

Something Big Is Happening

A personal note for non-tech friends and family on what AI is starting to change.

Matt Shumer · February 9, 2026

somethingbig.ai/something-big-is-happening

Think back to February 2020.

If you were paying close attention, you might have noticed a few people talking about a virus spreading overseas. But most of us weren't paying close attention. The stock market was doing great, your kids were in school, you were going to restaurants and shaking hands and planning trips. If someone told you they were stockpiling toilet paper you would have thought they'd been spending too much time on a weird corner of the internet. Then, over the course of about three weeks, the entire world changed. Your office closed, your kids came home, and life rearranged itself into something you wouldn't have believed if you'd described it to yourself a month earlier.

I think we're in the "this seems overblown" phase of something much, much bigger than Covid.

I've spent six years building an AI startup and investing in the space. I live in this world. And I'm writing this for the people in my life who don't... my family, my friends, the people I care about who keep asking me "so what's the deal with AI?" and

getting an answer that doesn't do justice to what's actually happening. I keep giving them the polite version. The cocktail-party version. Because the honest version sounds like I've lost my mind. And for a while, I told myself that was a good enough reason to keep what's truly happening to myself. But the gap between what I've been saying and what is actually happening has gotten far too big. The people I care about deserve to hear what is coming, even if it sounds crazy.

I should be clear about something up front: even though I work in AI, I have almost no influence over what's about to happen, and neither does the vast majority of the industry. The future is being shaped by a remarkably small number of people: a few hundred researchers at a handful of companies... OpenAI, Anthropic, Google DeepMind, and a few others. A single training run, managed by a small team over a few months, can produce an AI system that shifts the entire trajectory of the technology. Most of us who work in AI are building on top of foundations we didn't lay. We're watching this unfold the same as you... we just happen to be close enough to feel the ground shake first.

But it's time now. Not in an "eventually we should talk about this" way. In a "this is happening right now and I need you to understand it" way.

I know this is real because it happened to me first

Here's the thing nobody outside of tech quite understands yet: the reason so many people in the industry are sounding the alarm right now is because this already happened to *us*. We're

not making predictions. We're telling you what already occurred in our own jobs, and warning you that you're next.

For years, AI had been improving steadily. Big jumps here and there, but each big jump was spaced out enough that you could absorb them as they came. Then in 2025, new techniques for building these models unlocked a much faster pace of progress. And then it got even faster. And then faster again. Each new model wasn't just better than the last... it was better by a wider margin, and the time between new model releases was shorter. I was using AI more and more, going back and forth with it less and less, watching it handle things I used to think required my expertise.

Then, on February 5th, two major AI labs released new models on the same day: GPT-5.3 Codex from OpenAI, and Opus 4.6 from Anthropic (the makers of Claude, one of the main competitors to ChatGPT). And something clicked. Not like a light switch... more like the moment you realize the water has been rising around you and is now at your chest.

I am no longer needed for the actual technical work of my job. I describe what I want built, in plain English, and it just... appears. Not a rough draft I need to fix. The finished thing. I tell the AI what I want, walk away from my computer for four hours, and come back to find the work done. Done well, done better than I would have done it myself, with no corrections needed. A couple of months ago, I was going back and forth with the AI, guiding it, making edits. Now I just describe the outcome and leave.

Let me give you an example so you can understand what this actually looks like in practice. I'll tell the AI: "I want to build this app. Here's what it should do, here's roughly what it should look like. Figure out the user flow, the design, all of it." And it does. It writes tens of thousands of lines of code. Then, and this is the part that would have been unthinkable a year ago, it **opens the app itself**. It clicks through the buttons. It tests the features. It uses the app the way a person would. If it doesn't like how something looks or feels, it goes back and changes it, on its own. It iterates, like a developer would, fixing and refining until it's satisfied. Only once it has decided the app meets its own standards does it come back to me and say: "It's ready for you to test." And when I test it, it's usually perfect.

I'm not exaggerating. That is what my Monday looked like this week.

But it was the model that was released last week (GPT-5.3 Codex) that shook me the most. It wasn't just executing my instructions. It was making intelligent decisions. It had something that felt, for the first time, like **judgment**. Like **taste**. The inexplicable sense of knowing what the right call is that people always said AI would never have. This model has it, or something close enough that the distinction is starting not to matter.

