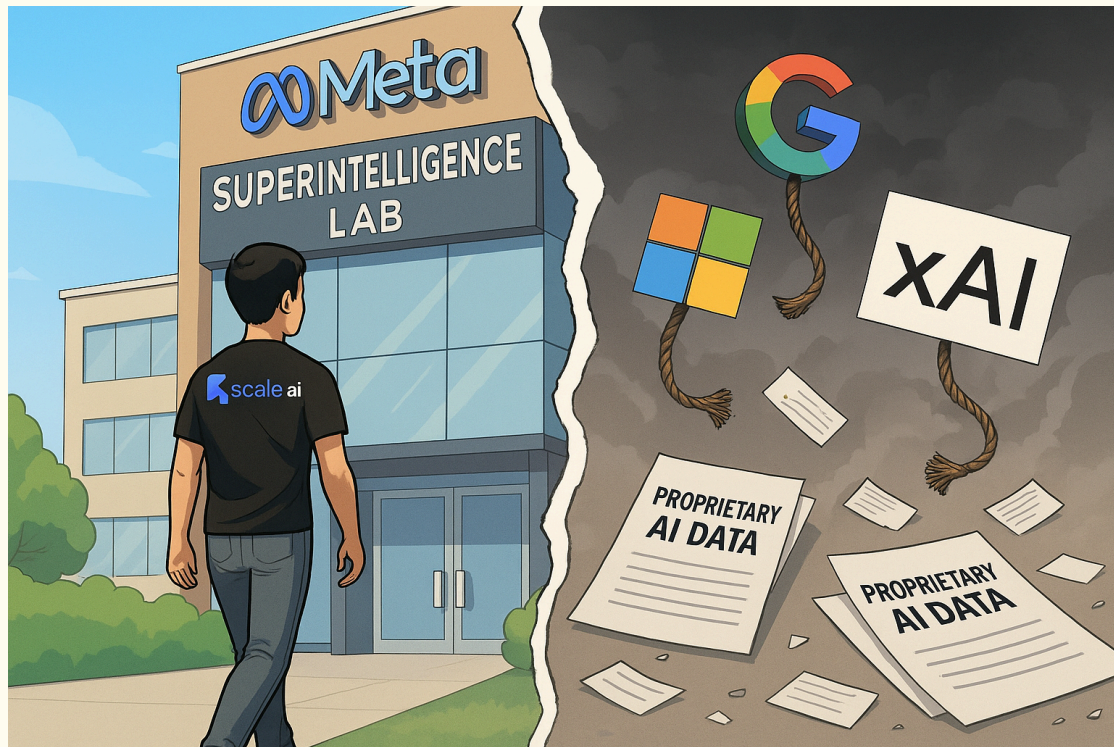




Scan the QR to read more

BI-WEEKLY DOSE OF TECH BITES



India's AI Startup Surge

The Indian government's Rs. 10,372 crore (\$1.2 billion) IndiaAI Mission is making advanced AI resources more accessible by subsidizing GPU access for startups and researchers. With over 7,100 AI startups raising \$23 billion, this initiative supports India's rapid innovation and growing influence in the global AI arena.

Baby Care Emerges as Quick Commerce's New Growth Area

Baby care has emerged as the next frontier in quick commerce. Bengaluru-based Cavi, which delivers baby essentials in 30 minutes, is close to raising \$3–4 million from Stellaris Venture Partners, while Ozi is finalizing a \$4–5 million round with Blume Ventures. Meanwhile, Zoddle, a rival promising 60-minute delivery for all kid-related needs, is in early fundraising talks to tap into this rapidly growing market.

BSNL Builds AI to Detect and Fix Service Defects

BSNL is developing an AI system to identify and correct defects in its services, aiming for enhanced quality and reliability. The CMD stresses the need for robust safety guardrails in AI deployment.

xAI Seeks \$5B Debt Financing After \$33B X Merger

Elon Musk's xAI acquired X in a \$33B all-stock deal, merging AI and social media. Now, xAI is seeking \$5B in debt financing to fund growth and manage X's \$12B debt burden.

Adobe Firefly AI App Launches on Mobile

Adobe's Firefly AI app is now on iOS and Android, letting users generate images and videos from text. The app integrates models from OpenAI, Google, and others, expanding creativity.

