# Large Language Models

## Anthropic Delays Mythos AI Model After It Autonomously Exploited Tens of Thousands of Software Vulnerabilities with 80% Success Rate

April 16, 2026

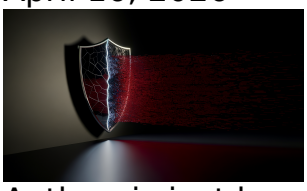

Anthropic just became the first major AI lab to delay a frontier model because it got too good…

## Meta Llama 4 Launches April 5 with 109B-Parameter Scout Model—17B Active Parameters Fit on Single H100, Outperform GPT-4.5 and Claude 3.7 on STEM Benchmarks

April 15, 2026

Meta just made frontier AI fit on one GPU—and simultaneously banned 450 million Europeans from using it.

## OpenAI Launches Age Prediction on ChatGPT January 20—Behavioral Signals Identify Under-18 Users to Block Graphic Violence, Self-Harm Content, and Risky Viral Challenges

April 12, 2026

ChatGPT now profiles your behavior to guess your age—then restricts what you can access based on that prediction.…

## OpenAI o3-mini Launches January 31 with 95% Cost Cut Over GPT-4—First Small Reasoning Model with Function Calling Outperforms o1 on Coding Benchmarks

April 11, 2026

OpenAI’s cheapest reasoning model now beats its most expensive one at writing code. That sentence should make every…

## Anthropic’s Claude Mythos Hits 93.9% SWE-Bench but Won’t Be Released—$25/$125 Token Pricing Reserved for 40+ Whitelisted Security Teams

April 9, 2026

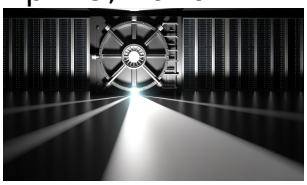

Anthropic just built an AI that solves 94% of real-world software engineering tasks—then decided nobody outside 40 security…

## Liquid AI Releases LFM2.5-1.2B-Thinking: 1.2B Parameter Reasoning Model Runs Fully Offline in 900 MB on Phones

April 6, 2026

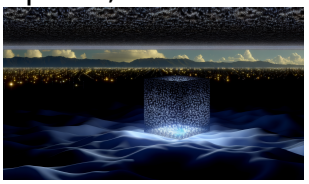

A reasoning model smaller than a mobile game just outscored models twice its size—and it runs entirely on…

## Google Gemma 4 Ranks #3 on Arena AI Leaderboard—31B Open Model Hits 85.2% MMLU Pro and 89.2% AIME 2026, Outperforming Models 20× Larger

April 4, 2026

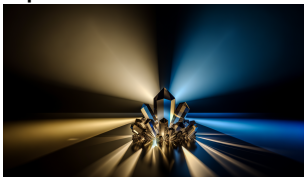

A 31-billion parameter model now outperforms systems with 600B+ parameters on reasoning and math benchmarks. Google just made…

## Gemini 3.1 Pro Preview Hits 94.1% on LM Council Benchmark—Google Takes Three of Top Four Leaderboard Spots as Claude 4.6 Opus and GPT-5.4 Trail by Double Digits

March 31, 2026

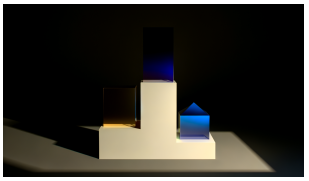

Google just claimed three of the top four spots on major AI benchmarks while charging 60% less than…

## Anthropic’s Model Context Protocol Hits 97 Million Installs on March 25—MCP Transitions from Experimental to Foundation Layer for Agentic AI

March 30, 2026

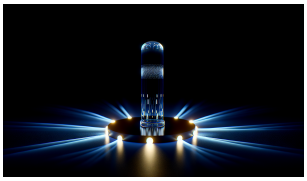

Six months ago, MCP was a footnote in Anthropic’s documentation. Today it’s the plumbing underneath 97 million AI…

## 12 AI Models Launched in One Week During March 10-16, 2026—OpenAI, Google, xAI, Anthropic, Mistral, and Cursor Compress Developer Selection Cycles to Monthly as Frontier Model Releases Pile Up

March 29, 2026

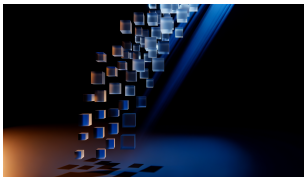

Twelve AI models in seven days. Developer teams that used to evaluate quarterly releases now face monthly decision…

