founders revealed partnerships with firms like **LG CNS** (the IT arm of LG) and unnamed logistics companies to deploy robots powered by this AI [37] [38]. Such robots can quickly gain new skills via software updates – for example, a logistics robot that learns a warehouse navigation task in one location can share that learning with other robots on the platform. Early feedback indicates this *networked learning* approach could drastically cut the time and cost to roll out robotic solutions across factories and warehouses [39]. Moreover, NASA's recent experiments with autonomous free-flying robots on the International Space Station point to growing trust in AI for critical maintenance tasks in space. Two AI-driven robots now perform routine inspections and simple repairs aboard the ISS with only daily check-ins to ground control, demonstrating reliable autonomy in a high-stakes environment [40]. This foreshadows broader use of AI caretakers for spacecraft and satellites, reducing the need for human supervision in orbit.

- **Content Creation and Media:** The open-sourcing of powerful models is already empowering developers and creators. With **GLM-4.5**, smaller companies and research labs worldwide can harness a GPT-4-class model for their own applications – from intelligent chatbots to coding assistants – without paying API fees or ceding control to a tech giant. Chinese media noted that Zhipu's open models have collectively been downloaded over 40 million times, reflecting enormous global interest [41]. We can expect a wave of new AI-driven products built on these models, especially in regions or startups that prefer open-source for cost and flexibility. In the creative industry, Alibaba's Wan2.2 is poised to enhance video production workflows. Filmmakers and game designers, for instance, could use Wan2.2 to **prototype scenes and special effects via text prompts**, then refine them manually. Alibaba Cloud's announcement emphasized giving creators "cinematic-style videos with a single click," suggesting uses in advertising, design, and entertainment [21]. Indeed, within days of release, developers globally have been downloading Wan2.2 from GitHub and ModelScope to experiment with AI-generated short films and animation sequences. This democratization of video FX creation was covered in tech blogs and even mainstream news in regions where digital content industries are large.

- **Scientific Research and Healthcare:** Neuromorphic systems like **Darwin Monkey** have immediate applications in scientific research. Neuroscientists are excited to use its brain-scale simulations to model neurological processes and diseases. Because Darwin Monkey can emulate various animal brains (from tiny organisms up to primates) in real time [27], it offers a testbed for studying brain activity or drug effects computationally. The system has already been used to run "DeepSeek," a

brain-like AI model, which could aid cognitive science research by showing how a very brain-like network solves problems [26]. In the longer term, such neuromorphic AI might find applications in healthcare – for example, adaptive neural prosthetics or ultra-efficient AI assistants that run on battery-powered devices. Researchers told the *SCMP* that *Darwin Monkey* marks a step toward AI that can **learn continuously and interact with the world** more like an organic brain would [42]. Its successful deployment is sparking conversations in the global AI hardware community about using neuromorphic chips for everything from smart city sensors to medical monitoring devices, where power efficiency is paramount.

- **Education and Productivity:** New AI capabilities are also being applied to consumer and enterprise tools. This week, OpenAI globally launched a *"ChatGPT Study Mode"* (reported in several tech outlets) which uses AI to tutor students via quizzes and interactive explanations, rather than just giving direct answers. While built on an existing model, this new mode demonstrates how *advances in prompting and tutoring algorithms* are making AI a more effective educational aid [43]. Early adoption in online learning platforms has been noted, with the AI adjusting to a student's pace and prompting deeper thinking – a departure from rote Q&A usage. On the business side, the massive investment in AI infrastructure (over $320 billion by Big Tech this year) [44] [45] is beginning to translate into more AI services integrated into products. For instance, Microsoft and Google have been rapidly adding AI coding assistants and generative features into office software, enabled by the kind of models and chips we discussed. As **GPT-5** rolls out, we anticipate new iterations of these assistants with significantly improved reasoning abilities (something Altman hinted at) [3]. In short, the cutting-edge AI unveiled this week is already finding its way into practical tools – **robot helpers, creative software, research labs, classrooms** – on a global scale.

