








NEWMIND AI JOURNAL WEEKLY CHRONICLES




1.4.2025 - 7.4.2025



- The first week of April 2025 (April 1st-7th) showcased significant model releases and strategic industry maneuvers in AI, as detailed in the NewMind AI Journal Chronicle.
- Major players like **Alibaba** prepared the next-generation Qwen model, **Meta** unveiled the diverse Llama 4 family, and **Cohere** launched the enterprise-focused Command R+.
- **Intel's** foundry ambitions engaged **TSMC** and **Arm**, reflecting the intense global scramble for semiconductor supremacy.
- AI's influence reached further into society, with **Adobe** enhancing video editing tools, the **US Social Security Administration** exploring AI applications, and **Chinese state media** using AI for political commentary.
- As AI matures, the focus is shifting toward specialized, enterprise-ready models designed to address specific business needs.
- Models like **Cohere's Command R+** excel in efficient information retrieval from company data (RAG) and seamless integration into existing workflows.
- Competitive advantage will increasingly stem from strategically adopting AI tailored to operational challenges and opportunities, rather than relying on general capabilities.
- The ongoing development of hardware supply chains remains critical to sustaining AI progress.



Models					
#	Highlights	Summary	Author	Source	Date
1.1	Alibaba prepares for flagship AI model release as soon as April	Alibaba is reportedly preparing to launch an upgraded version of its flagship AI model, Qwen, later in April 2025, according to Bloomberg. This move highlights Alibaba's ongoing efforts to stay competitive in China's rapidly evolving AI sector. The upcoming Qwen model is expected to offer significant improvements in performance and capabilities, reflecting a broader push by Chinese tech giants to rival global AI leaders like OpenAI and Google. Alibaba had previously released earlier versions of Qwen under its DAMO Academy as part of its research and innovation strategy.	By Reuters		April 1, 2025
1.2	Command A: An Enterprise-Ready Large Language Model	Command A is a state-of-the-art large language model purpose-built for enterprise applications, offering multilingual capabilities across 23 global business languages and excelling in Retrieval-Augmented Generation (RAG), tool use, and complex workflows. With a hybrid architecture optimized for efficiency and performance, it achieves best-in-class results across academic, reasoning, code, and multilingual benchmarks while requiring minimal computational resources. Innovations include self-refinement algorithms, model merging techniques, and decentralized training. Released under a non-commercial license, Command A sets a new standard for enterprise-ready AI,	By Cohere		April 1, 2025


 Models					
#	Highlights	Summary	Author	Source	Date
		balancing versatility, efficiency, and privacy-preserving deployment.			
1.3	IBM and Tokyo Electron Renew Partnership to Advance AI Chip Manufacturing	IBM and Tokyo Electron have renewed their collaboration to accelerate research in advanced semiconductor manufacturing, focusing on scaling technologies critical for AI workloads. The partnership will enhance development in areas like extreme ultraviolet (EUV) lithography and advanced packaging—key to producing high-performance, energy-efficient AI chips. This move supports IBM’s ongoing AI hardware roadmap and Japan’s ambition to strengthen domestic chip production. The renewed agreement continues a decade-long relationship and reflects growing geopolitical and industrial efforts to secure next-generation chip capabilities amid rising global AI demand.	By Bethany McCarthy		April 2, 2025
1.4	Open-Qwen2VL: Compute-Efficient Pre-Training of Fully-Open Multimodal LLMs on Academic Resources	This paper introduces Open-Qwen2VL, a 2-billion-parameter Multimodal Large Language Model (MLLM) designed for efficient pre-training on limited academic resources. Utilizing approximately 5 billion high-quality image-text pairs—merely 0.36% of the 1.4 trillion tokens used in comparable models—Open-Qwen2VL achieves state-of-the-art performance on benchmarks like MMBench, SEEDBench, MMStar, and	By Weizhi Wang et al.		April 2, 2025




Models					
#	Highlights	Summary	Author	Source	Date
		MathVista. Key innovations include dynamic image resolution, multimodal sequence packing, and advanced data filtering techniques. The project is fully open-source, providing training code, data filtering methods, sequence packing scripts, and model checkpoints, promoting transparency and reproducibility in MLLM research.			
1.5	Efficient LLaMA-3.2-Vision by Trimming Cross-attended Visual Features	The paper "Efficient LLaMA-3.2-Vision by Trimming Cross-attended Visual Features" introduces a method to enhance the efficiency of large vision-language models (LVLMs) using cross-attention architectures, such as LLaMA-3.2-Vision. The authors identify that the key-value (KV) cache size for image tokens in self-attention layers is much larger than that for text tokens in cross-attention layers, creating a computational bottleneck. To address this, they propose "Trimmed Llama," a training-free approach that exploits sparsity in cross-attention maps to prune redundant visual features during inference. This method reduces KV cache demands, decreasing inference latency and memory usage while maintaining performance on benchmark tasks.	By Jewon Lee et al.		April 1, 2025
1.6	Nomic Embed Multimodal: Open Source Multimodal Embedding Models for	Nomic AI has introduced Nomic Embed Multimodal , an open-source embedding model that unifies text, images, PDFs, and charts into a shared vector space. Designed for real-world applications, the model excels at understanding documents	By Nomic Team		April 2, 2025



 Models					
#	Highlights	Summary	Author	Source	Date
	Text, Images, PDFs, and Charts	combining visuals and text, such as research papers, annotated screenshots, and reports. It enables powerful search and classification across complex data types. Trained on over 100 million diverse documents, it outperforms previous models on benchmarks like Vidore-v2, achieving a 62.7 NDCG@5 score. With native support in the Nomic Atlas platform, it offers scalable, multimodal understanding for enterprise and research use cases.			
1.7	Midjourney releases V7, its first new AI image model in nearly a year	Midjourney has recently released Version 7 (V7), its first new AI image generation model in nearly a year. This latest version introduces significant enhancements in coherence, realism, and prompt control, offering users finer-grained editing capabilities and higher fidelity images. V7 allows for more complex compositions through direct text prompting and includes advanced editing tools like inpainting and outpainting. This release reflects Midjourney's commitment to providing greater creative flexibility and positions the platform as a strong competitor in the evolving generative image landscape.	By Kyle Wiggers		April 4, 2025
1.8	Scaling Analysis of Interleaved Speech-Text Language Models	Existing Speech Language Model (SLM) scaling analysis suggests that SLMs require significantly more compute and data compared to text-based models, raising concerns about their feasibility. However, modern SLMs often leverage pre-trained TextLMs using speech-text interleaving for knowledge transfer. This paper demonstrates that interleaved SLMs scale more	By Gallil Maimon, Michael Hassid, Amit Roth, Yossi Adi		April 3, 2025