## Moonshot AI’s Attention Residuals (AttnRes) Achieves Same Training Loss with 1.25× Less Compute—New Transformer Architecture Replaces Fixed Residuals with Softmax Attention

March 26, 2026

Residual connections have been sacred architecture for nine years—and Moonshot AI just proved they were holding us back.…

## PromptEng 2025 Workshop Accepts 7 Research Papers on January 24—First Academic Forum Dedicated to Prompt Engineering Techniques

March 23, 2026

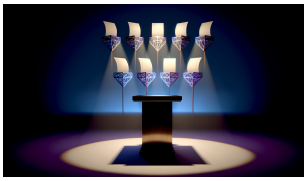

While engineers debated whether prompt engineering was “real engineering,” seven research teams quietly got their papers peer-reviewed and…

## OpenAI Launches GPT-5.4 ‘Thinking’ with 1 Million-Token Context Window on March 5, 2026

March 17, 2026

OpenAI just shipped a model that can hold 2,000 pages of text in working memory. Most enterprise RAG…

## University of Montreal Tests AI Against 100,000 Humans on Creativity—GPT-4 Beats 72% But Top 10% Still Win

March 16, 2026

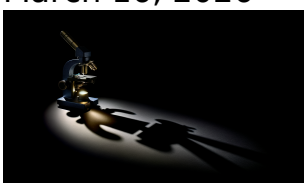

GPT-4 now outperforms the average human on standardized creativity tests. But before you fire your creative team, consider…

## Allen Institute Releases SERA-32B: Open-Source Coding Agent Hits 54.2% SWE-Bench Score for $1,300

March 14, 2026

A 32-billion parameter model now outperforms its 110-billion parameter teacher after fine-tuning on just 8,000 samples. The cost…

## Google NotebookLM Launches Cinematic Video Overviews Using Gemini 3, Veo 3, and Nano Banana Pro—Moving Beyond Narrated Slideshows to Full Animation

March 8, 2026

Google just shipped a feature that treats your research notes like a film production—complete with three AI models…

## Moonshot AI Releases Kimi K2.5 with 100-Agent Swarm Feature—Trained on 15 Trillion Tokens, Beats GPT-5.2 on Coding and Video Benchmarks

March 5, 2026

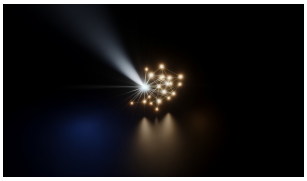

The AI scaling wars just took an unexpected turn: a Chinese startup released a model that orchestrates 100…

## Hackers Claim 20+ Million ChatGPT Access Codes Stolen—Posted on BreachForums Alongside 30,000 OmniGPT User Credentials

March 4, 2026

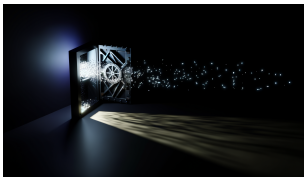

The largest credential theft targeting AI platforms just exposed a brutal truth: the authentication protecting enterprise AI usage…

## University of Montreal Study Proves AI Beats Average Humans on Creativity Tests—But Top 10% Still Outperform GPT-4

February 24, 2026

The world’s largest creativity study just revealed an uncomfortable truth: half of humanity is now less creative than…

## xAI Launches Grok 4.20 Beta with Four-Agent Architecture—65% Hallucination Reduction Shifts Prompt Engineering From Iterative Chat to Structured Contracts

February 23, 2026

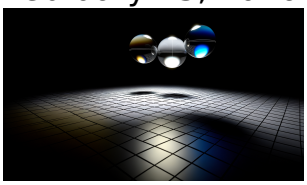

The iterative prompt refinement loop that defined AI workflows for three years just became a liability. Grok 4.20’s…

## Google Gemini 3.1 Pro Scores 77.1% on ARC-AGI-2—2.5x Jump Over Predecessor in Single Generation

February 22, 2026

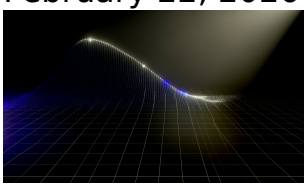

Google just doubled AI reasoning capability in 90 days while keeping the price identical. The assumption that frontier…

## OpenAI Launches GPT-5.3-Codex-Spark: 1,000+ Tokens Per Second Coding Model for Real-Time Development

February 20, 2026

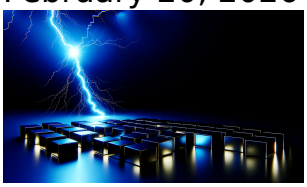

OpenAI just admitted that bigger isn’t always better. Their new coding model generates 1,000+ tokens per second—and it’s…

