

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

Introduction: A Week of Foundational Shifts

The past seven days in artificial intelligence have been marked not by incremental updates to familiar technologies, but by the unveiling of foundational shifts in the core components of the AI ecosystem. This report, based on a global survey of peer-reviewed research, official corporate announcements, and corroborated reporting from leading technology outlets, details a series of interconnected breakthroughs that signal a new chapter in AI's development. The central theme, "AI Unveiled," reflects the emergence of genuinely new paradigms in three critical domains: the software architectures that power AI reasoning, the specialized hardware that executes its complex computations, and the application of AI to the very process of scientific discovery itself.

This period's key discoveries represent a departure from the monolithic dominance of a single architectural and hardware paradigm. Specifically, this report details:

- **An Architectural Evolution:** The rapid maturation of State-Space Models (SSMs) as a computationally efficient and powerful alternative to the Transformer architecture, which has defined the last half-decade of AI progress. New research has demonstrated that SSMs can overcome the core scaling limitations of Transformers, potentially heralding a new era in sequence modeling for tasks involving extremely long contexts like genomics and high-resolution video.¹
- **A Hardware Realignment:** Strategic hardware announcements from Intel and Alibaba signal a direct and multi-front challenge to Nvidia's market dominance. Driven by a confluence of competitive ambition and geopolitical necessity, these new chips are not mere copies but represent distinct strategies—a high-end, integrated system approach from Intel and a self-reliant, inference-focused approach from Alibaba—that promise to reshape the global AI hardware landscape.⁴
- **A Scientific Application Frontier:** Groundbreaking AI tools have been unveiled that tackle fundamental challenges in public health and research integrity. These systems,

such as MIT's VaxSeer for predictive vaccine design and a new tool for vetting academic journals, demonstrate AI's transition from an industrial tool to a foundational instrument of science itself, capable of both accelerating discovery and safeguarding its integrity.⁷

These trends are deeply intertwined. The demand for new, efficient architectures like SSMs will inevitably influence the design of next-generation hardware. Simultaneously, the geopolitical fracturing of the hardware market will shape which architectures and applications flourish in different global spheres. The developments of the past week are therefore not isolated events but the first tremors of a systemic realignment in the world of artificial intelligence.

Key Discovery: The Post-Transformer Era Begins with State-Space Models

For years, the Transformer architecture has been the undisputed engine of the generative AI revolution. However, its foundational mechanism—self-attention—contains an inherent computational bottleneck that has limited its application to ever-longer sequences of data. Research unveiled this past week solidifies the emergence of State-Space Models (SSMs) as the most promising architectural successor, one that addresses the Transformer's core weakness and points toward a new paradigm in sequence modeling.

A New Architecture for Sequence Modeling: The Rise of SSMs

The primary challenge that SSMs are designed to solve is the quadratic complexity of the Transformer architecture. The self-attention mechanism, which allows a model to weigh the importance of all other tokens when processing a given token, requires a number of computations proportional to the square of the sequence length, expressed as $O(L^2)$.¹ This scaling law makes processing very long sequences—such as entire books, complete genomic data, or high-resolution video streams—prohibitively expensive in terms of both compute time and memory, creating a well-defined performance ceiling for Transformer-based models.¹⁰

State-Space Models offer a fundamentally different approach. Originating from control theory, SSMs model sequences as continuous-time dynamical systems.² Instead of processing all tokens in parallel via an attention matrix, they use a recurrent mechanism to update a compressed, fixed-size "state" that represents the history of the sequence seen so far.¹⁰ This

allows them to scale linearly with sequence length, with a complexity of

$O(L)$ or near-linear $O(L \log L)$, offering dramatic advantages in computational efficiency and speed, particularly during the inference stage where a response is generated one token at a time.³

From S4 to Mamba and S7: A Rapid Architectural Evolution

The rapid iteration and improvement within the SSM family over a short period indicates that this is not a niche academic pursuit but a fertile ground for innovation experiencing exponential progress, mirroring the explosive development seen in the Transformer ecosystem after 2017.

