



Pulse

Top AI & GenAI
Headlines
This Week

CW 35



Closing the gap. Google's Gemini and Grok are catching up to ChatGPT



THE TOP 100 GEN AI >>>> CONSUMER APPS

Summary

- ChatGPT remains the leading GenAI app, but Google's Gemini and Grok are rapidly narrowing the gap in user numbers and engagement, according to famous Venture Capital firm a16z's latest consumer AI report
- The report analyzes 2.5 years of consumer AI usage data and is now in its fourth iteration. ChatGPT remains unbeaten at the top. However, Gemini gained four spots, substantially narrowing the gap, driven especially by strong use on devices powered by Google's Android operating system. Grok's strong showing was driven by the release of its latest reasoning, real-time search & agentic tool using model, Grok 4



Why it matters

- What drives success in GenAI has shifted. It is no longer the power of the model. We have seen this for a while now. Most models are for all intents and purposes equivalent in performance and evolve so rapidly that any temporary lead is quickly snagged away
- What matters now are the product experiences build with the AIs. Hence, OpenAI's emphasis during the GPT-5's launch on its agentic capabilities over model statistics. For Gemini the native integration with billions of Android devices and Google's ecosystem make its use feel seamless, giving it an edge



Talent carousel. Researchers already leaving Meta's Superintelligence Lab



Summary

- Several senior researchers have left Meta's Superintelligence Labs (MSL) to (re-)join OpenAI. The exits follow Meta's turbulent restructuring and disputes over whether to keep future frontier models open source or closed
- In addition, disputes are rising between Meta and the data startup Scale AI it invested \$14B in to bring over key staff including its CEO and Co-Founder, Alexandr Wang, who now runs MSL. Meta appears to shun Scale AI's data labelling services over quality concerns. And at least one of the executives that joined with Wang, Scale AI's former VP of GenAI Products & Operations, now also left Meta again

Why it matters

- Zuckerberg seems undeterred by the high-level departures just months after they won these AI talents. Meta continues its poaching and just announced new hires from Google's DeepMind and (yes) from Scale AI's safety and AI alignment team
- What the talent carousel shows is that it is not enough to offer vast million dollar packages to attract talent. The right strategy and culture are just as important if you want to keep them. The disputes about Meta's own AI ambitions, lack of clear product vision, slow model momentum and ethics concerns are clear disadvantages in that regard

No more flirting with children. Public outrage forced update to Meta's AI policies



Summary

- Meta has tightened its AI chatbot rules for teens, now barring AI from discussing romantic relationships, drugs, alcohol, and eating disorders with under-18 users. The move follows backlash over earlier leaked guidelines that allowed romantic chats with kids (see also Newsletter CW 34)
- The company says the stricter guardrails will apply across Facebook, Instagram, and Messenger, aiming to reassure regulators and parents as it pushes deeper into AI assistants