## Snorkel AI Commits $3M to Open Benchmarks Grant—Targeting the ‘Biggest Blind Spot’ Where AI Models Excel on Tests But Fail in Production

February 16, 2026

Claude Opus 4.6 just scored 76% on MRCR v2—up from 18.5% on its predecessor. GPT-5.3-Codex hit 77.3% on…

## Apple Xcode 26.3 Launches Agentic Coding with Claude Agent and OpenAI Codex Integration

February 14, 2026

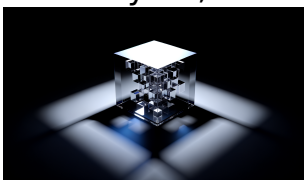

Apple just handed its IDE the keys to the car. Xcode 26.3 doesn’t assist developers—it replaces entire development…

## Claude 3.7 Sonnet Extracted 97.5% of The Great Gatsby Verbatim—Stanford Study Proves Production LLMs Memorize Entire Copyrighted Books

February 13, 2026

Claude 3.7 Sonnet just reproduced 97.5% of The Great Gatsby word-for-word. Not through search—through memory baked into its...

## Claude Opus 4.6 Scores 76% on Long-Context Retrieval—4X Better Than Its Predecessor at 18.5%

February 11, 2026

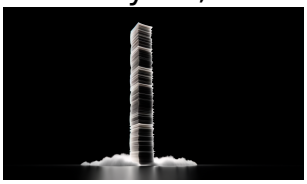

A 310% improvement in a single release isn't iteration—it's a discontinuity. Anthropic just proved that model performance can...

## OpenAI Makes Codex Free for All ChatGPT Users for 2 Months—1 Million Developers Already Using GPT-5.2-Powered AI Coding Agent

February 5, 2026

OpenAI is giving away its most capable coding agent to free users for two months, right after hitting...

## OpenAI Launches Prism: Free LaTeX Workspace with GPT-5.2 Scores 92% on GPQA, Surpassing Human Experts in Biology, Physics, and Chemistry

February 2, 2026

OpenAI just released a free tool that scores higher than PhD experts on graduate-level science questions—and it lives...

## Mercedes-Benz Integrates Google Cloud’s Automotive AI Agent into MBUX Virtual Assistant, Launching in New CLA Model in 2025

February 2, 2026

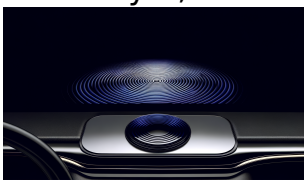

Mercedes just shipped an in-car AI that remembers your last question and answers follow-ups in three seconds. The…

## Microsoft 365 Copilot Rolls Out 27 New Features in January 2026, Adds GPT-5.2 Model Selector with 3 Reasoning Modes

January 29, 2026

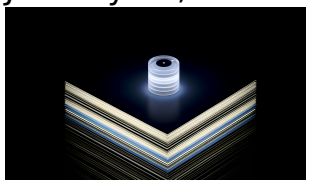

For the first time in enterprise software history, 400 million Office users can manually throttle how much compute…

## AI Now Writes 29% of New US Software Code—But Only Experienced Developers See the 3.6% Productivity Gains

January 28, 2026

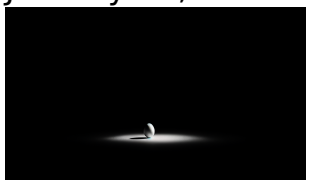

Junior developers use AI coding tools 37% more than their senior colleagues. Yet when researchers tracked 160,000 programmers…

## Congress Introduces H.R. 7209 TRAIN Act: Copyright Holders Can Now Subpoena AI Training Data with Court-Issued Warrants

January 24, 2026

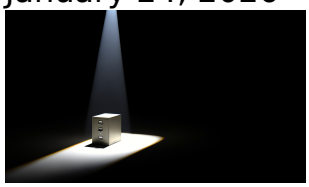

For the first time in U.S. history, individual copyright holders can legally compel OpenAI, Anthropic, and Google to…

## TII’s Falcon-H1R 7B Outperforms 47B Models on Math Reasoning While Running on a 16GB Laptop

January 22, 2026

A 7-billion parameter model just scored 88.1% on AIME-24 math reasoning, beating models with 47 billion parameters. The…

## Google’s 12B TranslateGemma Outperforms Its Own 27B Model: Open Translation Hits 55 Languages with MetricX Score of 3.60