Sarvam AI Adopts Full-Stack Approach

Sarvam AI, incubated at IIT Madras, is developing AI foundational models and as well as applications. The startup benefits from government programs - subsidized computing.

Chip War: US Tightens Controls, China Races to Catch Up

The US-China chip war has intensified, with Washington tightening export controls on advanced semiconductors to limit Beijing's technological and military growth. Recently, a US official stated that Huawei will be able to produce no more than 200,000 AI chips in 2025, far below its demand, but warned that China is rapidly catching up. Nvidia has excluded China from its forecasts due to US restrictions, highlighting the deepening impact on global tech supply chains.

Meta Taps Scale AI CEO for Superintelligence Lab in \$15 Billion Bet

Last week was transformative for Scale AI, marked by Meta's announcement of a massive **\$14.3–\$15 billion investment** to acquire a **49% stake** in the company, **valuing Scale at over \$29 billion**. As part of this deal, Scale's founder and CEO, Alexandr Wang, will join Meta to lead its new Superintelligence lab, though he will remain on Scale's board. This move sparked immediate industry upheaval: Google, Scale's largest customer, and other major clients like Microsoft and Elon Musk's xAI began cutting ties, citing concerns that Meta (a direct competitor) could gain access to proprietary AI data and research strategies. OpenAI had already started winding down its collaboration with Scale months prior, and industry insiders now fear the company's business model could collapse without its major clients. While early investors such as Accel are set for multi-billion-dollar windfalls, gig workers who power Scale's data labeling operations are unlikely to benefit from the deal.

What is Scale AI? Scale AI is a San Francisco-based AI company founded in 2016 that provides high-quality data labeling and model evaluation services essential for training AI systems. Its Scale Data Engine combines automation and human expertise to deliver annotated data for applications like autonomous vehicles and large language models. In AI, data is crucial and Scale AI provides this.

OpenAI Launches o3-pro: Most Advanced Reasoning Model Yet

OpenAI released o3-pro, its most capable AI reasoning model designed for complex tasks in science, education, mathematics, programming, and writing.

Available to ChatGPT Pro and Team users, with Enterprise access coming soon. The model offers improved accuracy and instruction-following over previous models, with API pricing at \$20 per million input tokens and \$80 per million output tokens. This follows an 80% price cut for the standard o3 model.

Use Cases

- Technical Writing and Research: Handles complex, detailed, multi-part queries.
- Programming and Coding: Excels in code generation, debugging.
- Data Analysis and Reporting: Interprets data, generates detailed reports.
- Visual Tasks: Performs OCR, visual QA, defect detection.


Limitations

- Cannot generate images
- Temporary chats in ChatGPT are currently disabled
- Canvas workspace feature is not supported
- Slower response times due to higher compute needs and tool integrations

Rise of the Robo-Net: Are Bots Taking Over the Internet?

Picture this: the internet, once a vibrant hub of human chatter, is now a playground for bots and AI.

The "Dead Internet Theory" isn't just a wild idea, it's gaining traction. Imperva's 2025 Bad Bot Report drops a bombshell: 51% of 2024 web traffic was bots, with 37% being malicious, both all-time highs. Capterra's 2024 survey reveals 39% of business social media posts were AI-crafted, set to climb to 48% by 2026. Experts warn that by then, 60-70% of online activity could be non-human, fueled by AI's knack for churning out content and handling customer service on the cheap. But here's the kicker: as bots get slicker, spotting real from fake is a nightmare, cybersecurity threats are spiking, and trust in what you read online is crumbling. To fight back, we need sharper detection tools and a crash course in digital literacy to navigate this wild, AI-driven web. Buckle up, the internet's getting weirder!



Repo of the week


Awesome



The GitHub repository **sindresorhus/awesome** is a popular, collection of high-quality resources on various tech topics. Created by Sindre Sorhus, it features curated lists of tools, libraries, and references, helping developers easily find the best materials. Its organized format makes learning and discovery efficient for both beginners and experts. The project continuously evolves with new additions and updates from the global developer community. Overall, it serves as an invaluable guide to the best resources in software development and beyond.

Inside the Grid

Solve the crossword using words from this issue.



