

Moore's Law has shaped the tech industry for decades.

It's the idea that computing power doubles roughly every two years, making our devices faster and cheaper over time.

But what if I told you there's now a "Moore's Law for AI agents" — and it's happening *more than three times faster*?

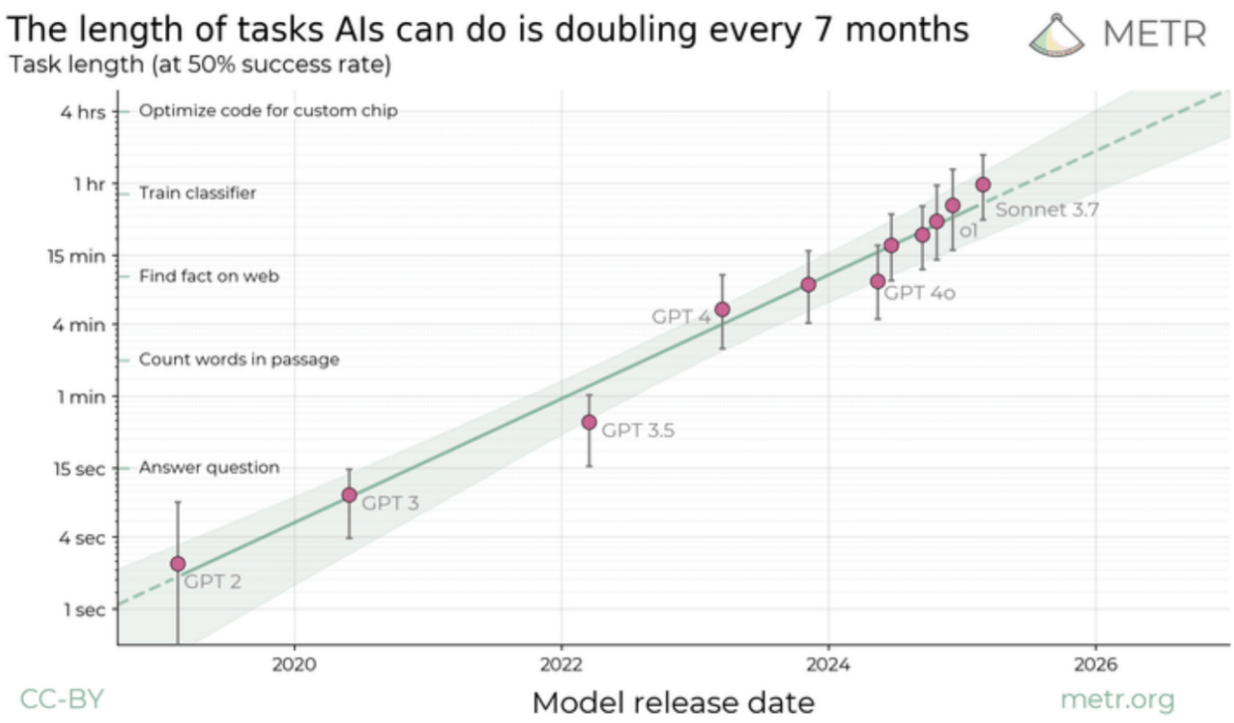
Well, a new study just dropped proving that:

- The capabilities of AI (and agents) are doubling every 7 months (not every 2 years!)
- This trend has been consistent since 2019
- Today's top models (like Claude 3.7 Sonnet) can complete tasks that take skilled humans nearly an hour
- At current rates, AI systems will tackle month-long human projects by 2030

In today's newsletter, I'll explain what this research tells us, why it matters, and what you should do now that you know about it.

—*@dharmesh*

The Exponential Growth of AI Capabilities



Screenshot from metr.org

First, let's talk about what exponential growth *actually* looks like in AI.

Moore's Law (the original one) predicts that computing power doubles roughly every two years, which is why our devices get faster and cheaper over time.

But this new research from METR reveals something much more dramatic happening with AI capabilities:

1. The researchers tracked human and AI performance across 170 different software tasks, ranging from quick 2-second decisions to complex 8-hour engineering challenges.
2. They measured how long each task took a skilled human to complete, then tested whether various AI models could complete the same tasks autonomously.
3. The results showed a clear pattern: The length of tasks that AI agents can reliably complete with at least 50% success rate has been doubling every 7 months since 2019 (see image above).

To put this in perspective, the progress has been remarkable.