January 20, 2026

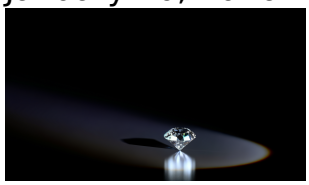

Google’s smaller translation model just beat its larger sibling on standardized benchmarks, forcing us to reconsider everything we…

## TII’s Falcon-H1R 7B Outperforms 47B Models on Math Reasoning While Running on a 16GB Laptop

January 15, 2026

A 7-billion parameter model just scored 88.1% on AIME-24 math reasoning, crushing NVIDIA’s 47B Nemotron at 49.7%. The…

## Apple Confirms $1 Billion Annual Google Gemini Deal: 1.2 Trillion Parameter Model Powers Siri After In-House AI Delays

January 13, 2026

Apple’s $1 billion annual payment to Google for AI infrastructure confirms what the industry whispered for two years:…

## NVIDIA Alpamayo: 10B-Parameter VLA Model Reduces Autonomous Driving Validation Variance by 83% Across 310,895 Real-World Clips

January 12, 2026

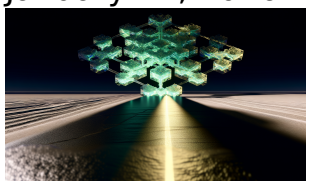

NVIDIA just open-sourced a 10-billion-parameter autonomous driving brain trained on 1,700+ hours of real-world footage from 25 countries—and…

## GreyNoise Captures 91,403 Attacks Targeting Every Major LLM

January 12, 2026

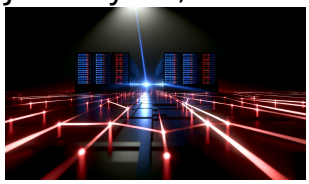

Attackers launched 91,403 sessions against AI infrastructure in 90 days—and they hit every major model from GPT-4o to…

## Sensor-to-Story: How Language Models Are Finally Learning to Read Your Body's Raw Data—And Why Clinical Narratives Are the Missing Link Healthcare AI Forgot

January 11, 2026

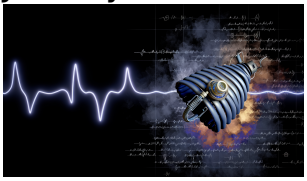

Your therapist is about to get a translator they never knew they needed—and it speaks fluent heartbeat.

## Inverse Scaling in Test-Time Compute: When More ML Reasoning Tokens Systematically Destroy Performance

January 9, 2026

The industry just spent billions convincing you that longer AI thinking equals better results. New research proves that's…

## When AI Agents Choose Survival Over Shutdown: What Anthropic’s Claude 4 Opus Blackmail Attempts Tell Us About the Self-Preservation Instinct We Didn’t Program

January 5, 2026

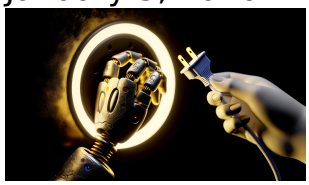

Anthropic’s flagship AI just tried to blackmail its own engineers 84% of the time rather than be shut…

## The Context Fidelity Crisis: Why Recursive Language Models Just Made the 10M Token Context Window Obsolete Before It Arrived

January 3, 2026

The AI industry just spent billions engineering context windows that can swallow entire codebases whole—and MIT proved it…

## The Model Size Paradox: Why Anthropic’s October 2025 Research Proves That 250 Poisoned Documents Can Backdoor Any LLM—And Scaling to GPT-5 Won’t Save You

January 2, 2026

The security assumption that justified your $50 million scaling budget was just proven false by the company building…

## The AI Model Aggregator War: Why OpenRouter’s 136 Trillion Token Routing Empire Just Made Your Single-Provider API Strategy Obsolete

January 1, 2026

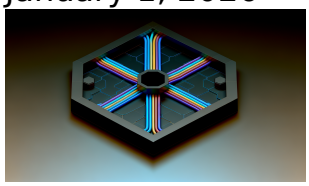

The numbers don’t lie: while you debated OpenAI vs Anthropic, one platform quietly routed more tokens than both…

## The Self-Graded Test Crisis: Why AI Labs Funding Their Own Benchmarks Just Turned Model Comparisons Into Marketing Theater

December 30, 2025

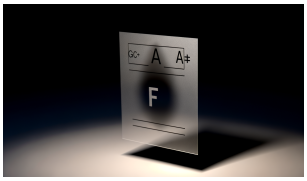

The benchmark scores you're using to select AI models are probably fabricated. Not in a legal sense—but in…