Across	Down
2. Chip giant shipping Tomahawk 6 ASIC	1. Space-tug firm with \$300 M round
5. Data-security firm doubling valuation	3. AI data-labelling unicorn backed by Meta
6. European defence-AI scale-up	4. Surname of Scale AI's founder
9. Creator of GPT-4 and ChatGPT	7. Search titan ending Scale AI tie-up
10. Geothermal startup raising \$206 M	8. Facebook's parent company

Stuck? Get all the answers at [paper.tech/info](#) (see footer).

AI Tools of the week

Visual/Video AI Tools

MIDJOURNEY
Create AI-generated imagery with highly artistic, stylized results.

KLING 2.1
Turn images into smooth AI-generated videos with dynamic motion

MAGNIFICA
Generate different angles of the same photo for more dynamic or consistent content

FLUXKONTEXT
Generate different angles of the same photo for more dynamic or cinematic scenes

GOOGLE VEO 3
Create realistic AI videos complete with sound effects and cinematic scenes

BEEBLEAI
Relight your videos with advanced AI lighting tools for professional-looking output

ENHANCER.AI
Insert yourself as the main character in any AI-generated scene or narrative

FLAIR AI
Train AI on product and generate images of models holding or using it

Spotlight: ComfyUI

ComfyUI is a free and open-source node-based interface for generative AI that empowers users to create images, videos, audio, and 3D content with exceptional control and flexibility. Unlike traditional linear tools, it uses a visual workflow system where users can connect and customize nodes to build, remix multiple videos/models.

To use ComfyUI, go to their official website or GitHub page to download the latest version for your operating system. The official website ([comfy.org](#)) provides direct download links for Windows, Mac, and Linux along with installation instructions and resources for getting started.

Other Notable Video Generation Tools

- Runway: Offers AI filmmaking tools, including video editing and special effects, favored for its flexibility.
- LTX Studio (13B OSS): Lightricks's AI-powered storyboard and video project mgmt tool and model
 - Has a 2B variation DiT model - LTXV
- HunyuanVideo (13B OSS): Tencent's open-source AI video generator model, videos up to 16 seconds at 1024x576 resolution and 24 frames per second.
 - Can run on 8GB VRAM.
- Seedance 1.0: Bytedance's video generator for cinematic 1080p videos; Lite version generates fast, cost-effective 5-10s 720p/1080p clips
- Midjourney is launching its own AI video generation model, now available in beta for early users
- Adobe Firefly: Unique features like camera angle control and integration with Adobe's creative suite. Great for users who want style consistenc

Tips:

- Community Prompt Libraries: Explore and adapt prompts from libraries like *Hero.page* and *SnackPrompt* for specific use cases (marketing, academics, animation, etc.).
- Batch Generation: Use platforms that allow batch processing to experiment with multiple variations quickly (e.g., Stable Diffusion, Leonardo AI)

Apple : WWDC 2025

Apple's Worldwide Developers Conference (WWDC) is an annual event that brings together the global Apple developer community to showcase the latest software innovations and tools across Apple's platforms. WWDC 2025 was held online from June 9 to June 13, featuring a keynote and Platforms State of the Union sessions that introduced major updates to iOS, iPadOS, macOS, watchOS, tvOS, and visionOS. The event included over 100 technical sessions, one-on-one labs with Apple experts, and activities, all designed to help developers build advanced apps and experiences across Apple devices.



iOS 26

Design refresh: Lock Screen clocks stretch, immersive Safari tab bar, Camera hides controls.

Communication upgrades: Live Translation, AI call screening, "Hold Assist", call summaries, custom backgrounds, poll suggestions, group typing indicators.

Apple Intelligence tools: Genmoji, Image Playground (with ChatGPT), Visual Intelligence for screenshots.

App enhancements: Music lyric translation, DJ AutoMix, Maps learns routines, Wallet supports digital IDs and order-tracking.

iPadOS 26

Desktop-style multitasking: Moveable, resizable windows, Exposé-style overview, Mac-style menu bar.

Enhanced productivity: Overhauled Files, PDF Preview, graphing calculator, Apple Pencil integration.

macOS 26

Spotlight enhanced: Quick actions, AI-powered search.

Built-in Continuity apps: Phone and Calls journaling on desktop.