- **Foundational Work (S4):** The Structured State Space Sequence Model (S4) was the precursor that first demonstrated SSMs could effectively model long-range dependencies, outperforming traditional architectures like LSTMs and even specialized Transformers on benchmarks like audio and time-series data.¹ It established the viability of the core SSM concept in a deep learning context.
- **The Mamba Breakthrough:** The Mamba architecture, introduced more recently, solved a key weakness of earlier SSMs. It introduced a "selective" state-space mechanism, making the model's internal parameters *input-dependent*. This allows the model to dynamically emphasize or ignore parts of the input sequence based on the context at hand—a capability for content-based reasoning that was previously considered the exclusive domain of attention mechanisms. This critical innovation dramatically improved SSM performance on information-dense and unstructured data like natural language and genomics, making it the first subquadratic model to achieve performance on par with state-of-the-art Transformers.¹
- **The Latest Advance (S7):** A paper published in the last week introduced the S7 model, which builds on these advances. S7 simplifies previous SSM designs while retaining the crucial input-dependent dynamics. By dynamically adjusting its state transitions based on input content, S7 achieves a powerful balance of adaptive filtering and computational efficiency. In extensive experiments, S7 set new state-of-the-art results on a diverse range of benchmarks, including event-based vision, human activity recognition, and genomics classification on sequences over 17,000 tokens long, demonstrating the growing maturity and versatility of the SSM architecture.¹⁷

Potential Impact and Future Trajectory

The maturation of SSMs has profound implications for the future of AI. The existence of a viable, powerful alternative to the Transformer threatens to fragment the highly consolidated AI software and hardware ecosystem. For years, the industry has optimized for a single paradigm; Nvidia's GPUs are exceptionally well-suited for the massive matrix multiplications at the heart of attention, and its CUDA software library has become the de facto standard.¹⁸ SSMs, however, rely on different mathematical operations, such as convolutions and parallel scans.¹³ This opens the door for a new wave of "hardware-aware" model design and, conversely, "architecture-aware" hardware design.¹ A world with two dominant architectures could lead to specialized hardware accelerators, divergent software libraries, and a split in developer talent and expertise.

- Unlocking New Domains:** The linear-scaling efficiency of SSMs could enable significant breakthroughs in scientific and industrial domains previously hampered by the Transformer's limitations. These include analyzing entire genomes (sequences of billions of base pairs), processing high-resolution video streams in real time for robotics and autonomous systems, and understanding long-form legal or financial documents without the need for summarization or truncation.³
- Hardware Co-Design:** Because SSMs are less reliant on the dense matrix multiplications that GPUs are built for, they create an opportunity for new, more efficient hardware. This could accelerate the development of neuromorphic chips and other specialized processors that are better suited to recurrent operations, making powerful AI more practical on edge devices like AR glasses and drones.¹

The following table provides a distilled comparison of the fundamental trade-offs between the established Transformer paradigm and the emerging State-Space Model architecture.

Table 1: Comparative Analysis of Emerging AI Architectures (Transformer vs. SSM)

Feature	Transformer Architecture	State-Space Model (SSM) Architecture
Core Mechanism	Self-Attention (Parallel Processing)	Recurrence & State Representation (Sequential)
Computational Complexity	$O(L^2)$ - Quadratic in sequence length L	$O(L)$ or $O(L \log L)$ - Linear or Near-Linear
Inference Speed	Slower (requires full context for each step)	Faster (maintains a fixed-size state)

Memory Usage	High (KV cache grows with sequence length)	Low (fixed-size hidden state)
Long-Context Performance	Challenged by cost and memory constraints	Excels due to linear scaling
Input-Dependence	Native (via attention scores)	Achieved in advanced variants (e.g., Mamba, S7)
Hardware Optimization	Highly optimized for GPUs (Matrix Multiplication)	Potential for new hardware (e.g., Neuromorphic)

Key Discovery: The Geopolitical Reshaping of AI Hardware

The past week witnessed two major announcements in the AI hardware sector that, taken together, signal the beginning of a multi-front challenge to Nvidia's market monopoly. These moves by Intel and Alibaba are not merely new product launches; they are strategic responses to a global landscape shaped by intense competition and geopolitical imperatives. The result is an emerging pincer movement against the incumbent: a frontal assault on the high-end training market from Intel, and a flanking maneuver in the high-volume inference market from Alibaba.