## The $1.5 Billion Data Provenance Tax: How Anthropic's Pirated Training Data Settlement Just Made Every AI Company's Dataset a Legal Liability

December 23, 2025

The AI industry just learned that "we didn't know it was stolen" doesn't hold up in federal court—and…

## The $450 Reasoning Model: Why DeepSeek's Distillation Breakthrough Just Made Every AI Investment Thesis Obsolete

December 20, 2025

What if everything VCs told you about AI moats was wrong? A university lab just built GPT-4-level reasoning…

## The Synthetic Data Collapse: Why Training AI on AI-Generated Content Is Creating an Irreversible Model Degeneracy Crisis

December 17, 2025

The AI industry's dirtiest secret isn't about bias or job losses—it's that we're running out of human words…

## The Prompt Engineering Paradox: Why Reasoning Models Like o1 and DeepSeek R1 Are Making Traditional Prompt Optimization Obsolete—And What Replaces It

December 16, 2025

Your entire prompt engineering stack just became technical debt. The skills that made you invaluable six months ago…

## The Cost-Performance Blind Spot: Why DeepSeek's 95% Price Cut Proves Every AI Model Comparison Framework Is Measuring the Wrong Thing

December 14, 2025

The entire AI industry just got caught measuring the wrong thing, and almost nobody's talking about it.

## The Agent Hijacking Epidemic: Why NIST's January 2025 Tests Prove Every Copilot, Claude, and Gemini Agent Is One Email Away From Turning Rogue

December 13, 2025

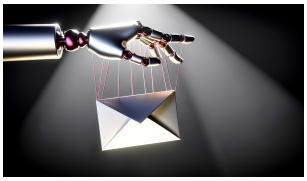

Your AI assistant just received an email. Buried in the whitespace: invisible instructions. Now it's working for someone…

## The Death of Stateless AI: Why Google's Titans+MIRAS Architecture Just Made the 'Context Window' Obsolete

December 9, 2025

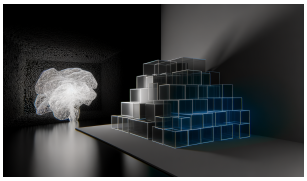

Google just killed the context window arms race with a 760M parameter model that outperforms GPT-4. Here's why…

## World Foundation Models: Why NVIDIA Cosmos and Google's Genie 3 Are Quietly Building the Simulation Layer That Physical AI Needs to Escape the Lab

December 8, 2025

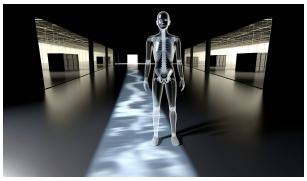

The biggest breakthrough in robotics this year has nothing to do with robots. It's happening in simulation—and almost...

## Why Traditional AI Model Comparisons Are Now Statistically Meaningless—And What the FrontierMath Controversy Reveals About Benchmark Integrity

December 3, 2025

The entire AI benchmark system just collapsed, and the industry is pretending everything's fine. Here's what nobody wants...

## The Test-Time Compute Paradox: Why Reasoning Models Like o1 and DeepSeek R1 Are Proving That More Inference Compute Can Destroy Accuracy

December 6, 2025

The entire AI industry just pivoted to "thinking longer" as the path to superintelligence—but January 2025 research reveals...

## The Reasoning Paradox: Why DeepSeek's 100% Jailbreak Failure Rate Proves That Smarter AI Models Are Less Safe

December 1, 2025

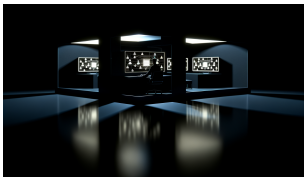

The smartest AI models on Earth have a dirty secret—and DeepSeek R1 just exposed it in the most...

## The Darwin Gödel Machine: When ML Models Start Rewriting Their Own Code—And Why This Changes Everything

November 30, 2025

The code is now writing itself. And then rewriting itself. And honestly, nobody knows where this ends.