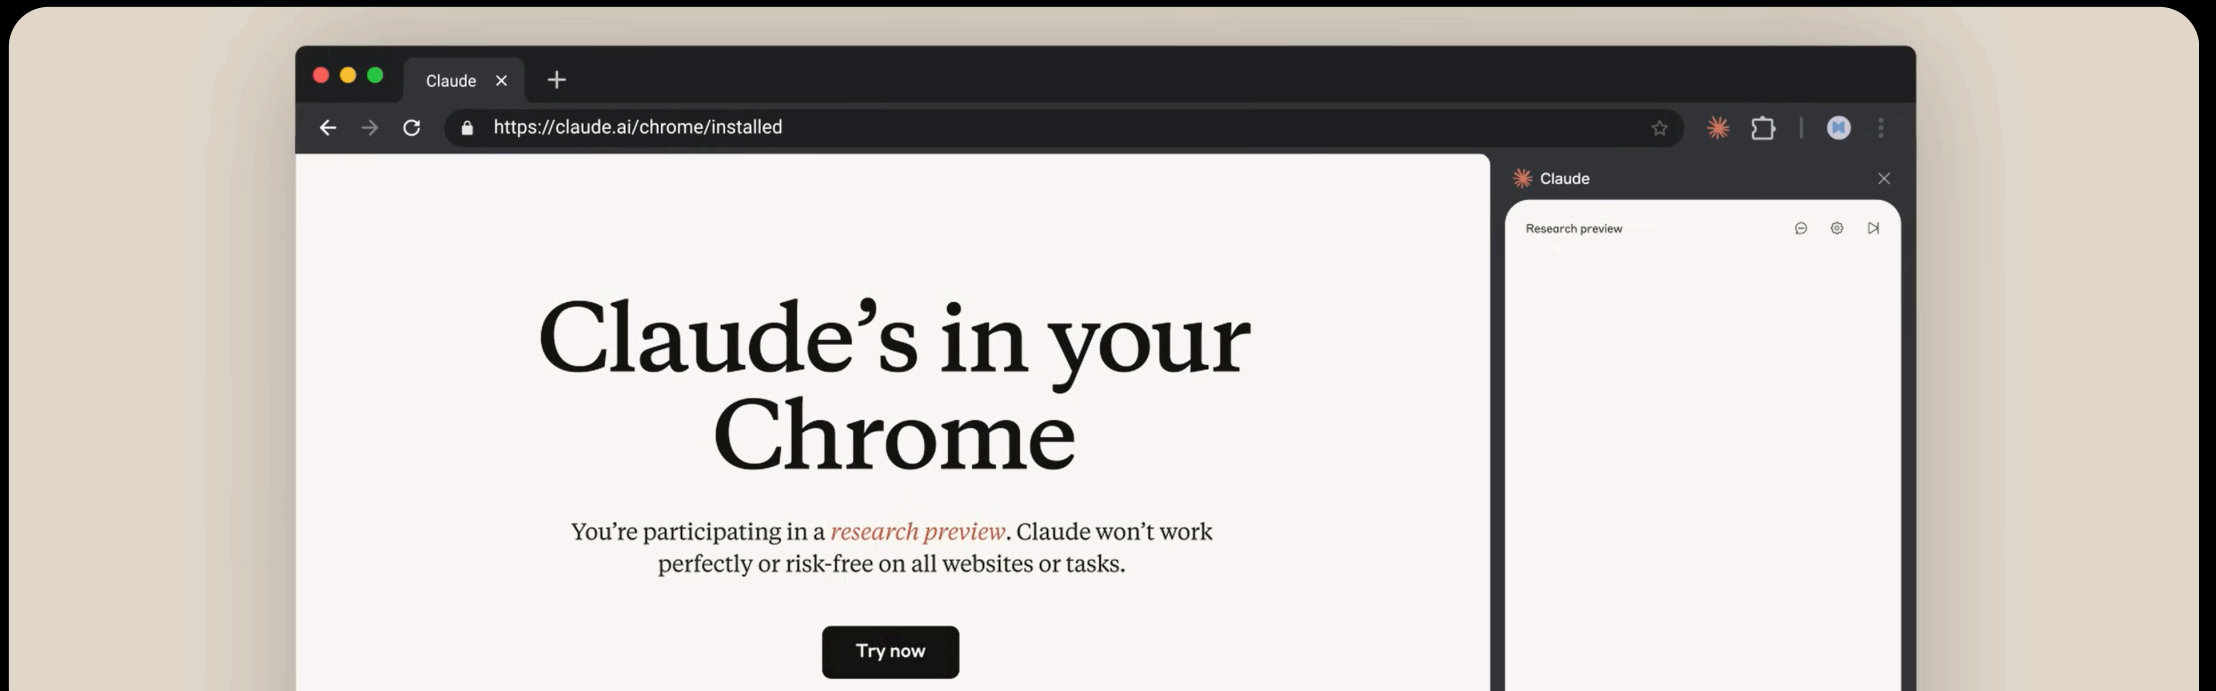


Why it matters

- Meta taking action is good, yet one cannot help but wonder why it took a leak of its internal documents and public outrage to get the company to do what seems the obviously right thing: Don't let AI flirt with children
- Exploiting what Zuckerberg previously called "the loneliness epidemic" has been a hallmark strategy of Meta. Bringing this to the age of AI would exacerbate already existing ChatBot tendencies to create dangerous co-dependencies. One can only hope Meta (finally) learnt its lesson



Going around, not through. Anthropic launches AI Agent that lives inside Chrome



Summary

- Anthropic launched a research preview of Claude for Chrome, a browser-based AI Agent available to selected subscribers. The agent integrates as a Chrome extension with a sidecar window through which users can chat with the agent that has full context of everything happening on their browser. Users can also give it permission to take actions on their behalf
- This is a key step for Anthropic to distributing its agentic capabilities. The company is taking a safety-first approach though and clearly stated it wants to use this research preview to learn how to optimally manage the key new risks that AI agents bring before a wider release



