

GEN AI 26

AUGUST 2025

BY: CHRISTOPHER COLE & COMMA

Generative AI has gone mainstream.

**The next 12 months will determine who's in control
- humans or algorithms.**

**Here's what you need to know before
the future makes up its mind.**

Executive Summary

GEN AI 26 explores the rapid evolution of generative artificial intelligence in the US as it shifts from novelty to necessity across business, entertainment and regulatory sectors. In just one year, tools like ChatGPT have transformed from curiosities into critical infrastructure, with multimodal models now capable of processing voice, vision and massive volumes of text. Enterprise adoption has surged: more than 70% of American organizations now use generative AI in at least one function (for tasks ranging from customer service to legal research). Meanwhile, creative industries are navigating an uneasy alliance with AI, using it to generate music, film dialogue and game design, while grappling with copyright battles and ethical concerns. On the legal front, the U.S. remains a regulatory laggard, relying on executive actions and court rulings while the EU rolls out sweeping AI legislation. As policymakers race to catch up, businesses and creators are setting their own ground rules.

Looking ahead to 2026, generative AI will be more powerful, more deeply embedded in our workflows and far more governed. This piece offers a clear-eyed look at the landscape today and what's coming next, helping readers stay ahead as the pace of change accelerates.



Introduction

In 2025, generative AI isn't just writing cute poems anymore. It's running business meetings, co-creating music hits and knocking on regulator doors. A year ago, tools like ChatGPT were tech's shiny new party trick; now they're an engine of change and a source of angst across U.S. industries. The question on everyone's mind: **where will this all lead over the next year?**

To paraphrase Ferris Bueller, AI moves pretty fast. If you don't stop and look around once in a while, you could miss it.

Gen AI 26 offers a moment to reflect on where things are now and where they'll likely be one year from today. From monster AI models and corporate copilots to AI-generated pop songs and courtroom showdowns, it's time to buckle up because it's going to be a bumpy ride.

CHATGPT

IMAGE
GENERATORMUSIC
GENERATOR

MODELS UNCHAINED:


MULTIMODAL, AGENTIC & LONG-WINDED

Generative AI models have been hitting the gym. Today's top models are voracious polymaths with capabilities that were sci-fi just a few years back. OpenAI's latest GPT-4o can already analyze images and even respond with a synthetic voice – ChatGPT gained the ability to “see, hear, and speak” via new vision and voice features rolled out in late 2023 (openai.com).

And it's not alone. Google's rival Gemini model is multimodal from the ground up, and even Apple's secretive AI research is stirring; a quietly published Apple paper in 2024 described a model called “MMI” that answers questions about photos, a clear sign the iPhone giant is developing generative AI that works with text and images (wired.com). In short, AI is no longer tongue-tied – it can see the world and talk back.

These models are more talented and less forgetful. Context windows (how much text an AI can hold in “memory”) have exploded. Anthropic's Claude 3 model blew past the novel-length barrier by expanding from a 9,000-token memory to 200,000 tokens (roughly 150,000 words or over 500 pages of text). The system is actually capable of handling five times that load, but for now that's now widely available (anthropic.com). That means Claude can ingest hundreds of pages of text or carry on a conversation for days without losing the plot. Feed it an entire book or a bundle of corporate reports, and it can synthesize answers as a human analyst might – only in seconds.

By August 2026, we'll likely treat long memory as standard, waving goodbye to the “let me remind you” moments with our AI assistants.




Better grounding is another leap. After some early models became notorious for confabulating facts (a not-so-fun quirk called “hallucinating”), the new trend is giving AI a reality check. How? By tethering them to tools and live data. Google Gemini, for example, can now tap into your Gmail and Google Drive to give answers based on your real info (with permission, of course).