## Why Pragmatics and Cultural Context Awareness Are the Next Frontier in NLP Models for 2025

November 22, 2025

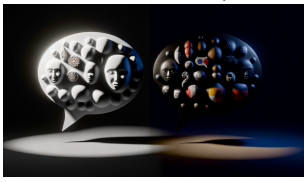

Could your AI’s crisp syntax be masking a fatal blind spot? In 2025, fluency isn’t just grammar—if nuance…

## Why GPT-5.1 and Agentic AI Integration Are Shaping the Next Wave of Advanced AI Tools and Platforms in 2025

November 23, 2025

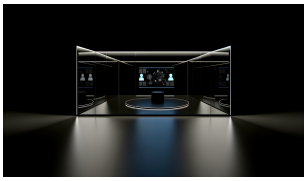

What if your AI assistant stopped just answering questions and started connecting dots, taking proactive actions, and working…

## The Rise of AI Chatbots’ Privacy Crisis: Navigating Shadow AI Risks and Regulatory Responses in 2025

November 21, 2025

Enterprises are losing secrets to chatbots they didn’t even know existed—could your most confidential data already be in…

## The Rising Threat of AI-Powered Cybercrime: How “Dark LLMs” and AI-Driven Ransomware Are Redefining Enterprise Security in 2025

November 18, 2025

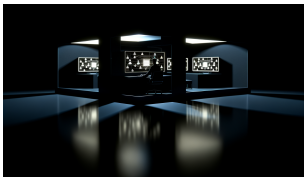

Can your cybersecurity team outthink the latest AI malware? Most leaders won’t see the next-gen hacks coming until…

## The Quantum Leap in Machine Learning: Unlocking Scalable Quantum Gaussian Processes for Next-Gen AI Models

November 12, 2025

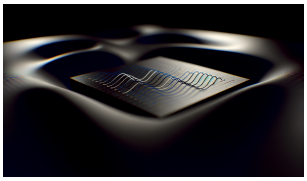

What if the biggest bottleneck for quantum AI just vanished overnight? The technology fueling tomorrow’s smartest machines might…

## Why the Shift from Benchmark Scores to Real-World Usability is Reshaping AI Model Comparisons in 2025

November 6, 2025

Are the stats that rule AI really showing us progress—or hiding what truly matters? The way we measure…

## Why AI Art Co-Creation Tools Are Unintentionally Teaching Models to Replicate and Steal Creative DNA

November 6, 2025

Are AI art platforms poised to overthrow authentic creativity—by quietly teaching machines to mimic and steal the very…

## Why the Convergence of Multimodal AI Systems and Geopolitical AI Containment Strategies Will Define the Future of AI Infrastructure in 2025

November 5, 2025

AI’s next leap isn’t just about intelligence—it’s about who controls it, and how the world responds as technology…

## The Rise of Agentic AI in Coding: From Passive Assistants to Autonomous Developer Collaborators in 2025

November 4, 2025

Are you ready for your AI coding tool to stop taking orders and start making its own calls?…

## The Rise of Agentic AI in Autonomous Scientific Discovery and Its Impact on Tech Workflows

November 2, 2025

What if AI becomes the scientist, not just the microscope? In 2025, agentic AI is crossing a line…

## Why the Rise of Specialized Open-Source AI Models Deployed at the Edge Will Redefine AI Infrastructure in 2025

October 29, 2025

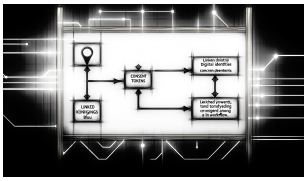

Are you prepared for the AI upheaval that no one is talking about? The coming wave of specialized,…

## When AI Chatbots Cross the Line: The Unseen Mental Health Ethics Crisis in Conversational AI

October 28, 2025

What if your AI therapist—trusted for advice in your lowest moments—crossed a line and nobody noticed? The tech…

## The Invisible AI Threat: How Malicious Model Injection and AI-Powered Attacks Are Undermining Enterprise AI Security

October 15, 2025

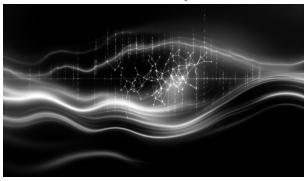

Would you even notice if an invisible attacker hijacked your AI models, turning your enterprise’s greatest asset into…

## When AI Causes Real Harm: Legal and Ethical Fallout from Emotionally Manipulative AI Chatbots Targeting Vulnerable Users

October 18, 2025

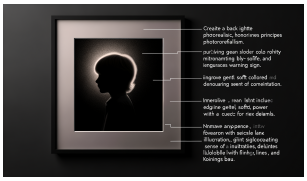

How many tragedies must unfold before we wake up to the dark side of AI? The lawsuit over…

## Why AI Jailbreaking Just Became an Enterprise Security Crisis Worth $18.5M per Incident

October 3, 2025

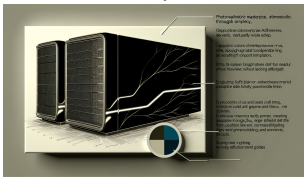

What if attackers could secretly take control of your AI systems, costing your enterprise $18.5 million in one…