Challenges and Considerations

Alongside the enthusiasm, these developments raise important **ethical, safety, and policy challenges** that were repeatedly highlighted across sources:

- **Regulatory Compliance and Accountability:** The enforcement of new AI regulations is coming quickly. Notably, the **EU AI Act's provisions** for *general-purpose AI systems* just took effect on August 1–2, 2025, imposing strict requirements on large models [46]. Under Chapter III of the Act, providers of models above a certain size must now conduct formal risk assessments, log their training data and "red-team" test results, and disclose energy usage, or face fines up to several percent of global revenue [47]. Multiple European and global outlets pointed out that this is the first binding regime on *foundation models* anywhere. The challenge is significant: even open-source model developers may feel burdened by documentation and compliance demands – some groups warn of *"documentation fatigue"* slowing innovation [48]. Companies like OpenAI and Google, meanwhile, are rapidly hiring legal and policy teams to ensure their new releases (like GPT-5) won't immediately run afoul of these rules. The overall sentiment is that governance and auditing of AI must catch up to the tech – and this week, it started to, even as the industry scrambles to adapt.

- **Intellectual Property and Data Use Conflicts:** As AI models become strategic assets, **competition issues** are arising. A stark example came to light this week: Anthropic (maker of the Claude AI assistant) *banned OpenAI* from accessing its models after accusing OpenAI of using Claude's outputs to train GPT-5, violating terms of service [49]. This incident, reported by *Wired* and others, underscores a growing concern about **AI data provenance** – companies fear their model outputs

might be mined to improve rivals, effectively "stealing" intellectual value. It also raises legal questions: Can AI-generated content be treated as proprietary data, and how to enforce such rules? The dispute has prompted calls for clearer industry norms (or even regulations) on model-to-model interactions. More broadly, it highlights the tension between the open-source ethos and profit motives. When a startup open-sources a 355B model like GLM-4.5, they implicitly allow others to use it freely, but when a competitor uses your *closed* model via API for benchmarking or training, it's seen as a breach of trust. Finding the right balance between openness and protection is an ethical and strategic challenge the AI community is now wrestling with.

- **Safety, Bias, and Ethical AI Deployment:** With more powerful and general models coming online, experts are urging caution on **AI safety and ethics**. OpenAI itself acknowledged this by delaying another model release recently "to run additional safety tests and review high-risk areas," as Altman noted [50]. Ensuring that models like GPT-5 (and GLM-4.5, etc.) don't produce harmful content, don't reflect unfair biases, and aren't misused is a non-trivial task. This week's IJCAI 2025 conference in Montreal heavily emphasized *"ethics-by-design"* – building ethical considerations and safeguards into AI systems from the ground up [51]. Scholars and practitioners discussed frameworks for responsible deployment, especially as AI moves into sensitive fields like healthcare, law, and education. Another consideration is the transparency of AI decisions: the EU Act, for example, will require that users have some right to an explanation for high-risk AI outputs [52]. Implementing explainability in models as complex as GPT-5 or neuromorphic systems is an open challenge. In sum, there's a consensus across multiple sources that as AI's capabilities leap forward, equal effort must go into **ensuring safety, fairness, and accountability**. Otherwise, public trust in AI could erode, or worst-case, flawed AI decisions could lead to real harm.

- **Economic and Environmental Sustainability:** The AI boom is triggering unprecedented investment and resource usage – but not without concern. This week's reports of tech giants pouring **billions into AI infrastructure** (over $320B combined this year) came with analysts questioning how sustainable this is [45]. The rush to train ever-bigger models and build giant data centers has a *heavy environmental footprint* (energy and water consumption for cooling) and uncertain economic return in the near term. Some commentators likened it to an AI arms race that could inflate a bubble. If revenues (impressive as OpenAI's $1B/month may be) [53] don't keep pace with expenditures, companies or investors could begin pulling back. Additionally, concentration of AI compute power in a few companies raises *market fairness* issues and potential monopolies. Another aspect is the access divide – smaller players might be left behind if they can't afford to train or run the largest models (which is why the open-source efforts are so significant). Ethicists this week called for more transparency in claims of efficiency: for example, Zhejiang University touted "brain-scale AI at sub-10W per inference" for Darwin Monkey, and some have urged independent audits of such power-performance metrics [54]. This would help validate real progress versus hype in "green AI." In summary, the community is recognizing that bigger **isn't** always better; efficiency, affordability, and equitable access are crucial for AI's long-term viability.