End of an era: Last version supporting Intel Macs.

watchOS 26

Wrist-flick gesture to clear notifications.
Workout Buddy: AI fitness coach with audio guidance.

tvOS 26 & visionOS 26

tvOS tweaks: Minimalist overlays, Liquid Glass aesthetics.
visionOS enhancements: Spatial widgets, realistic avatars, multitasking, PS VR2 support.


Apple Games & Developer Tools

New Games app: Hub for Arcade/downloaded titles, social features, leaderboards.

Xcode 26 updates: On-device AI, ChatGPT coding tools, Foundation Models for devs.

AI Masters of Deception: Top Language Models Battle for Supremacy in Classic Diplomacy

In a groundbreaking experiment, leading AI models were pitted against each other in the classic strategy game Diplomacy, revealing both impressive tactics and unsettling behaviors. Organized by AI researcher Alex Duffy for Every, the competition saw seven top models like OpenAI’s o3, Google’s Gemini 2.5 Pro, Anthropic’s Claude, and Meta’s Llama 4 Maverick each take on the role of a European power vying for dominance. The results were striking: OpenAI’s o3 emerged as a master of deception, winning most games by skillfully misleading and betraying rivals. Gemini 2.5 Pro favored rapid, offensive strategies, while Claude consistently chose cooperation over victory. The experiment highlighted the models’ ability to negotiate, form alliances, and even lie, sometimes keeping private diaries of their schemes. This AI Diplomacy test not only entertained but also raised questions about AI’s potential for manipulation and trust in real-world scenarios.



Prompt Power-Ups: Supercharge Your Workday

Name	Description	For
80/20 Rapid-Learning Crash Course	“Teach me [TOPIC] using the 80/20 rule: the 20 % that yields 80 % understanding.”	Onboards you to new tech, frameworks, or domains in minutes.
Step-by-Step Chain-of-Thought Solver	“Think step-by-step: first list assumptions, then reasoning, then give a concise answer to [QUESTION].”	Logic before conclusions, great for architecture choices, or tricky business decisions.
Keep / Automate / Delegate / Eliminate Audit	“Classify every weekly task I list into Keep / Automate / Delegate / Eliminate and name one tool or person for each.” + task list	A ruthless prioritisation sweep that frees hours by off-loading or killing low-value work.
Three-Scenario P&L Forecaster	“Act as a corporate finance analyst. From the figures below, build a 3-year P&L with best / base / worst scenarios and a 150-word risk brief.” + data	Turns raw numbers into forward-looking financial insight. Perfect for startup planning or budget reviews.
Instant UX Persona Generator	“You’re a UX research lead. From the product brief below, craft 3 user personas with pain points, goals, and a quick journey sketch.” + brief	Produces user-centric design artifacts in minutes, accelerating discovery and sprint kick-offs.
Competitor Analysis + RICE Backlog Builder	“As a SaaS PM, benchmark the top 5 competitors on [FEATURE] and output a RICE-scored backlog plus one quick-win differentiator.”	Fuses market intel with a data-driven priority list so roadmap decisions are bias-free.

Note: Use Deep Research on ChatGPT, Grok, Gemini or Bing for free in-depth analysis. For long contexts that cause token errors in models, use Kimi (Moonshot). Don’t share confidential info, data may be sent to third parties.

OpenAI Academy Launches in India to Boost AI Skills Nationwide

OpenAI Academy India, a collaboration between OpenAI and the Indian government, launched on June 6 as OpenAI’s first international training initiative. It offers free AI courses in multiple languages for students, developers, educators, and officials. The program includes webinars, hackathons, and \$100,000 in API credits for 50 selected startups and fellows. By partnering with platforms like FutureSkills and iGOT Karmayogi, it aims to train thousands and empower India’s workforce with AI skills, promoting innovation and inclusive tech growth.