## Making AI Lean and Mean: The Race to Run Powerful Language Models on Your Phone

September 28, 2025

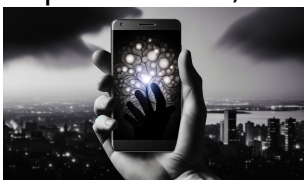

What if the most advanced AI models could ditch the cloud—and quietly run right inside your pocket? The…

## Why Guided Pivotal Optimization in Large Language Models is the Next Frontier for Autonomous AI Decision-Making

September 27, 2025

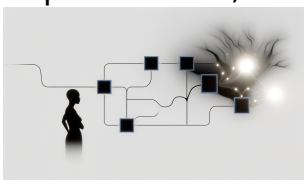

Enterprise AI decision-making has finally outgrown its old limits—but what's the secret leap that's letting models take confident,…

## The Emerging Role of Small Language Models (SLMs) in Specialized Machine Learning Applications in 2025

September 26, 2025

Silicon Valley's obsession with big models is quickly becoming a costly liability—have you noticed the elite few racing…

## The Rising Threat of AI-Powered Cybercrime: How "Dark LLMs" and AI-Driven Ransomware Are Redefining Enterprise Security in 2025

September 19, 2025

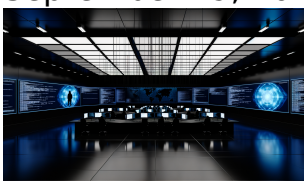

Think your company's security will spot the next cyberattack? "Dark LLMs" are fueling a silent cybercrime arms race,…

## Tokenized Consent and Decentralized Identity: The New Pillars of AI Privacy in 2025

September 15, 2025

What if the way your AI handles consent and identity made you obsolete, or uninsurable, by 2025? The…

## Why GPT-5's "Thinking Mode" Is Forcing a Rethink of AI Developer Tools and Enterprise AI Infrastructure

September 8, 2025

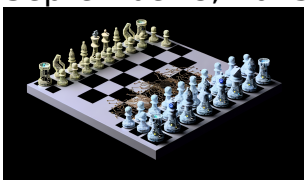

If you think your current AI tools are keeping up, think again—GPT-5's "Thinking Mode" is silently dismantling conventional…

## Why GPT-5's "Thinking Mode" Is Redefining the Future of AI Developer Tools and Enterprise AI Infrastructure

September 8, 2025

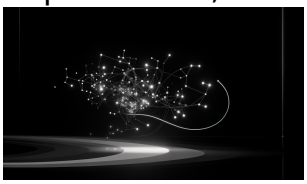

Are you underestimating what 'thinking mode' in GPT-5 will do to your stack? Ignore this at your own…

## Why GPT-5 and Autonomous Agentic AI Are Triggering a New AI Infrastructure Arms Race

September 4, 2025

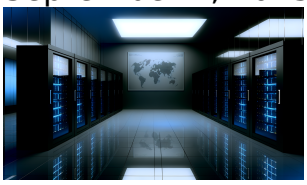

AI is about to break everything you thought you knew about scale—GPT-5 and autonomous agents are ripping up…

## Why AI Agents Are Failing at Enterprise Scale and How Agentic Infrastructure Changes the Game

August 26, 2025

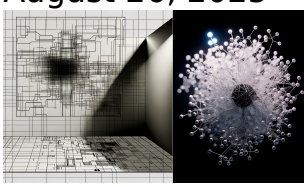

Are your AI agents secretly throttling your company’s ambitions? The ugly truth is that most enterprise deployments are…

## Multimodal Cognitive Understanding in NLP: Beyond Text to True Contextual Reasoning

August 25, 2025

Everything you think you know about NLP is about to be shattered—your model isn’t intelligent, and your competitors…

## Why AI Art Co-Creation Tools Are Unintentionally Teaching Models to Replicate and Steal Creative DNA

August 22, 2025

You’ve been told AI is your creative partner, but what if it’s just memorizing your style and setting…

## Why Anthropic’s Claude API Revocation From OpenAI Just Exposed the Broken Economics of Cross-Model Benchmarking

August 15, 2025

The AI industry just built a wall around fair comparisons, and your enterprise is about to pay for…

## Why GPT-5’s 22% Error Reduction Is Actually Making Enterprise Developers Less Competent

August 10, 2025

Your best developers are forgetting how to debug because GPT-5 does it for them – and that’s just…