Models					
#	Highlights	Summary	Author	Source	Date
		efficiently than textless-SLMs. Our scaling analysis, conducted across multiple models, shows that interleaved SLMs use compute more effectively. Additionally, we find that scaling dynamics differ notably from textless-SLMs, recommending more compute for model size rather than training tokens. Synthetic data and TextLM families play a key role in unlocking this efficiency.			
1.9	The Llama 4 herd: The beginning of a new era of natively multimodal AI innovation	Meta has unveiled two new AI models: Llama 4 Scout and Llama 4 Maverick. Llama 4 Scout is a compact model designed to operate on a single Nvidia H100 GPU, featuring a 10-million-token context window and outperforming competitors like Google's Gemma 3 and Mistral 3.1 across various benchmarks. Llama 4 Maverick is a larger model, comparable in performance to OpenAI's GPT-4o and DeepSeek-V3 in coding and reasoning tasks, while utilizing fewer active parameters. Additionally, Meta is developing Llama 4 Behemoth, boasting 288 billion active parameters, which is claimed to surpass models like GPT-4.5 and Claude Sonnet 3.7 on STEM benchmarks.	By Meta		April 5, 2025
1.10	Agentic Knowledgeable Self-awareness	Large Language Models (LLMs) excel in agentic planning tasks but often rely on a "flood irrigation" method, injecting gold trajectories and external feedback indiscriminately, ignoring human-like self-awareness in decision-making. To address this, we introduce KnowSelf, a data-centric approach that equips LLM agents with knowledgeable self-awareness. This allows agents to	By Shuofei Qiao, Zhisong Qiu, Baochang Ren, Xiaobin Wang, Xiangyuan Ru, Ningyu Zhang, Xiang		April 4, 2025




Models					
#	Highlights	Summary	Author	Source	Date
		autonomously regulate knowledge use by evaluating situations and strategically deciding when to use resources. We implement a two-stage training process with a heuristic system to mark special tokens for self-exploration. Experiments show that KnowSelf outperforms strong baselines while minimizing external knowledge usage.	Chen, Yong Jiang, Pengjun Xie, Fei Huang, Huajun Chen		
1.11	A Comprehensive Evaluation of Open Large Vision-Language Models	This paper evaluates 30 open-source Vision-Language Models (VLMs) across 10 diverse benchmarks including image captioning, visual question answering, referring expressions, and multimodal reasoning. The authors introduce OpenVLM, a unified framework and leaderboard for benchmarking, revealing significant performance gaps between open and closed VLMs. Notably, models like Kosmos-2, LLaVA-1.5, and X-VLM Plus show strong performance, but still lag behind proprietary models in zero-shot tasks. The study highlights challenges like hallucination and data quality issues, and provides a reproducible infrastructure to drive transparent VLM development.	By Xinyi Wang et al.		April 4, 2025
1.12	VISTA-OCR: Towards generative and interactive end to end OCR models	The paper introduces VISTA-OCR , a novel generative and interactive end-to-end OCR framework that enhances text recognition and correction. Unlike traditional OCR systems, VISTA-OCR leverages visual and linguistic context to improve accuracy, supporting user-in-the-loop corrections for	Laziz Hamdi et al.		April 4, 2025



Models					
#	Highlights	Summary	Author	Source	Date
		ambiguous or erroneous text. It integrates generative AI to predict and refine outputs dynamically, making it adaptable to diverse document types. Experiments show superior performance over conventional OCR models, particularly in noisy or complex layouts. The approach bridges the gap between automated recognition and human-involved refinement , offering a more robust and flexible OCR solution.			
1.13	ScholarCopilot: Training Large Language Models for Academic Writing with Accurate Citations	ScholarCopilot introduces a unified framework for enhancing large language models (LLMs) in academic writing by integrating dynamic citation retrieval within the generative process. Unlike traditional static retrieval-augmented generation (RAG) systems, ScholarCopilot dynamically generates retrieval tokens ([RET]) during text generation, enabling context-aware citation retrieval. Built on Qwen-2.5-7B and trained on 500K arXiv papers, it achieves a top-1 retrieval accuracy of 40.1%, outperforming baselines like E5-Mistral-7B-Instruct and BM25. Human studies show ScholarCopilot surpasses ChatGPT in citation quality (100% preference) and overall usefulness (70% preference), making it a pioneering tool for efficient, accurate, and coherent academic writing.	By Yubo Wang, Xueguang Ma et al.		April 3, 2025



AI Chips					
#	Highlights	Summary	Author	Source	Date
2.1	CoreWeave Surpasses IPO Price on Third Day of Trading	CoreWeave, a cloud computing company specializing in GPU-powered infrastructure for AI workloads, saw its shares rise above its initial public offering (IPO) price on the third day of trading. After pricing its IPO at \$16 per share, CoreWeave's stock climbed, reflecting growing investor confidence in the company's position within the AI-driven cloud infrastructure market. CoreWeave has been expanding rapidly, fueled by increased demand for AI model training and cloud services. The IPO's success marks a significant milestone as CoreWeave continues to capitalize on the booming AI and cloud industries.	By Reuters		April 2, 2025
2.2	Cerebras Systems, Ranovus win \$45 million US military deal to speed up chip connections	Cerebras Systems and Ranovus have secured a \$45 million DARPA contract to enhance chip communication. Cerebras, known for its large AI chips, aims to outperform traditional GPUs, while Ranovus specializes in optical data transfer for faster, energy-efficient communication. Their collaboration seeks to develop advanced computing systems for real-time battlefield simulations, promising a 150-fold speed increase while cutting power consumption from 30 watts to 3. The project will integrate new, undisclosed technologies to improve inter-chip connections, reinforcing U.S. military computing capabilities. Both companies are pushing innovation in AI and high-performance computing.	By Stephen Nellis		April 1, 2025
2.3	Intel's New CEO Outlines Recovery	Intel CEO Pat Gelsinger has reaffirmed the company's commitment to regaining leadership in semiconductor innovation by emphasizing	By Patrick Seitz		April 1, 2025

AI Chips					
#	Highlights	Summary	Author	Source	Date
	Strategy. Wall Street Remains Cautious	engineering excellence, customer collaboration, and the rollout of its Intel 18A process node. The 18A node, set for high-volume production later this year, is critical for Intel's AI chip ambitions amid fierce competition from Nvidia and AMD. Despite a recent stock dip and cautious investor sentiment, Intel aims to deliver next-gen chips with advanced performance and efficiency. Gelsinger framed the strategy as “the comeback of a lifetime” for the iconic chipmaker.			
2.4	Arm recently sought to acquire Alphawave for AI chip tech, sources say	Arm Holdings, owned by SoftBank, recently explored acquiring UK-based Alphawave to secure its advanced SerDes (serializer-deserializer) technology—crucial for AI chips that require high-speed data transmission. Alphawave, which was evaluating sale options, ultimately did not reach a deal with Arm. The move signaled Arm’s interest in expanding beyond licensing chip designs into producing its own AI processors. SerDes is vital in linking thousands of chips for AI workloads, and competitors like Broadcom and Nvidia already dominate this space. Alphawave’s shares rose 21% after the news, though geopolitical concerns related to its China joint venture remain a factor.	By Milana Vinn, Max A. Cherney and Amy-Jo Crowley		April 1, 2025
2.5	Qualcomm considers buying	Qualcomm is considering a takeover of UK-based semiconductor firm Alphawave, whose shares surged over 52% following the news. Alphawave specializes in SerDes technology—crucial for AI chips that demand fast, high-volume data processing—and supplies designs used in custom chips	By Reuters		April 1, 2025