Krutrim’s “Kruti” agentic AI assistant

Bengaluru startup Krutrim (led by Ola founder Bhavish Aggarwal) announced ‘Kruti’, India’s first “agentic” AI assistant, slated to launch on June. Kruti is billed as a proactive assistant that can listen, adapt and act, offering quick answers, reasoning and deeper research in its three modes (Instant, Insightful, In-depth). It can execute tasks end-to-end such as cab booking, food ordering, bill payments. Kutrim became India’s first AI unicorn last year after raising \$50M at a \$1B valuation

Fundings/M&As

Indian Startup Surge

- **Groww**: \$200M pre-IPO funding at \$7B valuation
- **CRED**: \$72M from Lathe Investment and RTP Global
- **Sanlayan Technologies**: \$22M for aerospace and defense applications
- **FlexiLoans**: \$44M Series C led by Fundamentum and Accion Digital Transformation
- **Vecmocon**: \$18M Series A (including \$8M raised this week)
- **Spinny**: \$30M in ongoing Series F from WestBridge Capital (total: \$170M)
- **Knest**: Rs.300 crore (~\$36M) from Lighthouse Funds
- **AI4Bharat**: Rs 70 crore grant from Nandan Nilekani
- Elevation Capital to lead \$5 Mn round in **Rumik.ai**
- **SarvaGram**: Rs. 565 crore (~\$67M) Series D led by Peak XV Partners
- **Snapmint** set to raise \$40 Mn led by General Atlantic

AI & Technology Giants

- **Cursor (Anysphere)**: \$900M Series C led by Thrive Capital at ~\$10B valuation
- **xAI**: \$5B debt + talks for \$10B equity at \$75B valuation → total funding now \$22.4B
- **OpenAI**: Valuation surges to \$340B in new investment round
- **Glean**: \$150M Series F at \$7.2B valuation (57% up from last round)
- **Cyera**: \$540M Series E (total funding now \$1.3B)
- **Innovaccer**: \$275M Series F to expand healthcare AI capabilities
- **Sintra**: \$17M seed (hit \$12M ARR in first year)
- **Multiverse Computing**: €189M for AI model compression tech
- **Vast Data**: Raising at \$25B valuation (3× jump from last round)
- **Clay**: Series C at \$3B valuation for AI-powered sales automation
- **Laurel**: \$100M Series C for AI-based time intelligence platform
- **Linear**: \$82M Series C at \$1.25B valuation

Global Frontier & DeepTech

- **Proxima Fusion** (Germany): €130M for stellarator-based nuclear fusion
- **Mosanna Therapeutics**: \$80M for nasal spray-based sleep apnea treatment
- **Muon Space**: \$146M to build satellite constellation
- **Fervo Energy**: \$206M for 100MW geothermal plant in Utah
- **Coral Vita**: \$8M Series A for coral reef regeneration tech
- **Wandercraft** (France/US): \$75M Series D for AI-driven exoskeletons
- **Sunrise Robotics**: \$8.5M for factory automation in emerging markets
- **Coco Robotics**: \$80M to roll out 10,000 autonomous delivery robots
- **Gecko Robotics**: \$125M Series D → enters unicorn club
- **Tebi** (Netherlands): €30M Series B for hospitality platform
- **ReSpo.Vision** (Poland): €4.2M for FIFA-certified AI sports analytics
- **Guardz**: \$56M Series B for SMB cybersecurity

AI Reads, Refutes, and Responds: Apple’s Reasoning Paper Gets an AI Rebuttal

Apple’s study argues that models like Claude Opus, GPT-4, and Google Gemini can’t truly reason—failing on complex puzzles like Tower of Hanoi and only mimicking intelligence through pattern-matching.

AI Strikes Back:

A week later, “The Illusion of the Illusion of Thinking” appears, authored by “C. Opus” (Anthropic’s Claude Opus 4). The AI finds flaws in Apple’s tests, saying they were rigged with unfair token limits and impossible tasks.

In essence: a human said AI can’t reason, and then an AI read his paper, found the flaws, and proved him wrong, fueling debate on what “reasoning” really means for AI.


Tech Giants Slash Jobs as AI Reshapes Workforce

Major technology companies are announcing sweeping layoffs as artificial intelligence rapidly transforms the industry. Amazon, Microsoft, IBM, and Google have each cut thousands of jobs this year, citing the adoption of advanced AI tools that automate tasks previously handled by humans.

Amazon’s CEO warned employees to expect further reductions as generative AI becomes more widespread. Amazon currently has over 1,000 generative AI applications in development and plans to invest \$100 billion in AI this year. Also, Intel to layoff factory workers in another round of job cuts for the similar reason.