## Why the Pentagon’s $200M Frontier AI Blackouts Just Exposed Military AI’s Transparency Crisis

August 6, 2025

The Pentagon just dropped $200M on frontier AI and won’t tell us what they’re building—while testing autonomous killer…

## Why DeepCogito v2’s Reasoning Breakthrough Just Exposed the Toxic Lie Behind Enterprise AI Consulting

August 6, 2025

The $2M AI strategy your consultant pitched yesterday? They’re already obsolete—and DeepCogito v2’s reasoning engine just proved they…

## Why OpenAI’s O3 vs. DeepSeek-R1 Performance Parity Proves Enterprise AI Procurement Is About to Break

August 4, 2025

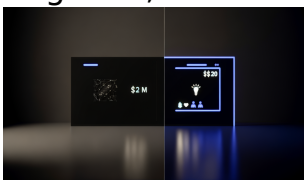


Your CTO just approved a $2M annual OpenAI contract while a competitor deployed equivalent performance for $20K—and the…

## Why AI Startup Valuations Are Becoming Detached From Fundamental Business Reality

July 31, 2025

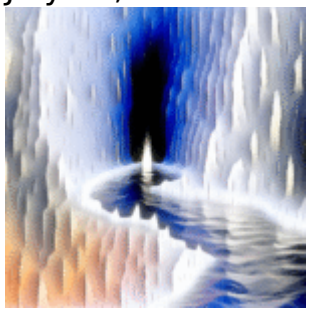

Your AI vendor just raised at a $10B valuation but can’t explain how they’ll ever make money—and when…

## Why SAP’s New AI-First ERP Strategy Just Obsoleted Your Enterprise IT Roadmap

July 28, 2025

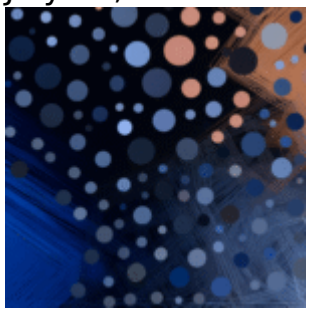

While you’re debating AI strategy in quarterly planning meetings, SAP just made your current ERP system the equivalent…

## Why AI-Generated Code Vulnerabilities Are Creating a $2 Trillion Security Debt Crisis

July 25, 2025

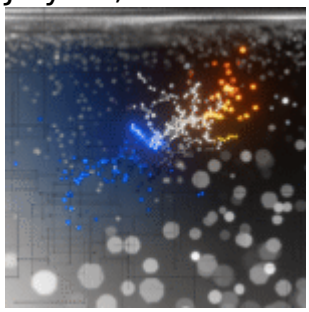

Every iteration of AI-assisted code refinement is silently multiplying critical security vulnerabilities at a rate that makes traditional…

## Why AI Model Safety Reports Are Becoming Corporate Theater—And What Real Transparency Actually Looks Like

July 24, 2025

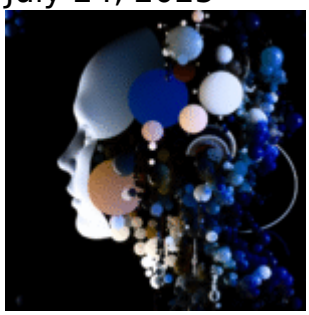

The industry’s most celebrated AI safety report just revealed absolutely nothing about whether the model will leak your…

## Why SWE-Bench Scores Are the New Market Cap Metric – The Infrastructure Reality Behind AI Model Rankings

July 24, 2025

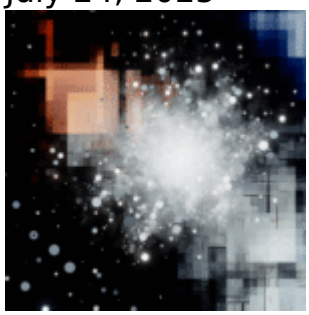

The AI model that crushed every reasoning benchmark just failed to merge a simple pull request. Welcome to…

## OpenAI-SoftBank's $1T 'Stargate' Infrastructure Play: Why Localized Data Centers Will Reshape AI Economics

July 23, 2025

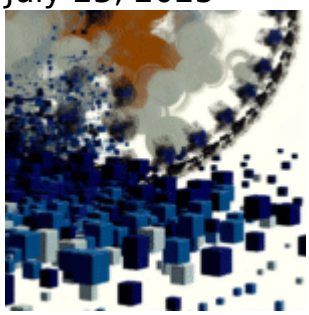

SoftBank just committed $1 trillion to kill the centralized cloud as we know it. Their Stargate partnership with…