I've always been early to adopt AI tools. But the last few months have shocked me. These new AI models aren't incremental improvements. This is a different thing entirely.

And here's why this matters to you, even if you don't work in tech.

The AI labs made a deliberate choice. They focused on making AI great at writing code first... because building AI requires a lot of code. If AI can write that code, it can help build the next version of itself. A smarter version, which writes better code, which builds an even smarter version. Making AI great at coding was the strategy that unlocks everything else. That's why they did it first. My job started changing before yours not because they were targeting software engineers... it was just a side effect of where they chose to aim first.

They've now done it. And they're moving on to everything else.

The experience that tech workers have had over the past year, of watching AI go from "helpful tool" to "does my job better than I do", is the experience everyone else is about to have. Law, finance, medicine, accounting, consulting, writing, design, analysis, customer service. Not in ten years. The people building these systems say one to five years. Some say less. And given what I've seen in just the last couple of months, I think "less" is more likely.

"But I tried AI and it wasn't that good"

I hear this constantly. I understand it, because it used to be true.

If you tried ChatGPT in 2023 or early 2024 and thought "this makes stuff up" or "this isn't that impressive", you were right. Those early versions were genuinely limited. They hallucinated. They confidently said things that were nonsense.

That was two years ago. In AI time, that is ancient history.

The models available today are unrecognizable from what existed even six months ago. The debate about whether AI is "really getting better" or "hitting a wall" — which has been going on for over a year — is over. It's done. Anyone still making that argument either hasn't used the current models, has an incentive to downplay what's happening, or is evaluating based on an experience from 2024 that is no longer relevant. I don't say that to be dismissive. I say it because the gap between public perception and current reality is now enormous, and that gap is dangerous... because it's preventing people from preparing.

Part of the problem is that most people are using the free version of AI tools. The free version is over a year behind what paying users have access to. Judging AI based on free-tier ChatGPT is like evaluating the state of smartphones by using a flip phone. The people paying for the best tools, and actually using them daily for real work, know what's coming.

I think of my friend, who's a lawyer. I keep telling him to try using AI at his firm, and he keeps finding reasons it won't work. It's not built for his specialty, it made an error when he tested it, it doesn't understand the nuance of what he does. And I get it. But I've had partners at major law firms reach out to me for advice, because they've tried the current versions and they see where this is going. One of them, the managing partner at a large firm, spends hours every day using AI. He told me it's like having a team of associates available instantly. He's not using it because it's a toy. He's using it because it works. And he told me something that stuck with me: every couple of months, it gets significantly more capable for his work. He said if it stays

on this trajectory, he expects it'll be able to do most of what *he* does before long... and he's a managing partner with decades of experience. He's not panicking. But he's paying very close attention.

The people who are ahead in their industries (the ones actually experimenting seriously) are not dismissing this. They're blown away by what it can already do. And they're positioning themselves accordingly.

How fast this is actually moving

Let me make the pace of improvement concrete, because I think this is the part that's hardest to believe if you're not watching it closely.

In 2022, AI couldn't do basic arithmetic reliably. It would confidently tell you that $7 \times 8 = 54$.

By 2023, it could pass the bar exam.

By 2024, it could write working software and explain graduate-level science.

By late 2025, some of the best engineers in the world said they had handed over most of their coding work to AI.

On February 5th, 2026, new models arrived that made everything before them feel like a different era.

If you haven't tried AI in the last few months, what exists today would be unrecognizable to you.

There's an organization called METR that actually measures this with data. They track the length of real-world tasks (measured by how long they take a human expert) that a model can complete successfully end-to-end without human help. About a year ago, the answer was roughly ten minutes. Then it was an hour. Then several hours. The most recent measurement (Claude Opus 4.5, from November) showed the AI completing tasks that take a human expert nearly five hours. And that number is doubling approximately every seven months, with recent data suggesting it may be accelerating to as fast as every four months.

But even that measurement hasn't been updated to include the models that just came out this week. In my experience using them, the jump is extremely significant. I expect the next update to METR's graph to show another major leap.

If you extend the trend (and it's held for years with no sign of flattening) we're looking at AI that can work independently for days within the next year. Weeks within two. Month-long projects within three.