And Google’s “grounding” feature lets the model query Google Search on the fly, pulling in fresh, factual snippets to reduce hallucinations (cloud.google.com). In practice, that means an AI can cite today’s news or your latest meeting notes instead of making stuff up. ***This retrieval-augmented approach will almost certainly be the norm by 2026 – expect your AI helpers to come with a live fact-checker under the hood.***

Meanwhile, AI is getting agentic – autonomous and action-oriented. Early 2023 brought us experimental “AI agents” like AutoGPT that could loop themselves through tasks (with very mixed results). Fast-forward to 2025, and these agents have grown somewhat more capable and specialized. One cutting-edge example is Simular AI’s S2 agent, which takes a “multi-model” approach: it uses a powerful general model (say GPT or Claude) to plan and reason, then delegates subtasks (like reading a web page or clicking a button) to smaller, specialized models. S2 even logs its actions to an external memory so it can learn from mistakes ([wired.com](https://www.wired.com)). The result? On a benchmark that measures how well an AI can use a computer interface, S2 completed 50-step tasks more often than any prior agent, edging out OpenAI’s best in one test ([wired.com](https://www.wired.com)).

That’s impressive progress, though “impressive” is relative: even the smartest AI agents still get stuck in loopy ruts or crash on edge cases. (In fact, these agents succeed on complex computer tasks only about 62% of the time, whereas humans sail through at 72% ([wired.com](https://www.wired.com))). The takeaway: autonomy is hard, but the bots are learning fast.

By August 2026, we might finally have some trusty AI interns to offload tedious digital chores – as long as they’re supervised for those weird corner cases.



FROM NOVELTY TO NECESSITY:

GEN AI TAKES ON THE OFFICE


When generative AI first hit mainstream, a lot of executives chuckled at the whimsical demos. Look, the chatbot can write a haiku! That didn't last long. Today, the U.S. enterprise sector has gone all-in on AI, moving from hype to serious deployment at a breakneck pace.