AI Chips					
#	Highlights	Summary	Author	Source	Date
	UK semiconductor firm Alphawave	by Broadcom and Marvell. The potential deal highlights Qualcomm's interest in expanding its AI hardware capabilities. Qualcomm has until April 29 to make a formal offer under UK takeover rules. This follows earlier reports that Arm, owned by SoftBank, also explored acquiring Alphawave but walked away. Both companies see Alphawave's IP as strategically vital for AI chip development.			
2.6	Chinese Tech Firms Reportedly Order \$1.6 Billion in Nvidia AI Chips	Chinese companies have reportedly placed a \$1.6 billion order for Nvidia's new AI chips, according to The Information. The order includes specially modified H20 chips, designed to comply with U.S. export controls while still delivering high AI performance. Tech giants such as Alibaba and Tencent are among the buyers, seeking to bolster their AI capabilities amid tightening U.S.-China tech restrictions. Shipments are expected to begin in Q2 2025. This order highlights the continued demand for advanced AI hardware in China despite regulatory hurdles.	By Reuters		April 2, 2025
2.7	Parasail says its fleet of on-demand GPUs is larger than Oracle's entire cloud	The next-generation AI infrastructure startup Parasail has officially launched its GPU-focused platform, claiming that its on-demand GPU fleet is larger than Oracle's entire cloud infrastructure. The company aims to offer access to high-end GPUs like Nvidia's H100, A100, and H200 at lower costs, positioning itself as an alternative to major tech giants. Co-founders Tim Harris and Mike Henry believe AI infrastructure will become more modular and horizontally scalable. Parasail currently serves clients like Elicit, Weights & Biases, and Rasa, and raised \$10 million in funding in 2024. The company plans to compete directly with hyperscalers	By Rebecca Szkutak		April 2, 2025



AI Chips					
#	Highlights	Summary	Author	Source	Date
2.8	Intel and TSMC Reportedly Form Joint Venture to Advance Chipmaking	Intel and TSMC are reportedly forming a joint chipmaking venture to accelerate advanced semiconductor manufacturing amid rising AI demand. The collaboration could involve shared fabrication facilities and co-development of next-gen process technologies, aiming to boost production capacity and reduce geopolitical supply chain risks. Both companies are under pressure to meet surging global demand for AI-optimized chips while navigating U.S.-China tech tensions. If confirmed, the alliance would mark a significant shift in the semiconductor landscape, blending Intel's design capabilities with TSMC's foundry leadership to power future AI infrastructure.	By Rebecca Szkutak		April 3, 2025
2.9	Samsung Q1 Profits Fall 21% Amid Sluggish AI Chip Sales and Foundry Losses	Samsung Electronics reported a 21% decline in Q1 2025 operating profit, largely due to weaker AI chip sales and persistent losses in its foundry division. While the broader semiconductor market shows signs of recovery, Samsung continues to trail competitors like Nvidia in the fast-expanding AI chip sector. Its foundry unit also struggled with profitability, facing low yields and intense pricing pressure. Analysts warn that Samsung must accelerate innovation and strengthen its AI semiconductor strategy to remain competitive amid soaring global demand for AI infrastructure and advanced chip technologies.	By Heekyong Yang		April 7, 2025
2.10	Foxconn Posts Record Q1 Revenue on Robust AI Server Demand	The Foxconn reported record-high Q1 2025 revenue, driven by surging demand for AI servers and a rebound in consumer electronics. The Taiwanese tech giant, a key supplier to companies like Apple and Nvidia, saw revenues climb to \$53.2 billion—up 12.6% year-over-year. The growth	By Reuters		April 5, 2025




AI Chips					
#	Highlights	Summary	Author	Source	Date
		was largely fueled by global investments in AI infrastructure, particularly high-performance servers used for training and deploying large language models. Foxconn's strong performance underscores the ongoing hardware boom underpinning the AI revolution, as cloud providers and enterprises to expand data center capabilities.			
2.11	Japan's Rapidus in Talks with Apple and Google to Mass-Produce Advanced AI Chips	Japanese semiconductor startup Rapidus is in talks with Apple and Google to mass-produce advanced 2-nanometer AI chips by 2027, according to a Nikkei report. Backed by the Japanese government, Rapidus aims to establish a domestic supply of high-performance chips critical for AI and other compute-intensive applications. The collaboration would help Apple and Google diversify their supply chains and reduce dependency on foreign foundries like TSMC. If finalized, the deal would mark a major milestone in Japan's push to reclaim a leading role in global chip manufacturing amid rising geopolitical tensions.	By Reuters		April 4, 2025
2.12	World's first photon-based NPU is 50x faster and uses 30x less power	Researchers have developed the world's first photon-based Neural Processing Unit (NPU), offering remarkable advancements in speed and energy efficiency. This innovative NPU is 50 times faster than traditional silicon-based processors and uses 30 times less power, making it an ideal solution for AI tasks requiring massive computational resources. By leveraging light rather than electricity, this technology overcomes the limitations of conventional electronic circuits, enabling faster data processing and lower energy consumption. The breakthrough in photon-based processing could revolutionize AI applications, particularly in areas	By Joshua Shavit		April 6, 2025



AI Chips					
#	Highlights	Summary	Author	Source	Date
		like machine learning, autonomous systems, and high-performance computing.			
2.13	BrainChip and ISL advance AI-powered radar for military and aerospace	BrainChip Holdings Ltd has partnered with Information Systems Laboratories (ISL) to advance AI-powered radar systems for military and aerospace applications. The collaboration leverages BrainChip's Akida neuromorphic processor, which enables high-performance, ultra-low-power edge computing. This AI-driven radar technology will enhance applications ranging from drones to large systems in defense. ISL's expertise in radar signal processing is complemented by Akida's capabilities, providing real-time intelligence for radar platforms. The partnership aims to improve the effectiveness and efficiency of radar systems, demonstrating the potential of AI chips in military and aerospace sectors.	By Stephen Mayhew		April 7, 2025
2.14	Lightmatter releases new photonics technology for AI chips	Lightmatter, a leader in photonic computing, has unveiled cutting-edge photonics technology for AI chips. By using optical connections through silicon photonics, this new technology improves data transfer speeds and reduces power consumption, offering a significant advancement in AI processing. This development positions Lightmatter as an important player in the AI chips sector, aiming to address the growing demands for more efficient and powerful computing solutions. The integration of photonics into AI chips has the potential to revolutionize the industry, providing enhanced performance and energy efficiency for AI applications across various sectors.	By Stephen Nellis		April 1, 2025