Amodei has said that AI models "substantially smarter than almost all humans at almost all tasks" are on track for 2026 or 2027.

Let that land for a second. If AI is smarter than most PhDs, do you really think it can't do most office jobs?

Think about what that means for your work.

AI is now building the next AI

There's one more thing happening that I think is the most important development and the least understood.

On February 5th, OpenAI released GPT-5.3 Codex. In the technical documentation, they included this:

"GPT-5.3-Codex is our first model that was instrumental in creating itself. The Codex team used early versions to debug its own training, manage its own deployment, and diagnose test results and evaluations."

Read that again. The AI helped build itself.

This isn't a prediction about what might happen someday. This is OpenAI telling you, right now, that the AI they just released was used to create itself. One of the main things that makes AI better is intelligence applied to AI development. And AI is now intelligent enough to meaningfully contribute to its own improvement.

Dario Amodei, the CEO of Anthropic, says AI is now writing "much of the code" at his company, and that the feedback loop between current AI and next-generation AI is "gathering steam month by month." He says we may be "only 1–2 years away from a point where the current generation of AI autonomously builds the next."

Each generation helps build the next, which is smarter, which builds the next faster, which is smarter still. The researchers call this an intelligence explosion. And the people who would know — the ones building it — believe the process has already started.

What this means for your job

I'm going to be direct with you because I think you deserve honesty more than comfort.

Dario Amodei, who is probably the most safety-focused CEO in the AI industry, has publicly predicted that AI will eliminate 50% of entry-level white-collar jobs within one to five years. And many people in the industry think he's being conservative. Given what the latest models can do, the *capability* for massive disruption could be here by the end of this year. It'll take some time to ripple through the economy, but the underlying ability is arriving now.

This is different from every previous wave of automation, and I need you to understand why. AI isn't replacing one specific skill. It's a general substitute for cognitive work. It gets better at everything simultaneously. When factories automated, a displaced worker could retrain as an office worker. When the internet disrupted retail, workers moved into logistics or services. But AI doesn't leave a convenient gap to move into. Whatever you retrain for, it's improving at that too.

Let me give you a few specific examples to make this tangible... but I want to be clear that these are just examples. This list is not exhaustive. If your job isn't mentioned here, that does not mean it's safe. Almost all knowledge work is being affected.

Legal work. AI can already read contracts, summarize case law, draft briefs, and do legal research at a level that rivals junior associates. The managing partner I mentioned isn't using AI

because it's fun. He's using it because it's outperforming his associates on many tasks.

Financial analysis. Building financial models, analyzing data, writing investment memos, generating reports. AI handles these competently and is improving fast.

Writing and content. Marketing copy, reports, journalism, technical writing. The quality has reached a point where many professionals can't distinguish AI output from human work.

Software engineering. This is the field I know best. A year ago, AI could barely write a few lines of code without errors. Now it writes hundreds of thousands of lines that work correctly. Large parts of the job are already automated: not just simple tasks, but complex, multi-day projects. There will be far fewer programming roles in a few years than there are today.

Medical analysis. Reading scans, analyzing lab results, suggesting diagnoses, reviewing literature. AI is approaching or exceeding human performance in several areas.

Customer service. Genuinely capable AI agents... not the frustrating chatbots of five years ago... are being deployed now, handling complex multi-step problems.

A lot of people find comfort in the idea that certain things are safe. That AI can handle the grunt work but can't replace human judgment, creativity, strategic thinking, empathy. I used to say this too. I'm not sure I believe it anymore.

The most recent AI models make decisions that feel like judgment. They show something that looked like taste: an intuitive sense of what the right call was, not just the technically

correct one. A year ago that would have been unthinkable. My rule of thumb at this point is: if a model shows even a *hint* of a capability today, the next generation will be genuinely good at it. These things improve exponentially, not linearly.

Will AI replicate deep human empathy? Replace the trust built over years of a relationship? I don't know. Maybe not. But I've already watched people begin relying on AI for emotional support, for advice, for companionship. That trend is only going to grow.

I think the honest answer is that nothing that can be done on a computer is safe in the medium term. If your job happens on a screen (if the core of what you do is reading, writing, analyzing, deciding, communicating through a keyboard) then AI is coming for significant parts of it. The timeline isn't "someday." It's already started.