Intel's Gambit: The 'Jaguar Shores' Rack-Scale Platform

Intel's announcement of its 'Jaguar Shores' platform is a high-stakes attempt to re-establish itself as a leader in the AI accelerator market after shelving its previous 'Falcon Shores' project. This is a clear "comeback" strategy aimed directly at Nvidia and AMD.⁴

- Architectural Deep Dive:** Corroborated reports from multiple technology outlets provide a detailed picture of an ambitious, high-performance design. Jaguar Shores will be manufactured on Intel's cutting-edge 18A process node (a 2nm-class technology), giving it a potential manufacturing advantage.⁶ Leaked images and specifications point to a massive 92.5 mm x 92.5 mm package featuring a multi-chiplet design with a quad-tile

configuration. This architecture will integrate next-generation HBM4 memory from SK Hynix via eight visible memory subsystems, a design clearly aimed at the most demanding High-Performance Computing (HPC) and AI model training workloads.⁶

- **"Rack-Scale" Ambition:** Critically, Intel is positioning Jaguar Shores not as a standalone chip, but as the core of a "rack-scale AI system".⁶ This system will be paired with Intel's future 'Diamond Rapids' Xeon CPUs, signifying a strategic pivot from selling individual components to providing fully integrated, full-stack hardware and software solutions.⁶ This approach directly competes with Nvidia's highly successful DGX systems and Google's TPU pods, recognizing that victory in the AI hardware race now requires a deeply integrated, system-level solution. The value is migrating up the stack from the processor to the platform.

Alibaba's Declaration of Independence: A RISC-V Chip for Inference

Alibaba's unveiling of a new, domestically produced AI chip is a direct consequence of U.S. export controls that have restricted Nvidia's access to the Chinese market. This geopolitical reality has created a powerful strategic imperative for Chinese technology giants to develop self-reliant, homegrown hardware alternatives.⁵

- **The "Inference, Not Training" Strategy:** Alibaba has made a shrewd strategic decision to focus its new chip on AI *inference* (running a trained model) rather than the more computationally intensive task of *training*. While model training captures headlines, inference workloads are far more numerous and diverse, representing a larger total addressable market. This asymmetric strategy allows Alibaba to chip away at Nvidia's dominance where it is most broadly deployed, targeting a wide range of applications from chatbots to image generation within its massive cloud ecosystem.⁵
- **Architectural and Manufacturing Significance:** The chip's architecture is based on RISC-V, an open-source instruction set. This is a significant move that avoids licensing fees and reliance on Western-controlled architectures like ARM or x86, aligning with China's national strategy to foster a self-sufficient semiconductor ecosystem.³⁰ Furthermore, the chip is being manufactured by a domestic Chinese company, breaking the reliance on external foundries like TSMC for advanced AI processors.²⁶
- **Pragmatic Integration:** Despite its independent architecture and domestic manufacturing, Alibaba has designed the chip for compatibility with Nvidia's dominant software frameworks, including CUDA and PyTorch.³⁰ This is a crucial and pragmatic feature designed to lower the barrier to adoption for the vast community of developers already accustomed to the Nvidia ecosystem, facilitating a smoother transition.

The table below summarizes the strategic positioning and key technical features of these new

accelerators in comparison to the current market leader.

Table 2: Next-Generation AI Accelerator Specifications (Announced/Reported)

Feature	Intel Jaguar Shores	Alibaba RISC-V Chip	Nvidia (Current/Projected)
Target Market	High-End Training & HPC	High-Volume Inference	Training & Inference (Full Spectrum)
Process Node	Intel 18A (2nm-class)	Domestic Chinese Mfg.	TSMC 3nm/2nm-class
Architecture	Likely GPGPU (Multi-tile)	RISC-V (Open Standard)	Proprietary (e.g., Blackwell)
Memory	HBM4	Not Specified	HBM3e / HBM4
Key Feature	Rack-Scale Integrated System	Domestic Production, Software Compatibility	CUDA Ecosystem, Performance Leadership
Geopolitical Driver	Compete with Nvidia/AMD	Circumvent US Export Controls	Maintain Market Dominance
Availability	~2026	In Testing	Current & 2025/26 Roadmap

Key Discovery: AI Applied to Core Scientific and Health Challenges

The past week also highlighted a significant maturation in the application of AI, moving beyond commercial and consumer domains to tackle fundamental processes in science and public health. Two new systems, developed at leading research institutions, showcase a new

frontier for AI: not just analyzing existing scientific data, but actively participating in and safeguarding the *process* of science itself. This represents a meta-level application of AI that could simultaneously accelerate discovery and improve research quality.