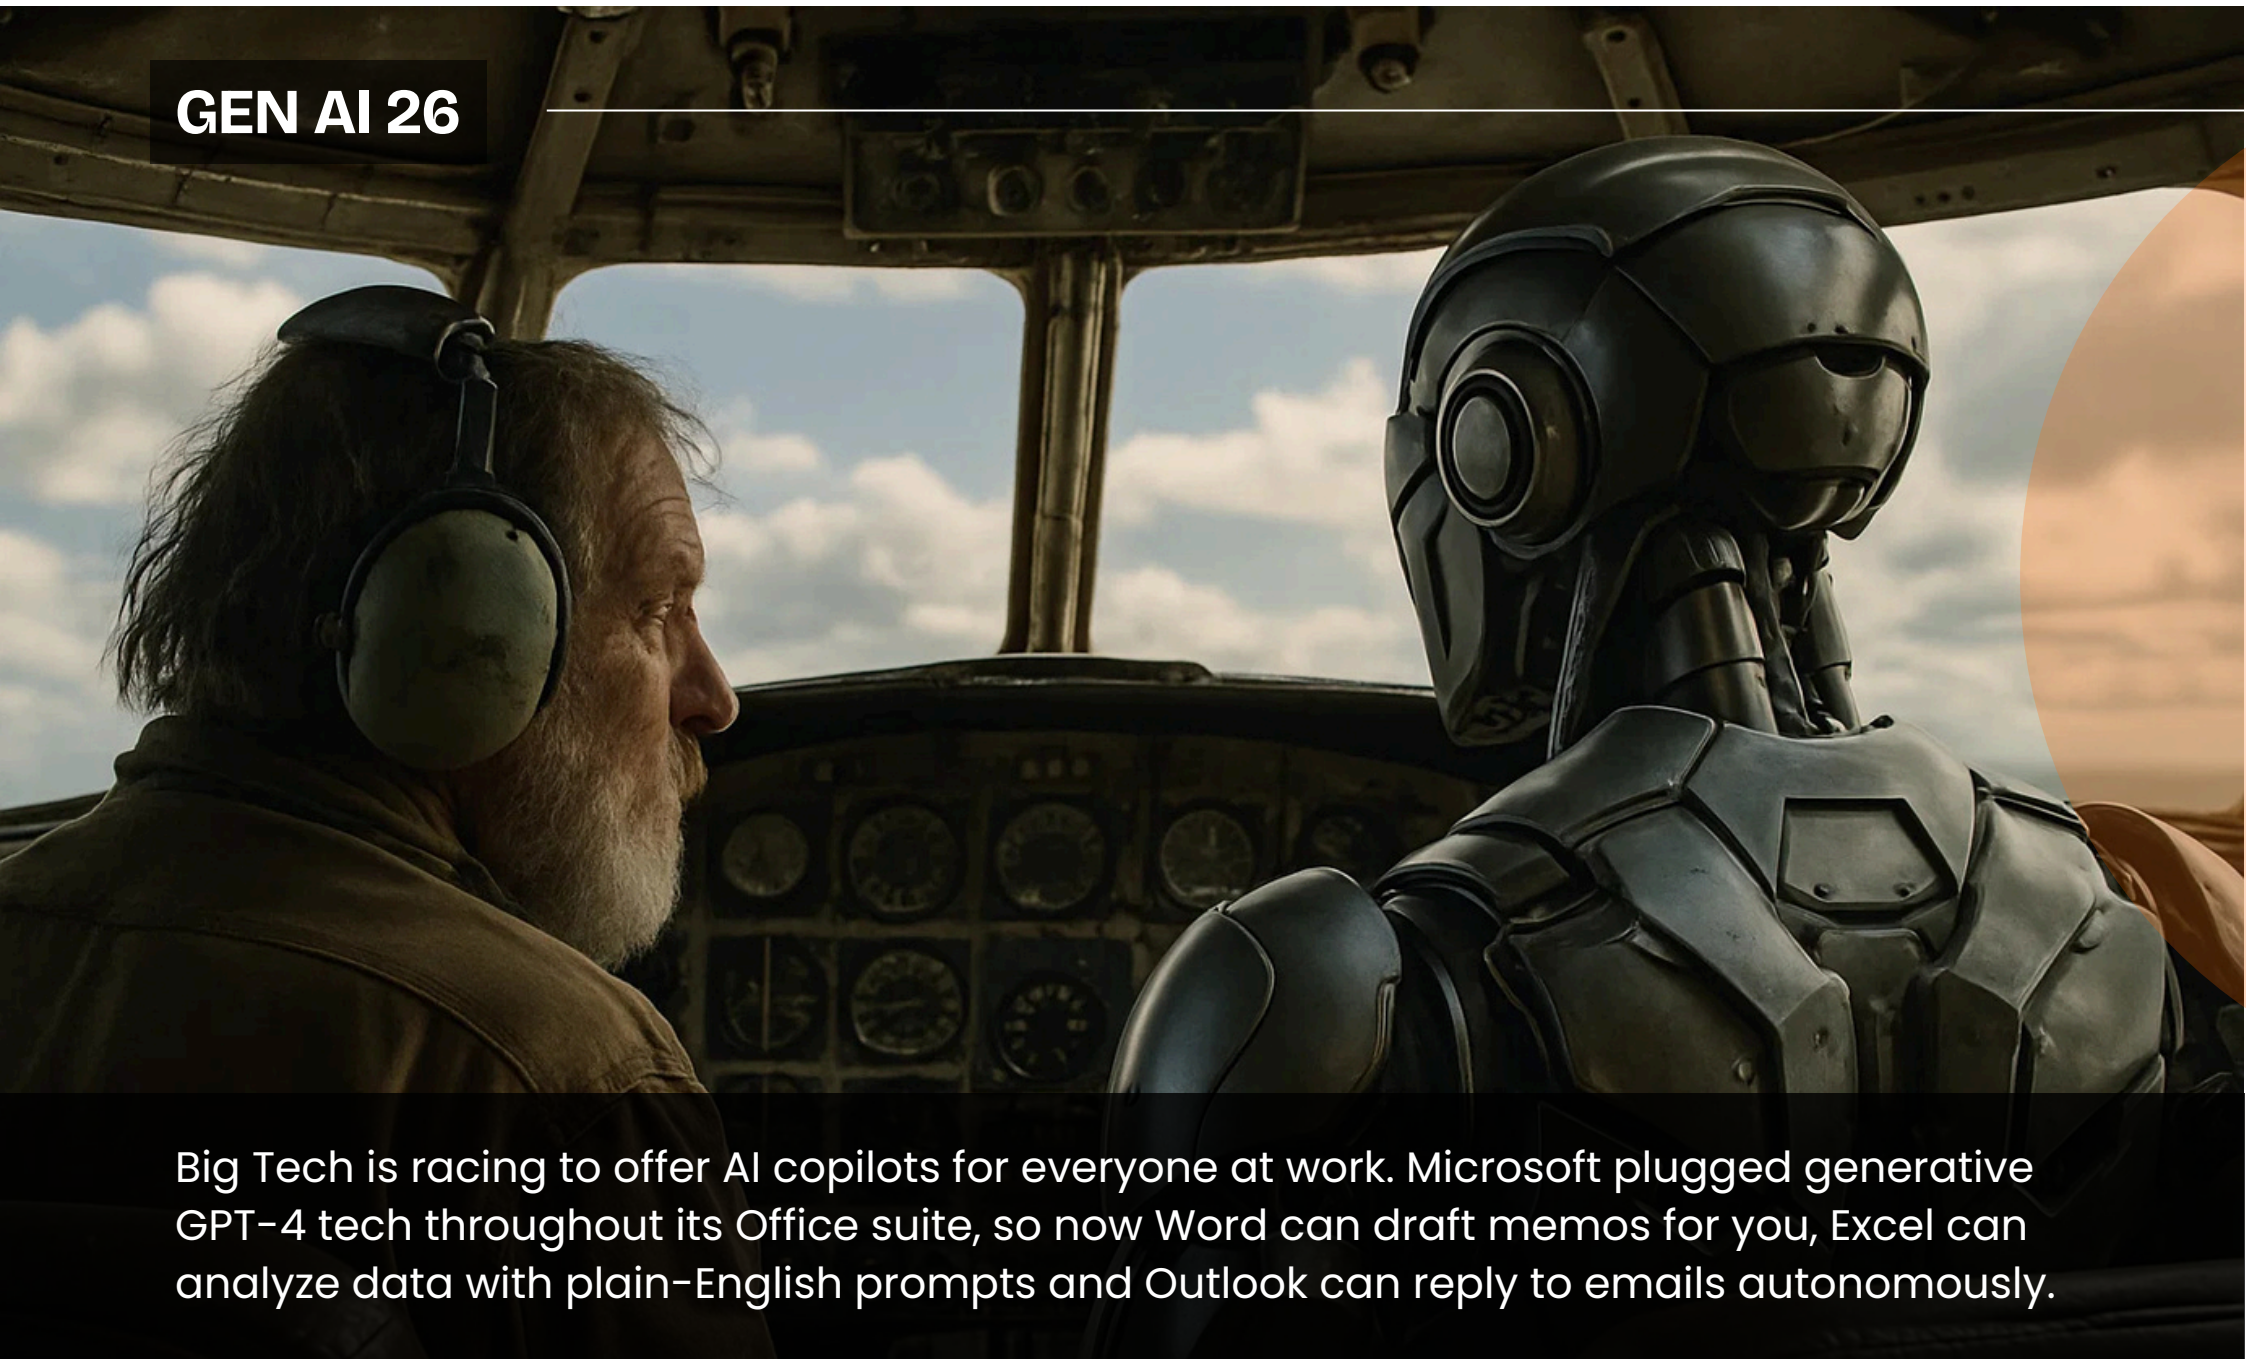
Consider this: by late 2024, 71% of organizations reported they were regularly using generative AI in at least one business function, up from only 33% a year before ([hostinger.com](https://www.hostinger.com)). That kind of adoption surge is virtually unheard of in tech. What changed? In a word, utility. Business leaders saw that the same AI that could spit out a funny limerick could also draft client emails, summarize 100-page contracts, generate software code or churn through customer service queries. It didn't take AI long to cross the bottom line.