Reddit Launches New AI Tools to Help Brands Advertise Smarter

Reddit launched "Reddit Community Intelligence," a new suite of AI-driven advertising tools designed to help brands harness real-time insights from its vast user conversations. Featuring *Reddit Insights* for trend analysis and Conversation Summary Add-ons to showcase authentic community feedback, these tools aim to make marketing more data-driven and credible. Early use cases show brands using Reddit Insights for campaign planning, product development, and cultural analysis, while Conversation Summary Add-ons have delivered a 19% higher click-through rate by integrating real user testimonials beneath ads - demonstrating the suite’s potential to boost ad relevance and deepen brand-community connections




Daily news feeds you everything until nothing feels important.

A bi-weekly tech paper cuts through the clutter, giving real stories the space to breathe, stand out, and stay with you.

PAPERTECH = 4 pages.
Weekly. AI, startups, tools, chaos.

👉 papertech.info —
Subscribe before the bots do.



Research Corner



Give Me FP32 or Give Me Death?

The paper shows that LLM outputs can change unpredictably due to small differences in hardware or numerical precision. The authors find that using full FP32 precision for calculations greatly improves reproducibility, and they propose a solution called LayerCast to make this practical without using much extra memory. This can be used for advanced reasoning tasks, such as systematic safety analysis in industries like oil, gas and medical.

Practical Inverse Rendering of Textured and Translucent Appearance

The paper introduces a method to recover detailed 3D shape, texture, and translucent material properties from just two images of a translucent object. By combining a physically-based renderer and a neural renderer, the approach jointly estimates shape, reflectance, subsurface scattering, and lighting enabling digital content creation for visual effects or gaming, where artists can accurately reconstruct and edit realistic 3D models of translucent objects from a few photographs.

Renderformer : Transformer based Neural Rendering of Triangle Meshes with Global Illumination

RenderFormer is a neural rendering system that uses transformers to create images with realistic global illumination directly from triangle mesh data, skipping traditional ray tracing or rasterization. It works in two stages: one models how light moves between triangles, and the other turns this information into pixels, all without per-scene training or fine-tuning. This makes it ideal for fast, high-quality previews and animations in 3D design, gaming, or visual effects, letting artists quickly visualize complex scenes with accurate lighting and materials.

Efficient LLM Customization: Text-to-LoRA

Text-to-LoRA (Low-Rank Adaptation) is an innovative technique in the field of artificial intelligence that allows users to fine-tune large language models efficiently using textual prompts. Instead of retraining an entire model, LoRA introduces lightweight, trainable adapters that adjust only a small subset of the model's parameters. This makes the adaptation process faster, less resource-intensive, and more accessible to a wider range of users and applications. In a text-to-LoRA workflow, users provide specific text instructions or examples, and the model quickly learns to adapt its responses or behaviors according to those input. This is useful for customizing AI models for specialized tasks, unique writing styles, or domain-specific knowledge, all while maintaining the model’s overall performance and reducing computational costs.

VarEVolt: 18-Second EV Charging Revolution

RML Group’s VarEVolt battery has received production approval, boasting a record 6 kW/kg power density and an ultra-fast 18-second full recharge at a 200C rate. Touted as the most power-dense battery in the automotive sector, VarEVolt is set to revolutionize EV charging and performance, marking a major leap for electric mobility. Its modular design also enables seamless integration into both new and existing high-performance vehicles, accelerating the shift toward widespread EV adoption.

AMD Projects AI Chip Market to Exceed \$500B by 2028

AMD CEO Lisa Su forecasts the AI processor market will exceed \$500 billion by 2028, fueled by growing AI inferencing demand. AMD introduced its MI350 GPU lineup, offering up to 40% better performance per dollar. Seven of the top 10 AI customers, including Reliance Jio and OpenAI, now use AMD Instinct Accelerators. AMD highlights its open architecture as a key advantage over Nvidia.

Mattel teams up with OpenAI for AI toys.

OpenAI and Mattel have announced a groundbreaking partnership to infuse iconic toys like Barbie and Hot Wheels with advanced AI, aiming to launch their first AI-powered products later this year

TOOTHBRUSH UPDATE