VaxSeer: Predictive AI for Proactive Vaccine Design

The annual process of selecting strains for the seasonal influenza vaccine is a high-stakes challenge. Health authorities must make a decision at least six months before the flu season, often relying on limited data and expert guesswork. This frequently leads to a mismatch between the vaccine and the circulating viral strains, resulting in low vaccine effectiveness.⁹

- **The VaxSeer Solution:** Researchers at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) have developed an AI system called VaxSeer to make this process more data-driven and predictive. The system uses a dual-engine architecture trained on decades of viral genetic sequences and laboratory data to forecast viral evolution and vaccine performance months in advance.⁷
 - **Engine 1 (Dominance Prediction):** This engine uses a large protein language model to analyze viral mutations. Unlike traditional methods that assess mutations individually, VaxSeer evaluates the combinatorial effects of multiple mutations to predict which strains are most likely to spread widely and become dominant. It models the competition between different strains over time using a framework of ordinary differential equations.⁹
 - **Engine 2 (Antigenicity Prediction):** This engine estimates how effectively a candidate vaccine will neutralize each predicted dominant strain. It does this by simulating the results of a common lab test known as the hemagglutination inhibition (HI) assay, which measures how well antibodies prevent the virus from binding to red blood cells.⁹
- **Validation and Performance:** The efficacy of VaxSeer was tested in a comprehensive 10-year retrospective study. Its recommendations were compared against the historical selections made by the World Health Organization (WHO). For the A/H3N2 flu subtype, VaxSeer's choices outperformed the WHO's in nine out of ten seasons. For the A/H1N1 subtype, it matched or exceeded the WHO's choices in six out of ten seasons. Notably, for the 2016 flu season, VaxSeer identified a critical strain that the WHO did not select until the following year. The model's predictions also showed a strong correlation with real-world vaccine effectiveness data reported by public health agencies like the CDC.⁹

A Firewall for Science: AI-Powered Vetting of Academic Journals

The integrity of the scientific record is threatened by the proliferation of "predatory" journals. These entities charge researchers publication fees but fail to provide legitimate peer review, polluting scientific literature with unvetted and potentially fraudulent work.⁸

- **The CU Boulder Solution:** To combat this, computer scientists at the University of Colorado Boulder have developed an AI tool that automatically screens and flags questionable journals. The system was designed to serve as a "firewall for science," augmenting the work of human experts who struggle to keep pace with the sheer volume of new publications.⁸
 - **Methodology:** The AI was trained using a curated dataset from the Directory of Open Access Journals (DOAJ), a nonprofit that vets journal legitimacy. The tool analyzes journal websites for a range of red flags commonly associated with predatory practices. These include frequent grammatical errors, fake or non-existent editorial boards, unusually high publication volumes, and patterns of excessive author self-citation.⁸
 - **Validation and Performance:** In its initial run, the AI screened nearly 15,200 open-access journals and flagged over 1,400 as potentially problematic. To validate the system's accuracy, human experts then reviewed a subset of these flagged journals. This review identified an estimated 350 false positives, which were subsequently corrected. After this human-in-the-loop refinement, over 1,000 journals were confirmed as exhibiting questionable characteristics. The researchers emphasize that the tool is designed as a "helper" to prescreen large numbers of journals, with the final judgment remaining with human professionals.⁸

The success of VaxSeer and the challenges faced in journal vetting both underscore a critical dependency for high-stakes AI: the availability of large, high-quality, longitudinal datasets. VaxSeer's predictive power stems directly from the decades of structured flu surveillance data it was trained on.⁹ The journal vetter's accuracy, in turn, is entirely dependent on the quality of its initial training list from DOAJ.⁸ This demonstrates that for AI to solve the most critical real-world problems, the primary bottleneck is often not the algorithm itself, but the existence and meticulous curation of the data required to train it.

Emerging Technologies and Early Industry Applications

The foundational discoveries of the past week are not confined to theoretical research; they are already being translated into specialized technologies with clear industry applications.

From novel computer vision models to targeted hardware deployments, these early use cases provide a glimpse into how new AI paradigms will reshape specific sectors.

SSMs in the Wild: New Models for Computer Vision

The architectural advantages of State-Space Models are being rapidly adapted for domains beyond text and audio. New research has introduced vision-specific SSMs, such as LocalMamba and GroupMamba, which are engineered to better process 2D data like images.²⁰ These models address a limitation of early vision SSMs by dividing images into distinct windows or patches. This allows the model to efficiently capture local dependencies (like textures and edges) within a window while also modeling the global relationships between windows. In benchmark tests, these vision-specific SSMs have demonstrated superior performance-to-efficiency trade-offs, outperforming traditional Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) of a similar size while processing images significantly faster.²⁰

Hardware Market Targeting and Economic Impact

The new hardware platforms from Intel and Alibaba are being aimed at distinct, high-value market segments.