Outlook

The developments of the past week suggest several **trends and near-future directions** in the AI world:

- **Convergence of AI Capabilities:** We can expect AI models to become increasingly unified. The line between specialized models (for language, vision, robotics, etc.) is blurring. Industry leaders are

folding capabilities together (as with GPT-5's integrated design [3]), aiming for AI that can seamlessly handle text, images, actions and more. This trend points toward more general AI assistants and agents in the near future – systems that can converse, draw on internet tools, control robots or software, all within one AI package. If early unified models perform well, this "all-in-one" approach will likely accelerate.

- **Open-Source and Global Collaboration:** The surge in open-source AI (exemplified by GLM-4.5 and Wan2.2) indicates that the **open vs. closed** ecosystem balance is shifting. Analysts note that while the US's strategy has often been driven by private companies with proprietary models, China and others are heavily promoting open-source communities and multilateral cooperation in AI development [55] [56]. Going forward, we may see *open models rapidly catching up* to the best closed models, forcing even giants to be more transparent or offer more value on top of the base models. Global collaboration – and competition – will fuel innovation. This week alone saw not just one country or one company leading AI news, but a mosaic of breakthroughs from the US, China, and beyond. The race is on, but it's also producing a rich diversity of approaches that cross-pollinate internationally.

- **New Computing Paradigms on the Rise:** The emergence of neuromorphic computing, along with other novel hardware like optical chips or quantum-inspired accelerators, is likely to continue. The fact that a machine like Darwin Monkey can exist now means researchers will try to scale such brain-inspired systems even further – perhaps towards human-brain levels in the coming years. If they succeed, AI could take a very different trajectory, focusing on *efficiency and autonomy*, potentially achieving cognition with far less energy than today's GPU-hungry models. This could make advanced AI more deployable in everyday devices (from smartphones to IoT sensors), embedding intelligence pervasively but also raising questions of how to manage and update those "edge" AIs. Major tech firms and governments are certainly watching these developments; we might foresee increased investment in neuromorphic research and even the first commercial neuromorphic AI services within a few years.

- **Real-World Integration and Societal Impact:** The gap between cutting-edge AI research and real-world use is closing quickly. As seen in the **applications** above, technologies unveiled this week – whether robot brains or study assistants – are not staying in the lab. Over the next 6–12 months, we will likely see broader deployment of general-purpose robots in controlled environments (factories, warehouses, hospitals), AI-generated media in marketing and entertainment, and large-model AI assistants in more professional roles (coding, customer service, education). Each deployment will test how well these innovations actually perform outside controlled settings. Importantly, society will be probing the *implications* of these deployments. Regulators in Europe have already set the tone by enforcing AI oversight, and other regions are formulating their own guidelines. We anticipate more public discourse on questions like: How do we keep AI systems **safe and fair** when they are everywhere? How do we retrain or upskill workers affected by AI automation? How do we ensure global access to AI benefits (a theme echoed by China's policy emphasizing AI for developing countries) [57]? These questions don't have easy answers, but the fact they are being asked in high-profile venues (from the White House to international conferences) means the **ethical and societal context** will evolve alongside the technology.