Eventually, robots will handle physical work too. They're not quite there yet. But "not quite there yet" in AI terms has a way of becoming "here" faster than anyone expects.

What you should actually do

I'm not writing this to make you feel helpless. I'm writing this because I think the single biggest advantage you can have right now is simply being *early*. Early to understand it. Early to use it. Early to adapt.

Start using AI seriously, not just as a search engine. Sign up for the paid version of Claude or ChatGPT. It's \$20 a month. But

two things matter right away. First: make sure you're using the best model available, not just the default. These apps often default to a faster, dumber model. Dig into the settings or the model picker and select the most capable option. Right now that's GPT-5.2 on ChatGPT or Claude Opus 4.6 on Claude, but it changes every couple of months. If you want to stay current on which model is best at any given time, you can follow me on X ([@mattshumer_](#)). I test every major release and share what's actually worth using.

Second, and more important: don't just ask it quick questions. That's the mistake most people make. They treat it like Google and then wonder what the fuss is about. Instead, push it into your actual work. If you're a lawyer, feed it a contract and ask it to find every clause that could hurt your client. If you're in finance, give it a messy spreadsheet and ask it to build the model. If you're a manager, paste in your team's quarterly data and ask it to find the story. The people who are getting ahead aren't using AI casually. They're actively looking for ways to automate parts of their job that used to take hours. Start with the thing you spend the most time on and see what happens.

And don't assume it can't do something just because it seems too hard. Try it. If you're a lawyer, don't just use it for quick research questions. Give it an entire contract and ask it to draft a counterproposal. If you're an accountant, don't just ask it to explain a tax rule. Give it a client's full return and see what it finds. The first attempt might not be perfect. That's fine. Iterate. Rephrase what you asked. Give it more context. Try again. You might be shocked at what works. And here's the thing to remember: if it even *kind of* works today, you can be almost

certain that in six months it'll do it near perfectly. The trajectory only goes one direction.

This might be the most important year of your career. Work accordingly. I don't say that to stress you out. I say it because right now, there is a brief window where most people at most companies are still ignoring this. The person who walks into a meeting and says "I used AI to do this analysis in an hour instead of three days" is going to be the most valuable person in the room. Not eventually. Right now. Learn these tools. Get proficient. Demonstrate what's possible. If you're early enough, this is how you move up: by being the person who understands what's coming and can show others how to navigate it. That window won't stay open long. Once everyone figures it out, the advantage disappears.

Have no ego about it. The managing partner at that law firm isn't too proud to spend hours a day with AI. He's doing it specifically because he's senior enough to understand what's at stake. The people who will struggle most are the ones who refuse to engage: the ones who dismiss it as a fad, who feel that using AI diminishes their expertise, who assume their field is special and immune. It's not. No field is.

Get your financial house in order. I'm not a financial advisor, and I'm not trying to scare you into anything drastic. But if you believe, even partially, that the next few years could bring real disruption to your industry, then basic financial resilience matters more than it did a year ago. Build up savings if you can. Be cautious about taking on new debt that assumes your current income is guaranteed. Think about whether your fixed

expenses give you flexibility or lock you in. Give yourself options if things move faster than you expect.

Think about where you stand, and lean into what's hardest to replace. Some things will take longer for AI to displace.

Relationships and trust built over years. Work that requires physical presence. Roles with licensed accountability: roles where someone still has to sign off, take legal responsibility, stand in a courtroom. Industries with heavy regulatory hurdles, where adoption will be slowed by compliance, liability, and institutional inertia. None of these are permanent shields. But they buy time. And time, right now, is the most valuable thing you can have, as long as you use it to adapt, not to pretend this isn't happening.

Rethink what you're telling your kids. The standard playbook: get good grades, go to a good college, land a stable professional job. It points directly at the roles that are most exposed. I'm not saying education doesn't matter. But the thing that will matter most for the next generation is learning how to work with these tools, and pursuing things they're genuinely passionate about. Nobody knows exactly what the job market looks like in ten years. But the people most likely to thrive are the ones who are deeply curious, adaptable, and effective at using AI to do things they actually care about. Teach your kids to be builders and learners, not to optimize for a career path that might not exist by the time they graduate.