✦ LLM Techniques & Metrics



#	Highlights	Summary	Author	Source	Date
3.1	Multi-Token Attention	Soft attention enables LLMs to identify relevant context, but standard mechanisms rely on a single query-key pair, limiting the depth of information used. To overcome this, researchers introduce <i>Multi-Token Attention (MTA)</i> , which allows attention weights to depend on multiple query and key vectors simultaneously. By applying convolution operations across queries, keys, and heads, MTA enables nearby tokens to influence one another, enriching attention with broader contextual signals. Evaluations show MTA outperforms standard Transformers on language modeling and long-context retrieval tasks, highlighting its ability to better capture nuanced information.	By Olga Golovneva, Tianlu Wang, Jason Weston, Sainbayar Sukhbaatar		April 1, 2025




✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
3.2	Video-T1: Test-Time Scaling for Video Generation	This study introduces <i>Test-Time Scaling (TTS)</i> for video generation, enabling models to improve output quality by leveraging more compute at inference rather than during training. Researchers frame video generation as a search problem—finding optimal trajectories from noise to the desired video guided by test-time verifiers and heuristics. They propose a method called <i>Tree-of-Frames (ToF)</i> , which efficiently generates videos by adaptively expanding and pruning branches. Experiments show that increasing test-time compute significantly enhances video quality across benchmarks, offering a scalable alternative to costly model training.	By Fangfu Liu, Hanyang Wang, Yimo Cai, Kaiyan Zhang, Xiaohang Zhan, Yueqi Duan		April 1, 2025
3.3	MixerMDM: Learnable Composition of Human Motion Diffusion Models	Generating human motion from textual descriptions is challenging due to the need for high-quality motion-condition datasets, especially for fine-grained control. Prior works combine multiple motion diffusion models pre-trained on different conditions but overlook the optimal integration of their generative processes. We introduce MixerMDM , the first learnable model composition technique for text-conditioned motion diffusion. Unlike previous approaches, MixerMDM dynamically learns to mix denoising steps based on the given conditions. By integrating single- and multi-person models, it enables precise control over individual dynamics and interactions. We also propose a novel evaluation method to assess alignment and adaptive mixing quality.	By Pablo Ruiz-Ponce, German Barquero, Cristina Palmero, Sergio Escalera, José García-Rodríguez		April 1, 2025



✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
3.4	New AI benchmarks test speed of running AI applications	MLCommons, an AI industry group, has launched two new benchmark tests to measure how efficiently advanced hardware and software can run AI workloads. One of the benchmarks uses Meta's Llama 3.1 model with 405 billion parameters to test how well systems handle complex queries and integrate data across tasks like question answering, coding, and math. Companies such as Nvidia and Dell Technologies participated, with Nvidia's latest Grace Blackwell AI servers showing 2.8 to 3.4 times faster performance than the previous generation. These benchmarks aim to advance the efficiency and performance of modern AI applications across the tech industry.	By Max A. Cherney, Stephen Nellis		April 2, 2025
3.5	PaperBench: Evaluating AI's Ability to Replicate AI Research	PaperBench, a benchmark designed to assess AI agents' ability to replicate cutting-edge AI research. The benchmark includes 20 ICML 2024 papers, requiring agents to interpret contributions, build codebases, and run experiments. Each task is broken into 8,316 sub-tasks using detailed rubrics co-developed with the original authors. To scale grading, we introduce an LLM-based judge evaluated on its own benchmark. Testing frontier models, Claude 3.5 Sonnet (with open-source scaffolding) achieved the top score at 21.0%. Comparisons with top ML PhDs show that AI models still lag behind human performance in replicating AI research.	By Giulio Starace et al.		April 2, 2025
3.6	MLPerf Unveils New Benchmarks to Measure Real-	MLCommons has released updated MLPerf inference benchmarks to assess how efficiently AI models perform in real-world scenarios. The benchmarks cover key applications like large language models, recommendation systems, and computer vision, tested across both data	By Max A. Cherney and Stephen Nellis		April 2, 2025

✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
	World AI Inference Speed	center and edge hardware. Major players such as Nvidia, Intel, Qualcomm, and Google submitted results, showcasing improvements in processing speed and energy efficiency. These standardized tests offer a reliable way to compare AI performance across platforms, reflecting real application demands rather than just theoretical capabilities.			
3.7	Qwen2.5-Max Vulnerability Assessment	Protect AI's assessment of Qwen2.5-Max, a large-scale mixture of experts model, revealed a medium risk score of 35 out of 100. The evaluation, conducted using Protect AI's Recon tool, tested the model against over 400 attack techniques, including evasion, system prompt leaks, prompt injections, jailbreaks, safety issues, and adversarial suffixes. The model was found to be most vulnerable to prompt injection and evasion attacks, with 140 successful breaches, over 94 of which were classified as critical or high severity. This highlights the need for enhanced security measures in large language models.	By Mukunth Madavan & Sailesh Mishra		April 2, 2025
3.8	Efficient Model Selection for Time Series Forecasting via LLMs	Model selection is essential for accurate time series forecasting, yet it often demands extensive evaluations across datasets—making it resource-intensive. Traditional meta-learning methods automate this process but rely heavily on costly performance matrices. In this study, we introduce a novel approach that uses Large Language Models (LLMs) like LLaMA, GPT, Gemini to streamline model selection. By tapping into their reasoning abilities, we eliminate the need for explicit performance matrices. Our experiments show that LLMs not only outperform conventional meta-learning and heuristic methods but also greatly reduce computational costs.	By Wang Wei, Tiankai Yang, Hongjie Chen, Ryan A. Rossi, Yue Zhao, Franck Démoncourt, Hoda Eldardiry		April 2, 2025


✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
		This highlights the promising role of LLMs in efficient time series forecasting.			
3.9	Rethinking RL Scaling for Vision Language Models: A Transparent, From-Scratch Framework and Comprehensive Evaluation Scheme	Reinforcement learning (RL) is enhancing reasoning in large language models and expanding to vision-language models (VLMs). However, current RL applications in VLMs often rely on complex frameworks, limiting reproducibility and lacking standardized evaluation. This work introduces a transparent, minimal four-step RL framework for VLMs, validated across models and datasets. A standardized evaluation scheme is proposed to analyze training dynamics and reflection. Experiments on visual reasoning reveal key insights: response length varies with random seeds, reflection links to output length, and RL surpasses supervised fine-tuning in generalization. This framework aims to establish a reproducible RL-based VLM baseline.	By Yan Ma, Steffi Chern, Xuyang Shen, Yiran Zhong, Pengfei Liu		April 3, 2025
3.10	OpenAI copyright lawsuits from authors, New York Times consolidated in Manhattan	Several copyright lawsuits against OpenAI and Microsoft, including one filed by The New York Times and another by a group of authors, have been consolidated in Manhattan federal court. These cases allege that OpenAI used copyrighted materials without permission to train its AI models, raising significant legal questions about generative AI and intellectual property rights. The consolidation aims to streamline proceedings and address common issues across the lawsuits. Plaintiffs argue that AI-generated content based on copyrighted works infringes upon their rights, while OpenAI and Microsoft maintain that their use falls under fair use and is transformative in nature.	By Blake Brittain		April 4, 2025




✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
3.11	Multi-SWE-bench: A Multilingual Benchmark for Issue Resolving	Multi-SWE-bench is a multilingual benchmark designed to evaluate LLMs on issue-resolving tasks across seven languages, including Java, C++, and Rust. Unlike prior benchmarks focused mainly on Python, it provides 1,632 expert-annotated instances from 2,456 candidates, enabling reliable evaluation. The study assesses top models using Agentless, SWE-agent, and OpenHands methods, offering valuable insights. Additionally, the authors introduce the Multi-SWE-RL open-source community and release 4,723 structured RL training instances. By open-sourcing the entire pipeline and tutorials, they encourage ongoing contributions. This work advances reinforcement learning in software engineering and pushes the field closer to AGI.	By Daoguang Zan, Zhirong Huang, et al.		April 3, 2025
3.12	MME-Unify: A Comprehensive Benchmark for Unified Multimodal Understanding and Generation Models	MME-Unify, a benchmark to evaluate Unified Multimodal Large Language Models (U-MLLMs), addressing key limitations in current benchmarks. It offers a standardized evaluation of 10 tasks and 30 subtasks from 12 datasets for fair comparisons. Additionally, it proposes five new tasks focused on multimodal reasoning, such as image editing, commonsense QA with image generation, and geometric reasoning. The authors benchmark 12 top U-MLLMs, including Janus-Pro and Gemini2-flash, as well as specialized models like Claude-3.5-Sonnet and DALL-E-3. Results show notable performance gaps, underscoring the need for more capable mixed-modality models.	By Wulin Xie, Yi-Fan Zhang, Chaoyou Fu, Yang Shi, Bingyan Nie, Hongkai Chen, Zhang Zhang, Liang Wang, Tieniu Tan		April 4, 2025



✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
3.13	MCTS-RAG Introduces Monte Carlo Tree Search to Improve Retrieval-Augmented Generation	The paper "MCTS-RAG: Enhance Retrieval-Augmented Generation with Monte Carlo Tree Search" proposes a novel method that integrates Monte Carlo Tree Search (MCTS) into Retrieval-Augmented Generation (RAG) pipelines. MCTS-RAG decomposes complex queries into sub-questions, explores reasoning paths, and selectively retrieves evidence to guide final answer generation. The approach improves factuality and retrieval precision by simulating multiple reasoning trajectories before committing to an answer. Evaluated on multi-hop QA datasets like HotpotQA and IIRC, MCTS-RAG achieves state-of-the-art performance with enhanced evidence grounding and lower hallucination rates.	By Université Paris-Dauphine, PSL University		April 4, 2025
3.14	ReACT+ Improves Reasoning by Combining Chain-of-Thought and Direct Answering	The paper "ReACT+: Efficient and Effective Reasoning with Unified CoT and Direct Answering" introduces ReACT+, a method that unifies Chain-of-Thought (CoT) prompting and direct answering to improve reasoning performance in LLMs. Unlike traditional approaches that commit to either CoT or direct responses, ReACT+ dynamically evaluates both paths and selects the superior answer using log probability scoring. This strategy enhances efficiency by reducing unnecessary CoT steps and improves accuracy across reasoning-intensive benchmarks like GSM8K and StrategyQA. ReACT+ outperforms previous techniques with fewer inference calls, offering a lightweight yet powerful reasoning enhancement.	By Tu Ao et al.		April 4, 2025
3.15	Z1: Efficient Test-time Scaling with Code	This paper introduces Z1, a method to enhance Large Language Models' (LLMs) efficiency in complex problem-solving by reducing excessive reasoning tokens during inference. The authors curate Z1-Code-	By Zhaojian Yu et al.		April 1, 2025f



✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
		Reasoning-107K, a dataset of coding problems with both short and long solution trajectories. They propose a Shifted Thinking Window technique, which removes context-delimiting tags and caps reasoning tokens, allowing the model to adjust its reasoning depth based on problem complexity. Trained with this approach, Z1-7B matches the performance of larger models like R1-Distill-Qwen-7B while using approximately 30% fewer reasoning tokens, demonstrating efficient test-time scaling across various reasoning tasks.			
3.16	Why do LLMs attend to the first token?	This paper investigates the phenomenon where Large Language Models (LLMs) disproportionately focus attention on the first token of the input, forming an “attention sink.” While previous studies noted this behavior, the authors provide a new theoretical and empirical explanation: attention to the first token helps prevent over-mixing of information across the sequence, maintaining better representation structure. Through controlled experiments, they show how factors like context length, model depth, and data packing influence this effect. The findings offer insights into the training dynamics and architecture choices of transformer-based models.	By Federico Barbero et al.		April 4, 2025
3.17	TransMamba: Flexibly Switching between Transformer and Mamba	The paper introduces TransMamba, a novel framework that unifies Transformer and Mamba models through shared parameter matrices, enabling dynamic switching between attention and state space model (SSM) mechanisms. This approach addresses the limitations of both models: Transformers' quadratic complexity in long sequences and Mamba's unstable contextual learning. TransMamba achieves superior	By Yixing Li2, Ruobing Xie et al.		March 31, 2025




✦ LLM Techniques & Metrics					
#	Highlights	Summary	Author	Source	Date
		training efficiency and performance across various tasks, validated through extensive experiments. The framework's Memory Converter ensures seamless information flow during transitions, while TransPoint scheduling optimizes the model's structure. TransMamba offers a scalable solution for next-generation sequence modeling, bridging the gap between efficiency and effectiveness in large language models.			



✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
4.1	Adobe Unveils Generative AI Tools for 4K Video Editing at NAB 2025	At NAB 2025, Adobe introduced new generative AI features for video editing, including Generative Extend, which can seamlessly add frames to extend clips in 4K resolution. Integrated into Adobe Premiere Pro, the updates also include AI-powered object addition, removal, and generative fill for video content. These tools aim to significantly speed up post-production workflows for filmmakers and content creators. Adobe's advancements reflect the expanding role of generative AI in creative industries, following the success of its Firefly models for still images. Public beta testing for the new features begins later this year.	By Tech Desk		April 2, 2025




✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
4.2	This Financial Firm Can Give Investment Advice in Gen Z Slang, No Cap	Arta Finance is introducing an AI investment assistant that offers financial advice tailored to user preferences—including Gen Z slang like “no cap.” Designed to feel conversational, the tool adjusts tone and language based on the user’s age, goals, and risk tolerance. Arta's AI evaluates financial data and behavior to recommend personalized investment strategies, such as stock portfolios or real estate allocations. The goal is to democratize wealth management, making sophisticated financial guidance accessible to everyday users. With this new approach, Arta aims to bridge the gap between traditional finance and digital-native audiences.	By Hannah Erin Lang		April 1, 2025
4.3	Staircase Studio AI is turning Hollywood’s anxiety into opportunity	Staircase Studios AI is redefining filmmaking by combining artificial intelligence with human creativity to produce high-quality, low-budget films. Founded by “Divergent” producer Pouya Shahbazian, the studio employs writers, directors, and actors while using AI for tasks like generating sets and visuals. Its first film, The Woman with Red Hair, involves real talent in every stage, from scriptwriting to voice acting. Unlike fully AI-generated content, Staircase prioritizes artistic involvement. With a projected production cost under \$500,000 per film, the studio aims to release up to eight features annually, transforming Hollywood’s AI fears into creative opportunity.	By Will Zimmerman		April 1, 2025
4.4	Greece vows to spend \$27B on armed forces overhaul centered	Greece has announced a €25 billion (\$27 billion) defense modernization plan over the next decade, signaling a major shift in its military strategy. At its core is the "Achilles Shield," a new air defense system aimed at countering regional threats, particularly from Turkey. The plan emphasizes	By Derek Gatopoulos		April 2, 2025

✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
	on high-tech warfare technology	mobile, AI-driven missile systems, drone warfare, and upgraded command infrastructure. It also includes next-gen soldier gear with integrated sensors and communication tools, along with the creation of satellite networks for secure wartime communication. Structural changes—such as merging units and closing outdated bases—highlight Greece’s broader efforts to enhance military efficiency and readiness.			
4.5	Prompting Medical Vision-Language Models to Mitigate Diagnosis Bias by Generating Realistic Dermoscopic Images	This study introduces the Dermatology Diffusion Transformer (DermDiT), a novel framework designed to address diagnostic biases in AI-based skin disease detection, particularly those related to underrepresented groups. DermDiT utilizes large vision-language models (VLMs) to generate detailed text prompts for dermoscopic images. These prompts guide a diffusion transformer model to produce synthetic, high-quality dermoscopic images that enhance the representation of minority subgroups in imbalanced datasets. Extensive experiments demonstrate that incorporating VLM-generated prompts significantly improves the quality and diversity of the generated images, thereby contributing to more equitable diagnostic outcomes.	By Nusrat Munia, Abdullah-Al-Zubaer Imran		April 2, 2025
4.6	LaLiga Embraces AI to Enhance Global Fan Engagement and Match Analysis	LaLiga is pioneering AI integration in football by launching AI-driven tools for global fan engagement, performance analytics, and real-time content creation. Collaborating with partners like Microsoft and sportstech firm WSC Sports, LaLiga now uses AI to automatically generate and distribute customized highlight reels and insights tailored to regional audiences. This strategy aims to grow the league’s international presence and deepen fan	By Ahmed El Khashab		April 2, 2025



✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
		loyalty. By merging data analytics with AI storytelling, LaLiga positions itself at the forefront of digital transformation in sports.			
4.7	Papa Johns wants AI to transform pizza ordering	Papa John's International is expanding its partnership with Google Cloud to integrate artificial intelligence into its ordering process. The initiative aims to personalize customer interactions by analyzing past behaviors to send targeted push notifications, marketing emails, and loyalty program offers. Additionally, the company plans to introduce an online chatbot and enable ordering through virtual assistants. This move reflects a broader trend in the fast-food industry to leverage AI for enhancing sales and customer service. CEO Todd Penegor and Chief Digital Officer Kevin Vasconi, both with prior experience implementing AI strategies at Wendy's and Domino's, are leading this technological advancement.	By Waylon Cunningham		April 3, 2025
4.8	Querying Hugging Face Datasets with the DuckDB UI	The article "Querying Hugging Face Datasets with the DuckDB UI" introduces the integration of DuckDB's local user interface with Hugging Face datasets. This integration allows users to efficiently query and analyze large datasets by leveraging their local machine's resources, overcoming browser limitations. The article outlines two primary methods for connecting to Hugging Face datasets: using the hf:// protocol within DuckDB's httpfs extension, and utilizing the "Copy for DuckDB CLI" feature in Hugging Face's Data Studio. These approaches aim to enhance data exploration and analysis capabilities for users working with extensive datasets.	By Caleb Fahlgren		April 3, 2025

✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
4.9	Amazon Integrates AI-Generated Recaps for Book Series on Kindle	Amazon has introduced an AI-powered recap feature for Kindle, designed to help readers of book series recall past plotlines before diving into the next installment. Available initially for select fiction titles in the U.S., the tool generates concise, spoiler-managed summaries using generative AI. Readers can access these recaps from the book's main menu, making it easier to reengage with complex narratives. The feature reflects Amazon's broader push to enhance Kindle's user experience through AI, offering more seamless and intelligent reading support.	By Aisha Malik		April 3, 2025
4.10	Amazon Unveils AI Shopping Agent That Buys from Third-Party Stores	Amazon has introduced a new AI-powered shopping agent capable of purchasing products from third-party retailers, not just Amazon's own store. Designed to act as a personalized shopping assistant, the agent can understand natural language requests, compare options across platforms, and complete purchases on behalf of users. This marks a strategic expansion of Amazon's AI ecosystem beyond its marketplace, aiming to dominate the broader e-commerce funnel. The move reflects growing consumer demand for intelligent agents that streamline complex online shopping journeys.	By Maxwell Zeff		April 4, 2025
4.11	Chinese AI-Generated Videos Mock U.S. Tariffs with Satirical	AI-generated videos featuring sarcastic robot characters are gaining traction on Chinese social media, mocking U.S. tariff policies in response to Donald Trump's proposals. Created by influencer "Sister Abalone" and others, the satirical content blends political commentary with generative AI tools like text-to-speech and animation. The videos portray American consumers and robots lamenting rising costs due to trade restrictions,	By Antoni Slodkowski		April 5, 2025


✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
		reflecting growing tech-enabled political discourse. Analysts say this trend shows how generative AI is becoming a tool for propaganda, humor, and dissent—highlighting its rising cultural and geopolitical influence.			
4.12	CMA CGM Partners with French AI Startup to Enhance Customer Service	Global shipping giant CMA CGM has partnered with French AI startup Target to improve its customer service using large language model (LLM) technology. The collaboration aims to automate and streamline responses to client inquiries across shipping, logistics, and supply chain services. By integrating Target’s AI platform, CMA CGM expects faster response times, greater personalization, and reduced workload for human agents. This partnership reflects a broader trend of AI adoption in traditional industries, leveraging generative AI to boost efficiency and customer satisfaction in complex operational environments.	By Gus Trompiz and Florence Loeve		April 6, 2025
4.13	APIGen-MT: Agentic Pipeline for Multi-Turn Data Generation via Simulated Agent-Human Interplay	APIGen-MT, a framework designed to generate high-quality multi-turn interaction data for training AI agents. It works in two phases: first, creating task blueprints with detailed actions using multiple Large Language Models (LLMs) and iterative feedback, and second, generating complete interaction trajectories through simulated agent-human interplay. This approach addresses challenges in maintaining complex dependencies and ensuring data quality for realistic conversations. Experiments show that models trained with APIGen-MT outperform existing benchmarks, demonstrating its effectiveness in improving agent capabilities. The framework offers a new solution for enhancing AI agent training with multi-turn dialogue data.	By Akshara Prabhakar et al.		April 4, 2025




✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
4.14	Shipping giant CMA CGM and French AI startup target customer service in tie-up	Shipping giant CMA CGM and French AI startup Mistral AI have entered a €100 million, five-year partnership to enhance customer service in shipping and logistics, as well as fact-checking for CMA CGM's French media assets like BFM TV. This collaboration increases CMA CGM's AI-related investment to €500 million. CEO Rodolphe Saade anticipates that initiatives, such as reducing response times for the one million weekly customer emails, will be implemented within 6 to 12 months. Mistral AI, recognized as a European AI leader, expects a tenfold revenue increase by December 2025.	By Gus Trompiz and Florence Loeve		April 6, 2025
4.15	Warner Bros. Discovery Launches Generative AI-Powered Cycling Central Intelligence Platform	Warner Bros. Discovery (WBD) Sports Europe has unveiled the Cycling Central Intelligence (CCI) platform, a generative AI-driven system developed with Amazon Web Services (AWS). Launched at the 2025 WHOOP UCI Mountain Bike World Series opener in Araxá, Brazil, CCI aims to enhance sports broadcasting by providing real-time data on riders, race history, and venues to commentators. Powered by AWS technologies like Amazon Bedrock and Anthropic's Claude 3.5, CCI allows for natural language queries and data synthesis, enriching the storytelling experience and improving viewer engagement in cycling events.	By George Winslow		April 5, 2025
4.16	Maison de Vapotage and Privately SA on schedule to perform millions of	Maison de Vapotage has partnered with Privately to deploy AgeAI, an AI-powered age verification system for French retailers. This system helps ensure compliance with regulations on age-restricted products like tobacco, vapes, and alcohol. Launched in over 40 stores by March, AgeAI handles more than 10,000 daily age checks. The collaboration enables tobacconists	By Privately SA		April 7, 2025

✦ AI Use Cases




#	Highlights	Summary	Author	Source	Date
	AI-based age checks quarterly supporting retailers in France	to automatically verify customer ages, providing a seamless, efficient solution that adheres to legal requirements. By leveraging AI, this system significantly improves the process of age verification in retail environments, making it faster and more reliable for both businesses and consumers.			
4.17	Dynamic Importance in Diffusion U-Net for Enhanced Image Synthesis	The paper presents Dynamic Importance in Diffusion U-Net (DIDU) , a novel approach to improve diffusion-based image synthesis by dynamically weighting feature importance in U-Net architectures. Unlike standard diffusion models, DIDU identifies and enhances critical spatial and channel features during denoising, leading to higher-quality outputs. The method introduces adaptive gating mechanisms to prioritize relevant regions, reducing artifacts and improving detail preservation. Experiments demonstrate superior performance in image fidelity and generation control compared to baseline models. DIDU offers a scalable way to refine diffusion models without extra computational overhead, making it efficient for high-resolution synthesis tasks.	By Xi Wang et al.		April 4, 2025
4.18	Advances and Challenges in Foundation Agents: From Brain-Inspired Intelligence to	The paper, <i>Advances and Challenges in Foundation Agents: From Brain-Inspired Intelligence to Evolutionary, Collaborative, and Safe Systems</i> , explores the transformative potential of Large Language Models (LLMs) in creating intelligent agents. It introduces a modular, brain-inspired architecture integrating principles from neuroscience and cognitive science	By Bang Liu, Xinfeng Li et al.		March 31, 2025



✦ AI Use Cases					
#	Highlights	Summary	Author	Source	Date
	Evolutionary, Collaborative, and Safe Systems	to enhance cognition, perception, and operational capabilities. The survey emphasizes self-evolution, multi-agent collaboration, and critical safety measures for real-world deployment. By addressing research gaps and proposing solutions, the paper provides a roadmap for building adaptive, secure, and socially beneficial AI systems, marking a significant step toward autonomous, human-like intelligence.			




🛡️ AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
5.1	China's love of open-source AI may shut down fast	China has rapidly embraced open-source AI as a strategic move amid U.S. tech sanctions. Companies like DeepSeek, Alibaba, and Baidu have released powerful, freely available models to close the AI gap with the West. Open-source innovation offers China flexibility while reducing reliance on foreign chips, aligning with President Xi's push for technological independence. However, this openness could be short-lived if geopolitical tensions rise. The European Union is also adopting a similar strategy, investing €200 billion in cooperative AI innovation. China's open-source push not only accelerates tech growth but also enhances its global influence and soft power.	By Robyn Mak		April 2, 2025




 AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
5.2	Tony Blair Institute Calls for Overhaul of UK AI Copyright Laws	<p>The Tony Blair Institute has urged the UK government to reform outdated copyright laws that hinder AI development, especially in comparison to more permissive U.S. regulations. A new report criticizes the UK's restrictive stance on text and data mining, arguing it impedes innovation and puts the country at a competitive disadvantage in the global AI race. The think tank advocates aligning UK policy with U.S. norms to promote AI research, while still ensuring fair compensation for creators. The debate reflects growing transatlantic tensions over intellectual property and AI training data access.</p>	By Dan Milmo		April 2, 2025
5.3	Taking a Responsible Path to AGI	<p>Google DeepMind has released a comprehensive 145-page strategy paper on AGI (Artificial General Intelligence) safety. The document suggests that AGI could be developed by 2030 and may pose existential risks to humanity. DeepMind proposes safeguards like access controls to prevent misuse and technical tools to better understand AI behavior. However, some experts argue that the assumptions in the paper lack scientific grounding, particularly around concepts like recursive self-improvement. While the paper aims to address serious safety concerns, it is unlikely to end ongoing debates about AGI's risks and nature. It highlights institutional strategies for responsible AI development.</p>	By Anca Dragan, Rohin Shah, Four Flynn and Shane Legg		April 2, 2025

🛡️ AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
5.4	Evaluating potential cybersecurity threats of advanced AI	DeepMind has published an in-depth analysis evaluating how advanced AI systems could pose cybersecurity threats. The new framework introduces an “offensive cyber capability benchmark” with 50 challenges covering all stages of a cyberattack. By examining over 12,000 AI-driven attack attempts across 20 countries, the study shows that AI can significantly reduce the cost and effort of executing phishing, malware, and DDoS attacks. While current AI systems are not yet breakthrough threats on their own, the scale and complexity of future risks are expected to grow. The report urges the security community to develop proactive defense strategies.	By Four Flynn, Mikel Rodriguez and Raluca Ada Popa		April 2, 2025
5.5	US plans to develop AI projects on Energy Department lands	The U.S. Department of Energy has identified 16 potential sites on its lands for developing data centers and power plants to support the rapid growth of artificial intelligence (AI). These sites, including the Idaho National Laboratory and facilities in Kentucky and Ohio, offer existing energy infrastructure and expedited permitting processes for new energy generation, such as nuclear reactors. Energy Secretary Chris Wright emphasized the importance of this initiative, likening the global AI competition to the "next Manhattan Project." The DOE is encouraging public-private partnerships to advance these developments.	By Timothy Gardner		April 3, 2025
5.6	OpenAI Makes First Cybersecurity Investment in Defender AI Startup	OpenAI has made its first cybersecurity investment by backing <i>Defend</i> , a startup focused on developing AI-native security tools that can autonomously detect and mitigate cyber threats in real time. <i>Defend</i> aims to replace traditional rule-based systems with intelligent agents capable of	By Charles Rollet		April 3, 2025

 AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
		<p>understanding complex attack patterns and adapting defenses dynamically. The investment underscores OpenAI's broader commitment to AI safety and the growing need to secure AI systems from misuse, data breaches, and adversarial attacks. As AI becomes more embedded in enterprise and critical infrastructure, this move highlights the convergence of cybersecurity and advanced AI development.</p>			
5.7	<p>U.S. Energy Department Designates 16 Sites for AI and Data Center Development</p>	<p>The U.S. Department of Energy has identified 16 federal sites across the country to host new data centers and AI infrastructure as part of a national effort to boost AI capabilities and energy innovation. These locations will support advanced computing for scientific research, national security, climate modeling, and clean energy development. The initiative aims to balance rising AI compute demands with sustainable infrastructure, including energy-efficient design and renewable integration. This marks a major federal push to build AI capacity within government while supporting technological competitiveness and climate goals.</p>	<p>By U.S Department of Energy</p>		<p>April 3, 2025</p>
5.8	<p>Trump Tariff Plan May Disrupt Big Tech's U.S. Data Center Expansion</p>	<p>Donald Trump's proposed 10% tariff on all imports could jeopardize major tech firms' aggressive spending on U.S. data centers, critical for AI infrastructure. Industry leaders like Amazon, Google, and Microsoft rely heavily on imported components—servers, chips, and networking gear—to build these facilities. Analysts warn that broad tariffs could raise costs, delay</p>	<p>By Big Tech</p>		<p>April 4, 2025</p>

AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
		construction, and reduce the competitiveness of U.S.-based AI development. The potential policy shift highlights the vulnerability of AI infrastructure to geopolitical decisions and trade policy, raising concerns across the tech sector about future expansion and innovation.			
5.9	Robots, fraught consumers star in China AI videos mocking tariffs	The Reuters article discusses how China's state media is using AI-generated videos to criticize U.S. tariffs. These videos, featuring dancing robots and anxious consumers, portray the negative economic impact of U.S. trade policies. For example, one video mocks the loss of cheap Chinese cars due to tariffs, while another shows a robot named "TARIFF" choosing self-destruction over trade war chaos. This reflects how AI is being used as a tool for political messaging and influence. In this case, it's a clear example of AI being used not just as a tech application, but as part of a political strategy.	By Antoni Slodkowski		April 5, 2025
5.10	Microsoft AI CEO's remarks interrupted by pro-Palestinian protester	During Microsoft's 50th anniversary event in Redmond, Washington, AI CEO Mustafa Suleyman was interrupted by employee Ibtihal Aboussad, who accused the company of profiting from war and using AI for genocide. Suleyman acknowledged the protest before Aboussad was escorted away. This incident follows reports that Microsoft's AI models were used by the Israeli military to select bombing targets in Gaza and Lebanon. Microsoft	By Kanishka Singh		April 5, 2025

 AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
		stated it supports free expression through proper channels. Aboussad and another protester reportedly lost access to their work accounts after the protest.			
5.11	Elon Musk's Lawsuit Against OpenAI Set for Jury Trial in Spring 2026	A California judge has scheduled a jury trial for spring 2026 in Elon Musk's lawsuit against OpenAI, alleging the company abandoned its nonprofit mission in favor of profit-driven partnerships, particularly with Microsoft. Musk, a co-founder of OpenAI, claims the organization breached its founding agreement by prioritizing closed-source development. OpenAI has denied the allegations, asserting that its current structure aligns with its long-term mission. The high-profile case is expected to spotlight tensions over AI governance, openness, and commercial influence in the rapidly evolving artificial intelligence landscape.	By Anna Tong		April 5, 2025
5.12	Meta Plans \$1 Billion AI-Optimized Data Center in Wisconsin	Meta is investing nearly \$1 billion to build a new data center in Mount Pleasant, Wisconsin, designed to support its growing AI infrastructure needs. The facility will be optimized for training and deploying large language models and other generative AI systems. Construction is expected to complete by 2026, with Meta citing the state's energy resources and talent pool as key factors. This move underscores Big Tech's race to	By Bloomberg News		April 4, 2025

 AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
		expand U.S.-based AI compute capacity while addressing concerns over energy use, localization, and long-term infrastructure resilience.			
5.13	Could AI fix Social Security? Its likely new boss thinks so — but critics have serious doubts	<p>The incoming leadership of the U.S. Social Security Administration is considering the use of artificial intelligence to improve internal processes such as handling disability claims and reducing administrative delays. The plan suggests integrating AI tools to automate decision-making and boost operational efficiency. While some argue this could modernize outdated systems and enhance service delivery, others express serious concerns about potential risks like wrongful denials, algorithmic bias, and lack of accountability. This development reflects a concrete example of how AI might be used in public services and raises important questions about responsible implementation in government systems.</p>	By Alessandra Malito		April 5, 2025
5.14	Chinese State Media Rebuke Trump's Tariffs With AI Song and Videos	<p>Chinese state media have embraced artificial intelligence (AI) to create content addressing U.S. trade policies. CGTN released an AI-generated music video, "Look What You Taxed Us Through," which critiques U.S. tariffs from the perspective of American consumers. Additionally, New China TV published a sci-fi short film, "T.A.R.I.F.F.," in which a robot administers tariffs and explores their broader societal effects. These AI-created productions illustrate how China is using AI to deliver political</p>	By Chad de Guzman		April 4, 2025

AI Policies Regulations & Strategies					
#	Highlights	Summary	Author	Source	Date
		messages and influence public perception, showcasing AI's potential in political communication and state media strategies.			

Conclusion

- The first week of April 2025 marks a narrative of simultaneous scaling and specialization in AI.
- The anticipation surrounding **Alibaba's** Qwen and **Meta's** Llama 4 highlights the ongoing quest for foundational model dominance.
- **Cohere's** Command R+ reflects the growing demand for tailored enterprise solutions.
- The strategic interactions between semiconductor giants like **Intel**, **TSMC**, and **Arm** emphasize that future AI progress is deeply intertwined with hardware innovation and complex supply chains.
- Tangible applications of AI include the **US Social Security Administration (SSA)** optimizing government functions.
- **Chinese state media** using AI-generated content for political narratives.
- AI's accelerating integration into governance and communication highlights the increasing focus on debates around capability, ethics, and regulation, which will continue to intensify in the coming weeks.