Why Most AI Breaks in the Real World — and What Founders Get Wrong

Opinions expressed by Entrepreneur contributors are their own.

Key Takeaways

- AI often fails outside of demos because it can’t learn from real-world mistakes or adapt to unpredictable users and systems.
- Founders who focus on AI that improves over time — not just executes commands — are the ones turning automation into real business results.

According to the internet, startups are running entire companies on AI. Founders have AI sales teams closing deals while they sleep. AI agents are supposedly replacing full departments overnight.

Meanwhile, your agents stall out. They make questionable tool calls, get stuck in loops and fail to complete tasks reliably.

That doesn’t mean you’re behind. It means you’re operating in the real world.

Your AI agents interact with real customers, real enterprise systems and real constraints. When they make mistakes, those mistakes don’t disappear into a demo — they cost time, money, and credibility.

You’re not alone

Research from MIT helps explain why this gap exists.

Tools like ChatGPT are now ubiquitous. MIT found that roughly 90% of employees in surveyed companies use large language models regularly at work. Coding agents such as Claude Code, Cursor and Codex have become standard in many developer workflows.

But the area with the most excitement is also the area with the least success: AI agents designed to automate tasks — and eventually entire business functions.

MIT’s research found that 95% of pilot projects involving task-specific or embedded generative AI failed to deliver sustained productivity or P&L impact once deployed to production.

Why? Because today’s AI works well for simple tasks but breaks down when the stakes are higher. Users turn to ChatGPT for quick answers, then abandon it for mission-critical work. What’s missing are systems that can adapt, remember, and improve over time.

Researchers are paying attention

This limitation hasn’t gone unnoticed.

Research teams from institutions including Stanford and the University of Illinois have published studies showing that most AI agents struggle to adapt based on their own experiences. Google DeepMind has explored the same problem through its work on Evo-Memory, which evaluates how well an agent learns and evolves while operating.

My own research has focused on this gap as well. In a research paper I co-authored with Virginia Tech’s Sanghani Center for

AI and Data Analytics, we proposed a new approach to agent memory called Hindsight. The research showed how using memory pathways to store and reflect on agent experiences allows agents to learn from those experiences.

Together, these efforts point to an important shift: the emergence of adaptive agent memory.

Why this matters in the real world

Today, when an AI agent fails, engineers fix it manually. They tweak prompts, rewrite instructions, change tool descriptions or add examples. These changes can help — but they don't scale.

Prompts grow longer and more fragile. Fixes for one issue can break something else that was working. And once an agent is live, the problem compounds.

Real users behave unpredictably. Interaction volumes increase. Failures become harder to track and diagnose. A single error is manageable. Dozens of failures a day are not.

Without a way for AI to learn from these interactions, progress remains incremental — and expensive.

Why memory is the missing piece

To understand why this matters, consider a simple question: what would Albert Einstein have accomplished if he had all his intelligence but no memory?

That's essentially the state of today's AI.

Modern language models are incredibly knowledgeable, yet they repeat the same mistakes because they don't learn from experience. A customer service agent that issues a refund incorrectly today is likely to make the same mistake tomorrow.

An agent that answers questions correctly 70% of the time has no understanding of why it fails the other 30%.

Early “memory” solutions didn’t solve this. They simply searched past conversations for context.

The next generation of adaptive agent memory is different. These systems allow agents to separate facts from experiences, reflect on outcomes, and ask a critical question: *How can I do better next time?*

The founder takeaway

For founders building an AI-powered workforce, this shift is significant.

The future isn’t just AI agents that execute instructions. It’s agents that improve themselves, reduce errors over time, and become more reliable the longer they operate.

That’s how AI moves from impressive demos to durable business impact — and how startups turn experimentation into a real competitive advantage.

Sign up for the Entrepreneur Daily newsletter to get the news and resources you need to know today to help you run your business better. Get it in your inbox.

Key Takeaways

- AI often fails outside of demos because it can’t learn from real-world mistakes or adapt to unpredictable users and systems.
- Founders who focus on AI that improves over time — not just executes commands — are the ones turning automation into real business results.

According to the internet, startups are running entire

companies on AI. Founders have AI sales teams closing deals while they sleep. AI agents are supposedly replacing full departments overnight.

Meanwhile, your agents stall out. They make questionable tool calls, get stuck in loops and fail to complete tasks reliably.

https://www.entrepreneur.com/science-technology/why-most-ai-breaks-in-the-real-world-and-what-founders/501202