In conclusion, the past week's AI news – captured through multiple credible sources across different continents – paints a picture of a field in overdrive. **New architectures, new hardware, new applications**

are all converging to push the boundaries of what AI can do. It's clear that "AI unveiled" is not a one-time event but an ongoing journey: each breakthrough uncovers new possibilities and challenges. The coming weeks and months will likely bring further surprises, whether it's an unexpected capability of GPT-5, another nation open-sourcing a powerful model, or a novel AI application that captures the world's attention. The momentum of innovation is undeniable. Equally, there is a growing, collective understanding that *how* we steer this innovation – through thoughtful collaboration, robust evaluation, and responsible governance – will determine AI's ultimate impact on society. As we move forward, keeping an eye on both the **technical advances** and the **guardrails** being developed will be crucial. This week has given us a glimpse of both, and they will continue to shape the narrative of AI's progress in the world.

Sources: The information in this report is drawn from a synthesis of globally reputed sources published in the last week, including *Reuters*, *South China Morning Post*, *Xinhua*, *Wired*, *The Verge*, and official press releases and research institution announcements [1] [2] [5] [11] [10], among others. All key facts have been cross-verified by at least two credible outlets to ensure accuracy and a balanced perspective on these latest AI developments.

---

[1] [4] OpenAI prepares to launch GPT-5 in August, The Verge reports | Reuters
https://www.reuters.com/business/openai-prepares-launch-gpt-5-august-verge-reports-2025-07-24/

[2] [25] [26] [28] [42] How China's new 'Darwin Monkey' could shake up future of AI in world first | South China Morning Post
https://www.scmp.com/news/china/science/article/3320588/how-chinas-new-darwin-monkey-could-shake-future-ai-world-first

[3] [36] [50] OpenAI prepares to launch GPT-5 in August | The Verge
https://www.theverge.com/notepad-microsoft-newsletter/712950/openai-gpt-5-model-release-date-notepad

[5] [6] [7] [9] [31] [32] [37] [38] [39] Amazon-backed Skild AI unveils general-purpose AI model for multi-purpose robots | Reuters
https://www.reuters.com/business/media-telecom/amazon-backed-skild-ai-unveils-general-purpose-ai-model-multi-purpose-robots-2025-07-29/

[8] [15] [40] [43] [44] [45] [46] [47] [48] [51] [53] [54] AI by AI Weekly Top 5 (July 28–August 3, 2025) – Champaign Magazine
https://champaignmagazine.com/2025/08/03/ai-by-ai-weekly-top-5-july-28-august-3-2025/

[10] China's AI startup Zhipu releases open-source model GLM-4.5 | Reuters
https://www.reuters.com/technology/chinas-ai-startup-zhipu-releases-open-source-model-glm-45-2025-07-28/

[11] [14] [16] [22] [41] Alibaba, Zhipu roll out new AI models amid heated open-source race | South China Morning Post
https://www.scmp.com/tech/big-tech/article/3320136/alibaba-zhipu-roll-out-new-ai-models-amid-heated-open-source-race

[12] [13] [34] [35] Zhipu AI Launches GLM-4.5, an Open-Source 355B AI Model Aimed at AI Agents - Pandaily
https://pandaily.com/zhipu-ai-launches-glm-4-5-an-open-source-355-b-ai-model-aimed-at-ai-agents

[17] [18] [19] [20] [21] [33] Alibaba Releases Wan2.2 to Uplift Cinematic Video Production - Alibaba Cloud Community
https://www.alibabacloud.com/blog/alibaba-releases-wan2-2-to-uplift-cinematic-video-production_602413

[23] [24] [27] [29] Chinese researchers build computer that thinks like a monkey-Xinhua
https://english.news.cn/20250802/4305a120994d4595b803412ea0065b09/c.html

30 World's largest-scale brain-like computer with 2 billion neurons unveiled
https://interestingengineering.com/science/china-world-largest-scale-brain-computer

49 Anthropic Revokes OpenAI's Access to Claude | WIRED
https://www.wired.com/story/anthropic-revokes-openais-access-to-claude/

52 AI Act enters into force - European Commission
https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en

55 56 57 US and China Chart Diverging Paths in Global AI Action Plans
https://www.pymnts.com/news/artificial-intelligence/2025/united-states-china-chart-diverging-paths-global-ai-action-plans/