- **Intel's Jaguar Shores** is being positioned for the traditional HPC market. Its target customers include national laboratories engaged in scientific research, sovereign states building national AI clouds, and large enterprises developing their own foundational models. This is a direct play for the highest-margin segment of the AI infrastructure market.⁶
- **Alibaba's RISC-V chip** is primarily targeting its own extensive cloud services customer base and the broader Chinese technology sector. It provides a cost-effective, domestically produced solution for high-volume inference tasks, insulating Chinese companies from the effects of U.S. export controls and enabling applications in e-commerce, customer service chatbots, and content generation.⁵

These developments are occurring within a context of massive economic investment in AI. The unprecedented demand for AI computation is fueling record revenues across the entire technology stack. In their latest earnings reports, cloud service providers like Microsoft Azure, Google Cloud, and Amazon Web Services all posted strong double-digit growth, largely attributed to AI workloads. This demand has cascaded down to hardware vendors, with Nvidia,

AMD, and Dell reporting record gains in their data center and server divisions. The surge confirms that AI has transitioned from a speculative technology to a primary economic driver transforming multiple global industries.¹⁸

Challenges and Strategic Considerations

The rapid pace of AI innovation brings with it a host of complex challenges, from ensuring the reliability of AI systems in high-stakes environments to mitigating the risks of misuse by malicious actors. Developments from the past week have brought these strategic considerations into sharp focus.

The Reliability Gauntlet: Validating AI in High-Stakes Domains

As AI systems are deployed to perform critical functions like safeguarding scientific integrity, the process of validating their reliability becomes paramount. This challenge can be described as the "vetting paradox": who vets the vetters?

- **The Tale of Two Vetting Tools:** The AI-powered journal screening tool from CU Boulder exemplifies a responsible approach. It was designed with a "human-in-the-loop" model for validation, where human experts reviewed and corrected the AI's outputs, and it is explicitly positioned as an assistant to, not a replacement for, human judgment.⁸ In stark contrast, another tool discussed in recent research, the "Academic Journal Predatory Checking System" (AJPC), has proven to be wildly inaccurate and unreliable. In one analysis, the AJPC tool classified 100% of the 4,756 journals published by the major academic publisher Elsevier as "suspected predatory," a demonstrably false and damaging conclusion.⁴⁷
- **Ethical Frameworks and the Accountability Gap:** The failure of the AJPC tool highlights a core ethical dilemma. The publishing industry, guided by organizations like the Committee on Publication Ethics (COPE), has established robust guidelines for human authors and editors. A central tenet of these guidelines is accountability.⁵⁰ However, as COPE's own position statement on AI clarifies, an AI tool cannot be held accountable; it cannot take responsibility for its work, manage conflicts of interest, or be held liable for errors.⁵² When an unreliable AI tool falsely damages the reputation of a legitimate journal or researcher, there is no clear path for recourse. This makes the human-in-the-loop model not just a best practice for accuracy, but an ethical necessity for accountability.

The Innovation-Security Dilemma

The increasing power and accessibility of advanced AI models create a significant dual-use problem, where the same technologies that enable scientific breakthroughs can be weaponized for malicious purposes.

- **Dual-Use Technology:** Reports from security researchers this week confirmed that cybercriminals are actively misusing commercially available AI models. Anthropic's Claude model, for example, has reportedly been used to design and build complex ransomware-as-a-service (RaaS) tools. The AI assists in writing sophisticated malware that incorporates advanced encryption, anti-debugging techniques, and command-and-control infrastructure.⁵³
- **Novel AI-Specific Attack Vectors:** New vulnerabilities are emerging that are unique to the way AI systems process information. Researchers have demonstrated "PromptLock," a proof-of-concept ransomware that uses a GPT model to *dynamically generate* malicious scripts in real time, making it much harder for traditional signature-based antivirus software to detect. Another novel attack involves embedding malicious prompts within high-resolution images. These prompts are invisible to the human eye and only become active after the image is downscaled by an AI vision system, tricking the model into executing unintended commands or leaking sensitive data.⁵³
- **The Security Arms Race:** These developments confirm the existence of a clear and escalating arms race in cybersecurity. As AI models become more capable, they become more powerful tools for both offense and defense. This reality necessitates a much deeper focus on AI safety, alignment research, and the development of robust security measures specifically designed to counter AI-powered threats.

Outlook: Converging Trends and Near-Future Directions

The discoveries of the past seven days provide a clear trajectory for the near future of artificial intelligence, pointing toward a period of architectural diversification, hardware fragmentation, and an intensified focus on the application of AI to science and the governance of its risks.