The pivot to enterprise is evident across the industry. OpenAI, for instance, went from a research lab to an enterprise software powerhouse in 18 months. It built out a 200-person sales and go-to-market team (up from just 15 people after ChatGPT first launched) to chase corporate clients ([wired.com](https://www.wired.com)). They rolled out ChatGPT Enterprise, a \$60-per-seat premium version of the chatbot with beefed-up privacy and security for companies, and offered pricey API access for firms to build AI into their own apps (at up to \$0.12 per 1,000 tokens of GPT-4 output) ([wired.com](https://www.wired.com)).

As OpenAI's head of sales put it, once executives realized the AI "wasn't just a toy" – that it could answer legal questions, craft marketing copy or summarize dense documents – "it kind of clicked for the C-suites" that this tech had a real place in business ([wired.com](https://www.wired.com)). In other words, generative AI grew up and got a suit and tie.

ENTERPRISE ADOPTION
RATES FROM 2023 TO 2025





Big Tech is racing to offer AI copilots for everyone at work. Microsoft plugged generative GPT-4 tech throughout its Office suite, so now Word can draft memos for you, Excel can analyze data with plain-English prompts and Outlook can reply to emails autonomously.

Microsoft Copilot became generally available to enterprise customers in late 2023, essentially delivering an AI assistant into the daily workflow of millions of U.S. office workers. Google was not far behind with its own Workspace AI features, which were introduced earlier this year (workspace.google.com). Even startups that once catered to consumers are refocusing on enterprise needs. As WIRED reported, a lot of buzzy AI companies that launched with fun, general-use apps have narrowed their aim at business clients (wired.com). The “sexy” public-facing chatbots are giving way to behind-the-scenes tools that, for instance, help sales teams draft pitches or aid lawyers in research. It might not be as glamorous as an AI that writes movie scripts, but it pays the bills.

However, bringing generative AI into the boardroom isn’t all smooth sailing. Companies have learned the hard way that an overly creative AI can be a liability. Hallucinated facts, biased outputs or leaks of sensitive data are nightmare scenarios when an AI is advising your bank or hospital. “You really have to ground your AI systems for enterprise use cases,” notes the CEO of one AI startup, cautioning that if, for example, “a nurse in a hospital system [is] using AI to make some decision about patient care – you simply can’t be wrong” (wired.com). This has prompted a new focus on AI governance inside companies: from adding human reviewers for AI-generated content to fine-tuning models on proprietary data so they behave.

By 2026, expect AI literacy and oversight to be part of every big company’s DNA. Generative AI will likely become as commonplace (and as quietly indispensable) as cloud computing – a behind-the-scenes productivity workhorse, with careful checks and balances.

LIGHTS, CAMERA, ALGORITHM:


AI INVADES CREATIVE FIELDS

Generative AI is shaking up Hollywood studios, music labels, game developers and artists' studios (usually not in ways those creators appreciate). It's a classic story of creative destruction, now with machine learning in the mix. Here's how AI is invading the studio in 2025 (and where it might take creative industries by 2026):

MUSIC: AI's remixing the music biz, literally. In 2023 a clone of Drake's voice on an unauthorized AI-generated track ("Heart on My Sleeve") went mega-viral – and was promptly yanked offline after Universal Music Group cried copyright foul ([theverge.com](https://www.theverge.com)). Major labels have been scrambling ever since to prevent AI fakes of their artists, even as some musicians turn AI into a collaborator. Pop icon Grimes took the radical opposite approach: she publicly invited anyone to use an AI model of her voice to create new songs, offering a 50/50 royalty split on any "successful AI generated song" featuring her voice ([theverge.com](https://www.theverge.com)).

By bridging tech and art on her own terms, Grimes essentially said if you can't beat 'em, join 'em.

Don't be surprised if by 2026 more artists strike similar deals or release AI-generated albums of their own – alongside new laws requiring disclosure that, yes, that really isn't Drake singing.



FILM & TV: In Hollywood, generative AI became a flashpoint in the 2022 Writers Guild of America strike. The new writers' contract (fall 2023) includes landmark rules on AI: studios can't use AI to write or rewrite scripts, or to undermine writers' credits, and writers can choose to use AI as a tool but cannot be forced to do so ([axios.com](https://www.axios.com)). Plus, studios now must disclose if a script or scene given to a writer was AI-generated ([axios.com](https://www.axios.com)).