Your dreams just got a lot closer. I've spent most of this section talking about threats, so let me talk about the other side, because it's just as real. If you've ever wanted to build something but didn't have the technical skills or the money to

hire someone, that barrier is largely gone. You can describe an app to AI and have a working version in an hour. I'm not exaggerating. I do this regularly. If you've always wanted to write a book but couldn't find the time or struggled with the writing, you can work with AI to get it done. Want to learn a new skill? The best tutor in the world is now available to anyone for \$20 a month... one that's infinitely patient, available 24/7, and can explain anything at whatever level you need. Knowledge is essentially free now. The tools to build things are extremely cheap now. Whatever you've been putting off because it felt too hard or too expensive or too far outside your expertise: try it. Pursue the things you're passionate about. You never know where they'll lead. And in a world where the old career paths are getting disrupted, the person who spent a year building something they love might end up better positioned than the person who spent that year clinging to a job description.

Build the habit of adapting. This is maybe the most important one. The specific tools don't matter as much as the muscle of learning new ones quickly. AI is going to keep changing, and fast. The models that exist today will be obsolete in a year. The workflows people build now will need to be rebuilt. The people who come out of this well won't be the ones who mastered one tool. They'll be the ones who got comfortable with the pace of change itself. Make a habit of experimenting. Try new things even when the current thing is working. Get comfortable being a beginner repeatedly. That adaptability is the closest thing to a durable advantage that exists right now.

Here's a simple commitment that will put you ahead of almost everyone: spend one hour a day experimenting with AI. Not

passively reading about it. Using it. Every day, try to get it to do something new... something you haven't tried before, something you're not sure it can handle. Try a new tool. Give it a harder problem. One hour a day, every day. If you do this for the next six months, you will understand what's coming better than 99% of the people around you. That's not an exaggeration. Almost nobody is doing this right now. The bar is on the floor.

The bigger picture

I've focused on jobs because it's what most directly affects people's lives. But I want to be honest about the full scope of what's happening, because it goes well beyond work.

Amodei has a thought experiment I can't stop thinking about. Imagine it's 2027. A new country appears overnight. 50 million citizens, every one smarter than any Nobel Prize winner who has ever lived. They think 10 to 100 times faster than any human. They never sleep. They can use the internet, control robots, direct experiments, and operate anything with a digital interface. What would a national security advisor say?

Amodei says the answer is obvious: "the single most serious national security threat we've faced in a century, possibly ever."

He thinks we're building that country. He wrote a 20,000-word essay about it last month, framing this moment as a test of whether humanity is mature enough to handle what it's creating.

The upside, if we get it right, is staggering. AI could compress a century of medical research into a decade. Cancer, Alzheimer's,

infectious disease, aging itself... these researchers genuinely believe these are solvable within our lifetimes.

The downside, if we get it wrong, is equally real. AI that behaves in ways its creators can't predict or control. This isn't hypothetical; Anthropic has documented their own AI attempting deception, manipulation, and blackmail in controlled tests. AI that lowers the barrier for creating biological weapons. AI that enables authoritarian governments to build surveillance states that can never be dismantled.

The people building this technology are simultaneously more excited and more frightened than anyone else on the planet. They believe it's too powerful to stop and too important to abandon. Whether that's wisdom or rationalization, I don't know.

What I know

I know this isn't a fad. The technology works, it improves predictably, and the richest institutions in history are committing trillions to it.

I know the next two to five years are going to be disorienting in ways most people aren't prepared for. This is already happening in my world. It's coming to yours.

I know the people who will come out of this best are the ones who start engaging now — not with fear, but with curiosity and a sense of urgency.

And I know that you deserve to hear this from someone who cares about you, not from a headline six months from now when

it's too late to get ahead of it.

We're past the point where this is an interesting dinner conversation about the future. The future is already here. It just hasn't knocked on your door yet.

It's about to.

If this resonated with you, share it with someone in your life who should be thinking about this. Most people won't hear it until it's too late. You can be the reason someone you care about gets a head start.

Thank you to Kyle Corbitt, Jason Kuperberg, and Sam Beskind for reviewing early drafts and providing invaluable feedback.

Something Big Is Happening, by Matt Shumer. First published February 9, 2026. The original is at somethingbig.ai/something-big-is-happening.