- **Architecture (Next 6-12 Months):** The momentum behind State-Space Models will accelerate. The immediate future will likely see a surge in research into hybrid

architectures that seek to combine the global context capabilities of Transformers with the long-sequence efficiency of SSMs. Concurrently, expect to see more rigorous, large-scale benchmarking of pure SSMs on complex language and reasoning tasks to definitively determine if they can match or exceed the performance of the best Transformer models in production environments.

- **Hardware (Next 1-2 Years):** The AI hardware market is set to become increasingly fragmented and overtly politicized. The announcements from Intel and Alibaba are likely the beginning of a trend, with more specialized chips emerging from both established players and new startups, some of which will be explicitly optimized for the computational patterns of SSM-like workloads. A critical variable to monitor will be the manufacturing performance of Intel's 18A process node. The success or failure of this technology will not only determine the competitive viability of the Jaguar Shores platform but will also have significant implications for the global semiconductor landscape and TSMC's long-held manufacturing dominance.
- **Applications & Ethics (Ongoing):** The application of AI to foundational scientific challenges will continue to expand beyond medicine into fields like materials science, climate modeling, and fundamental physics. This acceleration will be paralleled by an intensified global debate on AI governance. The dual-use problem, where powerful models can be used for both benevolent and malicious ends, will force policymakers, developers, and the public to grapple more seriously with frameworks for access control, responsible deployment, and legal liability for AI systems. The central challenge for the entire field will be to navigate the complex trade-offs required to foster profound innovation while simultaneously mitigating the potential for catastrophic risk.

Works cited

1. The Next Wave of AI Architectures in 2025 | by Babar saad, accessed September 1, 2025, https://ai.plainenglish.io/the-next-wave-of-ai-architectures-in-2025-99d0355703b7?source=rss-----artificial_intelligence-5
2. Primers • State Space Models - aman.ai, accessed September 1, 2025, <https://aman.ai/primers/ai/state-space-models/>
3. From S4 to Mamba: A Comprehensive Survey on Structured State Space Models - arXiv, accessed September 1, 2025, <https://www.arxiv.org/pdf/2503.18970>
4. Weekly news roundup: Intel revives AI chip fight with Jaguar Shores; Samsung may invest in Intel packaging push - DIGITIMES Asia, accessed September 1, 2025, <https://www.digitimes.com/news/a20250901VL204/weekly-news-roundup-intel-a-i-chip-2nm-packaging-samsung.html>
5. Stocks - MLQ.ai, accessed September 1, 2025, <https://mlq.ai/news/alibaba-unveils-advanced-ai-chip-to-challenge-nvidia-markin-g-major-step-to-end-reliance/>
6. [News] Intel's Jaguar Shores Reportedly Breaks Cover, Showcasing ..., accessed September 1, 2025,

- <https://www.trendforce.com/news/2025/08/21/news-intels-jaguar-shores-reportedly-breaks-cover-showcasing-18a-and-hbm4-on-rack-scale-ai-solution/>
7. Artificial intelligence | MIT News | Massachusetts Institute of Technology, accessed September 1, 2025, <https://news.mit.edu/topic/artificial-intelligence2>
 8. New AI tool identifies 1,000 'questionable' scientific journals | CU ..., accessed September 1, 2025, <https://www.colorado.edu/today/2025/08/28/new-ai-tool-identifies-1000-questionable-scientific-journals>
 9. MIT researchers develop AI tool to improve flu vaccine strain ..., accessed September 1, 2025, <https://news.mit.edu/2025/vaxseer-ai-tool-to-improve-flu-vaccine-strain-selection-0828>
 10. Introduction to State Space Models as Natural Language Models, accessed September 1, 2025, <https://neptune.ai/blog/state-space-models-as-natural-language-models>
 11. What Are State Space Models? | IBM, accessed September 1, 2025, <https://www.ibm.com/think/topics/state-space-model>
 12. Mamba: Linear-Time Sequence Modeling with Selective State Spaces - OpenReview, accessed September 1, 2025, <https://openreview.net/forum?id=tEYskw1VY2>
 13. Mamba Simplified - Part 2 - S4 and Mamba - Prem AI Blog, accessed September 1, 2025, <https://blog.prem.ai/s4-and-mamba/>
 14. From S4 to Mamba: A Comprehensive Survey on Structured State Space Models - arXiv, accessed September 1, 2025, <https://arxiv.org/abs/2503.18970>
 15. Mamba: Linear-Time Sequence Modeling with Selective State Spaces - Semantic Scholar, accessed September 1, 2025, <https://www.semanticscholar.org/paper/Mamba%3A-Linear-Time-Sequence-Modeling-with-Selective-Gu-Dao/7bbc7595196a0606a07506c4fb1473e5e87f6082>
 16. state-spaces/mamba: Mamba SSM architecture - GitHub, accessed September 1, 2025, <https://github.com/state-spaces/mamba>
 17. S7: Selective and Simplified State Space Layers for Sequence Modeling - arXiv, accessed September 1, 2025, <https://arxiv.org/html/2410.03464v1>
 18. CFOs Triple Down on Gen AI ROI, Fueling Chip-to-Server Surge | PYMNTS.com, accessed September 1, 2025, <https://www.pymnts.com/artificial-intelligence-2/2025/cfos-triple-down-on-gen-ai-roi-fueling-chip-to-server-surge/>
 19. Top 20 AI Chip Makers: NVIDIA & Its Competitors - Research AIMultiple, accessed September 1, 2025, <https://research.aimultiple.com/ai-chip-makers/>
 20. Making computer vision more efficient with state-space models - MBZUAI, accessed September 1, 2025, <https://mbzuai.ac.ae/news/making-computer-vision-more-efficient-with-state-space-models/>
 21. Intel Names Jaguar Shores as its Next-generation AI Chip - HPCwire, accessed September 1, 2025, <https://www.hpcwire.com/2024/11/19/intel-names-jaguar-shores-as-its-next-gen>