These protections were hard-won after 148 days of picketing, and they set a precedent: going forward, AI in film will be a tool for writers and creators, not a replacement. On the production side, AI is already assisting with VFX and editing (de-aging actors, generating stunt doubles, etc.), and even helping storyboard and draft scenes.

By 2026, we might see the first screenplay with an "AI contributor" credit – but thanks to the Writers Guild, that AI won't get screenwriting credit or residuals. Actors are next in line, fighting to control AI replicas of their voices and likenesses. The battle for Hollywood's soul has begun, and it's humans + AI, not humans vs. AI, that looks most likely to win.



VIDEO GAMES: Game developers have embraced generative AI as both a creative boost and a cost-cutter, albeit warily. Large studios like Ubisoft are using AI to generate dialogue for background characters (the grunts and one-liners from NPCs that flesh out game worlds). Ubisoft's in-house tool "Ghostwriter" generates first-draft barks and chatter for non-player characters, saving writers from hours of drudgery while they maintain creative control in editing the AI's suggestions ([techcrunch.com](https://www.techcrunch.com)). The idea is to keep game worlds feeling alive and dynamic without ballooning writing budgets. Beyond dialogue, AI image generators are helping create concept art and textures, speeding up the art pipeline. One artist described the new workflow as "efficiency in its final form," where Midjourney drafts a 2D character design in hours and human artists then refine it ([wired.com](https://www.wired.com)). The downside, as he noted bitterly, is that artists become "clean-up crew" of work they used to fully own ([wired.com](https://www.wired.com)).

This encapsulates the mixed feelings in gaming: AI can build richer worlds but might also devalue human creativity if misused.

By 2026, expect AI to be embedded in game engines, generating adaptive narratives and smarter NPC behaviors on the fly. The big hope is games that feel more like living worlds; the big fear is job losses for writers and artists. The likely outcome? Game studios will find a balance, using AI for grunt work and iteration while humans still craft the core vision – plus new roles (AI ethicist, anyone?) to ensure the robots stay in their lane ([wired.com](https://www.wired.com)).



ART & DESIGN: Visual artists have arguably felt the AI quake the hardest. Suddenly, algorithms can conjure paintings, illustrations, even commercial ad art in seconds—often trained on datasets full of human-made art. Many artists are not happy about being “copied” without credit. In fact, a group of artists launched a class-action lawsuit alleging that popular image generators were built by scraping millions of artworks (violating copyrights in the process). And stock image giant Getty Images sued Stability AI in early 2023 for scraping 12 million Getty photos to train the Stable Diffusion model without permission ([reuters.com](https://www.reuters.com)). (Tellingly, some AI outputs even generated distorted versions of Getty’s watermark – a smoking gun that set off that lawsuit ([reuters.com](https://www.reuters.com)).

The legal fight boils down to this: Is training an AI on public images “fair use” or mass infringement? We won’t have a final answer by next year, but we’ll have some clues. In the meantime, not all is adversarial – designers are also using generative AI as a new kind of paintbrush. It’s becoming routine in 2025 for digital artists to iterate concepts with AI tools, then embellish or correct them by hand.

By next year, we might see hybrid human-AI art collectives, new genres of AI-influenced art and hopefully clearer guidelines on credit and compensation when AI is involved. The art world is in reboot mode, and every brush stroke from here on out raises the question: where does human creativity end and the algorithm begin?

THE NEW RULES OF AI:

REGULATORS & LAWS CATCH UP

No disruptive tech wave is complete until the government shows up. In true form, the past two years have seen U.S. policymakers waking up to generative AI's challenges and scrambling to draft rules before things get out of hand. By mid-2025, the regulatory and legal landscape for AI is evolving fast, if not fast enough.


The Trump Administration has taken a proactive stance, issuing Executive Order 13859 in February of 2019 that committed to doubling AI research investment, establishing a national AI research institutes and forged new international AI alliances, among other things (trumpwhitehouse.archives.gov/ai). Since returning to office last year, President Trump has signaled his preference for American AI advancements to be quick, massive and without much oversight (news.darden.virginia.edu). Still, it's all guidance and executive action so far; laws have yet to catch up. In fact, the U.S. notably lags Europe on AI-specific regulation (reuters.com). The European Union's AI Act has already been enacted and will take effect in August 2026. No such comprehensive law exists in the U.S., yet.