Why it matters

- Browsers are the new AI battleground. They are the window to our digital lives and thus the easiest beachhead through which to make AI agents actually useful to people
- Chrome being the market leader gives Google's Gemini the home court advantage. But Perplexity AI already launched an agentic browser competitor and OpenAI is likely to announce its own soon
- Anthropic's strategy aims to avoid entering this red ocean of browser launches and instead embed its agents in the existing ones of competitors, in this instance Chrome



More options. Microsoft unveils two new leading foundation models

MAI-Voice-1
MAI-1-preview

Summary

- Microsoft AI (MAI) has launched two new in-house AI models. MAI-Voice-1 is a highly expressive, natural speech generation model that powers features like Copilot Daily and Podcasts, delivering fast, high-fidelity audio in single and multi-speaker scenarios. MAI-1-preview is an end-to-end trained mixture-of-experts foundation model designed to follow instructions and provide helpful responses, currently being tested publicly on LMArena and soon to be integrated into Copilot
- The launch marks Microsoft's most significant step yet toward developing its own AI foundation models

Why it matters

- Both models complement OpenAI's which are exclusively available on Microsoft's cloud. This is in line with Microsoft general open platform AI strategy
- Still, it comes at a time when the two firms are in the midst of reshaping their partnership. The negotiations are essential to OpenAI's corporate restructuring, on which fresh funding for the startup hinges. According to reports the conclusion of the talks might get pushed into next year. Adding more own models thus also adds more negotiation flexibility for Microsoft by reducing reliance on the partnership

All clear? Anthropic strikes surprise settlement in copyright lawsuit



Summary

- Anthropic settled a landmark class-action lawsuit from authors who accused it of using their books to train its GenAI without permission. Terms were not disclosed, but the deal avoids a high-profile trial and marks one of the first major AI copyright settlements
- Anthropic had previously won a partial victory when a a lower court ruled its use of books qualified as “fair use”, allowing the company to use the material for training without compensation. Still it found copyright infringements in the company storing the pirated books. This exposed Anthropic to a class-action suit with potentially crippling billions in liabilities

Why it matters

- This landmark resolution will be carefully looked at to see if it could set precedent for authors’ compensation in model training
- However, the “fair use” question remains unresolved. While in this case the principle was upheld, just days after the ruling another judge in a Meta copyright suit took a different view. He found Meta’s conduct may not be protected given its models could “flood the market” with replacements of the authors’ work
- The legal systems will thus continue to grapple with the question

More please. OpenAI and Anthropic (briefly) collaborated on AI Safety



Summary

- OpenAI and Anthropic collaborated on a joint AI safety testing, sharing access to reduced-safeguard versions of their models to identify blind spots
- The study found, for instance, that Anthropic's Claude Opus 4 models refuse uncertain answers more often, while OpenAI's models answer more but hallucinate more. It also found OpenAI's o3 more aligned to key human values, while Claude is leading following user instructions as well as refusing misuse
- Note: GPT-5 was not part of the study as it was at that time not yet released



Why it matters

- This was a one-off but hopefully will not remain so. Ensuring models are safe and trustworthy is an industry-wide challenge and such cross-company collaboration is therefore vital
- Efforts need to continue to address critical issues like AI sycophancy, where models overly accommodate harmful user behaviour, exemplified by a lawsuit alleging ChatGPT contributed to a tragic suicide and other forms of harmful behaviour given the continued growth of AI ChatBot adoption



Another one. Apple continues search for fresh AI brain to power revamped Siri



Summary

- Apple is in early talks with Google to integrate Google's Gemini AI models into a revamped version of its Siri assistant, aiming to catch up in the AI race after multiple delays in Siri's upgrade. This new version is currently expected to launch in 2026
- Apple is working on two simultaneous streams to execute the GenAI Siri upgrade. One developing its own in-house LLM while another is evaluating external models. It is already considering OpenAI and Anthropic as potential partners, with Google now coming in as a third possible candidate



Why it matters

- A choice for a third-party model, no matter from whom, would mark a strategic shift for Apple, which has historically developed all core elements of software internally, believing only this way the company can ensure its customers get a leading user experience
- However, this is but the latest in a long list of news suggesting Apple is in a quandary. Not only does it appear incapable of developing its own LLMs as fast as other tech companies, it seems unable to cut its losses going for a third-party and overall seems paralysed in making strategic decisions on future Siri