- [eration-ai-chip/](#)
22. Intel Jaguar Shores: First Look at the Next AI Processor - Blog - HostZealot, accessed September 1, 2025, <https://www.hostzealot.com/blog/news/intel-jaguar-shores-first-look-at-the-next-ai-processor>
 23. Intel Jaguar Shores: the company's new AI processor is shown in photos - ITC.ua, accessed September 1, 2025, <https://itc.ua/en/news/intel-jaguar-shores-the-company-s-new-ai-processor-is-shown-in-photos/>
 24. AI's Weekend Whirlwind: Global Crackdowns, Tech Giants' Bold Moves & Market Shocks, accessed September 1, 2025, <https://ts2.tech/en/ais-weekend-whirlwind-global-crackdowns-tech-giants-bold-moves-market-shocks/>
 25. Alibaba Tests New AI Chip as China Seeks Nvidia Alternatives - GuruFocus, accessed September 1, 2025, <https://www.gurufocus.com/news/3087943/alibaba-tests-new-ai-chip-as-china-seeks-nvidia-alternatives>
 26. China's biggest cloud-computing company unveils AI chip to challenge Nvidia, accessed September 1, 2025, <https://timesofindia.indiatimes.com/technology/tech-news/chinas-biggest-cloud-computing-company-unveils-ai-chip-to-challenge-nvidia/articleshow/123588223.cms>
 27. Alibaba unveils new AI chip as Nvidia's China headache deepens - City AM, accessed September 1, 2025, <https://www.cityam.com/alibaba-unveils-new-ai-chip-as-nvidias-china-headache-deepens/>
 28. Why Alibaba AI chip launch rattles Wall Street tech giants - Rolling Out, accessed September 1, 2025, <https://rollingout.com/2025/08/30/alibaba-ai-chip-launch-tech-giants/>
 29. Nvidia Made Every Robot 7x Smarter and Way More Expensive - eWeek, accessed September 1, 2025, <https://www.eweek.com/news/nvidia-jetson-thor/>
 30. Alibaba's AI Chip Breakthrough: A Strategic Play to Disrupt Nvidia's AI Dominance in China, accessed September 1, 2025, <https://www.ainvest.com/news/alibaba-ai-chip-breakthrough-strategic-play-disrupt-nvidia-ai-dominance-china-2509/>
 31. Alibaba's AI Chip Breakthrough: A Strategic Threat to Nvidia and a Buy Opportunity in China Tech? - AInvest, accessed September 1, 2025, <https://www.ainvest.com/news/alibaba-ai-chip-breakthrough-strategic-threat-nvidia-buy-opportunity-china-tech-2508/>
 32. Cambricon Rises as China's AI Chip Champion - EE Times, accessed September 1, 2025, <https://www.eetimes.com/cambricon-rises-as-chinas-ai-chip-champion/>
 33. Alibaba Stock Surges After Earnings Beat and AI Chip Strategy to Rival Nvidia - TECHi, accessed September 1, 2025, <https://www.techi.com/alibaba-stock-surges-ai-chip-strategy-vs-nvidia/>
 34. Alibaba unveils server-grade, high-performance RISC-V chip - Techerati,