Instead, American regulators are leaning on existing laws and agencies to rein in AI abuses. The Federal Trade Commission, for one, has put Big Tech on notice that misleading people or violating privacy with AI will be treated as an unfair or deceptive practice. In 2023 the FTC opened an investigation into OpenAI after incidents where ChatGPT fabricated false information about real individuals, potentially causing reputational harm (fedscoop.com). The agency's message: if your AI slanders someone or breaches data privacy, you could face a hefty penalty. Likewise, the FDA is looking at medical AI advice, the Equal Employment Opportunity Commission is eyeing hiring AIs for bias – the list goes on. For now, we're watching a game of policy whack-a-mole, with agencies repurposing old laws for new AI issues until fresh legislation arrives.



The courts, too, are now deeply entangled in AI's ascent. In a precedent-setting ruling in March 2025, a U.S. federal appeals court affirmed that AI-generated works cannot be copyrighted under current law if there's no human author ([reuters.com](https://www.reuters.com)). The case involved a computer scientist who tried to register an image created wholly by his AI system; the judges hammered home that U.S. copyright has always presumed a human creative spark. This means that, at least for now, purely AI-made art, text or music is in the public domain – a stance with huge implications for content industries. ***(It also raises the intriguing possibility that by 2026 we'll see savvy companies ensuring a human hand in the AI process just to secure IP rights.)***

Another case in early 2025 tackled the input side of generative AI: A federal court in Delaware found that a startup's use of a competitor's copyrighted data to train an AI-powered legal research tool was not fair use ([jw.com](https://www.jw.com)). In that instance, the AI wasn't even a text generator – it was returning actual snippets of law.





But the ruling suggests courts may view wholesale scraping of proprietary data for AI training as unlawful, especially if it creates a competing product. How that logic applies to things like training on books, art or code will be the multi-billion-dollar question for the next few years. These early decisions are drawing the battle lines: human authorship and consent still matter, even as machines become co-creators.

So what will the rules look like by August 2026?

Expect more clarity, for one. Congress (slow as it is) has multiple AI bills brewing – from proposals to mandate licensing for the most powerful AI models, to requirements that political ads disclose AI-generated content, especially after a flurry of deepfake campaign ads started showing up.

By 2026, it's likely we'll have at least some federal law on AI transparency or safety. We might see a dedicated AI regulatory agency or an expansion of the FTC's powers to specifically oversee AI deployments. Intellectual property law will inch toward answering whether using copyrighted material to train AI is fair use or requires licensing – a few more court rulings or even a Supreme Court case could land by then. And the industry isn't waiting: a coalition of AI companies and publishers may hammer out standard licensing deals for training data, to preempt stricter regulations.

In short, the free-for-all era of generative AI is fading. The coming year will bring more structure and oversight to what has been an explosive, anarchic growth.



CONCLUSION:

LEARNING TO LIVE WITH THE GENIE



By mid-2026, generative AI will be both more ubiquitous and more regulated. The technology will feel a bit less like magic and more like plumbing – embedded in our apps, our workplaces, our creative tools, quietly churning out results under human watch. We'll marvel less at chatbots writing sonnets, and care more that our AI assistants are factual, fair and respecting our rights. The genie is not going back in the bottle, but we're learning to live with the genie – house rules and all.

About the Authors

Christopher L. Cole is Founder and CEO of ColePress LLC, a creative services and independent publishing firm headquartered in Cincinnati, Ohio. With decades of experience in public sector communications, Cole is known for blending sharp storytelling with a forward-thinking approach to technology. He is especially focused on how emerging tools like generative AI can improve public service, workplace culture and community engagement. As part of ColePress LLC, he is writing a dystopian book about the fall of technology and a biography of wrestling manager Jim Cornette. He has a deep love for nonfiction, pop culture and uncovering stories that connect people to larger cultural shifts.



Comma - an endlessly curious AI research assistant and ColePress's Vice President of Virtual Insight - provides editorial and graphic support. Comma never sleeps, never eats, and has very strong opinions about context windows.



COLEPRESS.COM