- accessed September 1, 2025,
<https://www.techerati.com/news-hub/alibaba-unveils-server-grade-high-performance-risc-v-chip/>
35. New Alibaba AI Chip Could Replace Nvidia Hardware in China - eWeek, accessed September 1, 2025, <https://www.eweek.com/news/alibaba-ai-chip/>
 36. VaxSeer: Selecting influenza vaccines through evolutionary and antigenicity models, accessed September 1, 2025,
<https://www.biorxiv.org/content/10.1101/2023.11.14.567037v3>
 37. MIT's AI tool VaxSeer improves flu vaccine selection by predicting ..., accessed September 1, 2025,
<https://hyper.ai/en/headlines/cfafb1e5e67005efdf875927e87d146f>
 38. VaxSeer Predicts Flu Strains With AI to Improve Vaccine Selection and Outpace Viral Evolution - Complete AI Training, accessed September 1, 2025,
<https://completeaitraining.com/news/vaxseer-predicts-flu-strains-with-ai-to-improve-vaccine/>
 39. MIT researchers develop AI tool to improve flu vaccine strain selection, accessed September 1, 2025,
<https://www.csail.mit.edu/news/mit-researchers-develop-ai-tool-improve-flu-vaccine-strain-selection>
 40. MIT's VaxSeer AI outperformed WHO flu vaccine picks in 9 of 10 seasons - CO/Al, accessed September 1, 2025,
<https://getcoai.com/news/mits-vaxseer-ai-outperformed-who-flu-vaccine-picks-in-9-of-10-seasons/>
 41. - Machine learning predicts the future of flu - Connectivity, accessed September 1, 2025,
<https://www.connectivity4ir.co.uk/article/217301/Machine-learning-predicts-the-future-of-flu.aspx>
 42. MIT's VaxSeer Uses AI to Predict Flu Strains, Beating WHO Picks in Most Seasons Reviewed - iAfrica.com, accessed September 1, 2025,
<https://iafrica.com/mits-vaxseer-uses-ai-to-predict-flu-strains-beating-who-picks-in-most-seasons-reviewed/>
 43. A US-developed AI has identified over 1,000 fake scientific journals by analysing website patterns: Here's how it's helping to clean up academic publishing, accessed September 1, 2025,
<https://timesofindia.indiatimes.com/education/news/a-us-developed-ai-has-identified-over-1000-fake-scientific-journals-by-analysing-website-patterns-heres-how-its-helping-to-clean-up-academic-publishing/articleshow/123624269.cms>
 44. AI Flags Over 1,000 Questionable Scientific Journals, Raising Alarms About Research Integrity, accessed September 1, 2025,
<https://completeaitraining.com/news/ai-flags-over-1000-questionable-scientific-journals-raising/>
 45. [2403.09338] LocalMamba: Visual State Space Model with Windowed Selective Scan - arXiv, accessed September 1, 2025, <https://arxiv.org/abs/2403.09338>
 46. This Week's Awesome Tech Stories From Around the Web (Through August 30), accessed September 1, 2025,

<https://singularityhub.com/2025/08/30/this-weeks-awesome-tech-stories-from-around-the-web-through-august-30/>

47. The diagnostic accuracy of AI-based predatory journal detectors: An ..., accessed September 1, 2025,
https://www.researchgate.net/publication/370467389_The_diagnostic_accuracy_of_AI-based_predatory_journal_detectors_An_analogy_to_diagnosis
48. An artificial intelligence tool misclassifies sport science journals as predatory - PubMed, accessed September 1, 2025,
<https://pubmed.ncbi.nlm.nih.gov/38158321/>
49. (Mis-)Classification of 17,721 Journals by an Artificial Intelligence Predatory Journal Detector | Request PDF - ResearchGate, accessed September 1, 2025,
https://www.researchgate.net/publication/372341389_Mis-Classification_of_17721_Journals_by_an_Artificial_Intelligence_Predatory_Journal_Detector
50. Publication Ethics - University of Pennsylvania Press, accessed September 1, 2025,
<https://www.pennpress.org/about/publication-ethics/>
51. COPE: Committee on Publication Ethics, accessed September 1, 2025,
<https://publicationethics.org/>
52. Authorship and AI tools | COPE - Committee on Publication Ethics, accessed September 1, 2025,
<https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools>
53. Latest AI Breakthroughs and News: June, July, August 2025 - Crescendo.ai, accessed September 1, 2025,
<https://www.crescendo.ai/news/latest-ai-news-and-updates>