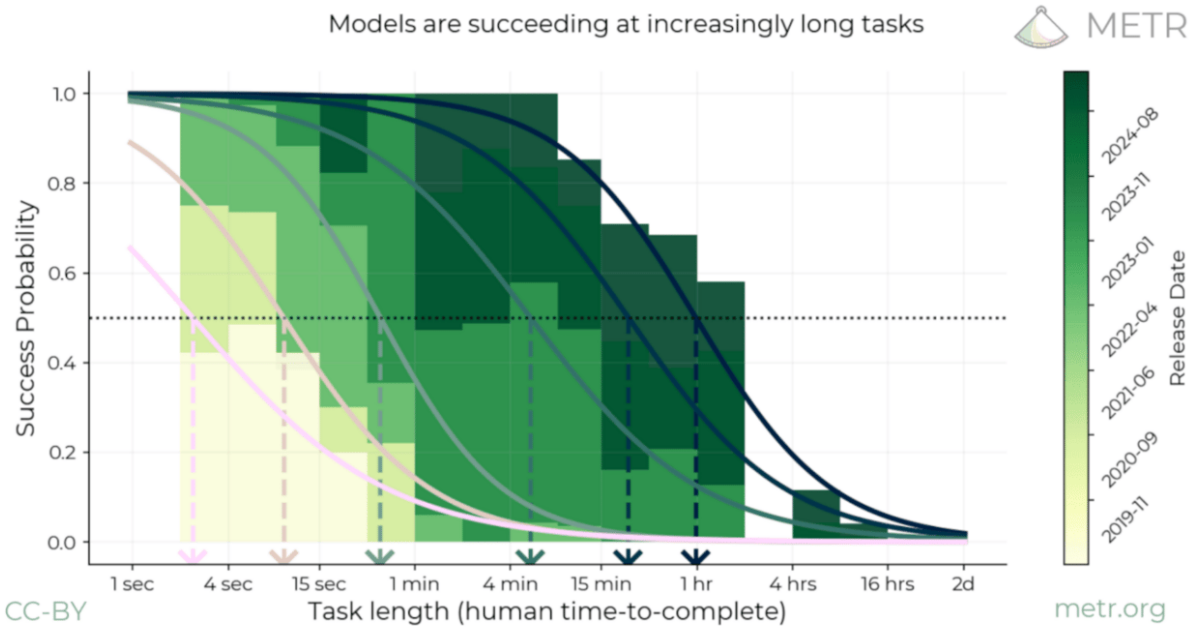
State-of-the-art models in 2019 could only handle tasks requiring humans a few seconds. Fast forward to GPT-4, which can reliably complete tasks that would take humans 8-15 minutes. Now look at Claude 3.7 Sonnet with its "time horizon" of 59 minutes — meaning it can complete tasks that would take a skilled human almost an hour.

If this exponential trend continues — and five years of data suggest it will — by 2030, AI systems will be tackling month-long human-equivalent projects with reasonable reliability.

For those of us building AI agents today, this isn't just interesting data. It's a roadmap.



What This Means for Agent Builders



Screenshot from metr.org

This research beautifully explains something many of us have been experiencing firsthand: the rapidly expanding capabilities of AI agents.

As someone deep in the trenches building agent.ai, I'm seeing these patterns play out in real-time.

Here's why I think this matters:

1. **We're approaching an inflection point**. If the doubling trend continues, we'll soon reach a threshold where AI agents can handle tasks that take humans days or weeks. That's when we'll see truly transformative applications emerge.
2. **The blueprint for progress is clear**. The limitation isn't knowledge or skills (models already know more than most humans), but rather maintaining coherence across longer sequences of actions. This is exactly what we're optimizing for in the agent.ai Agent Builder (the low-code tool for anyone to create and share agents).

3. **Planning horizons predict success**. When an agent fails, it's usually because the task exceeds its "time horizon." This explains why seemingly simple requests sometimes fail - they require longer coherent sequences than the model can maintain.

What I find most fascinating is how perfectly this explains our current reality: agents are incredibly useful for singular tasks but still can't fully replace human workers.

I think the explanation is relatively simple: Most valuable human work isn't about one brilliant insight — it's about sustained sequences of coherent action over time.



What You Should Do About It

If you're building AI agents today, this research should inform your strategy in several important ways:

1. **Focus on what's possible now**. Build for the current ~1 hour time horizon, not for theoretical future capabilities. The most successful agents today solve specific problems that fit within existing model limitations.
2. **Design for extensibility**. Create agents that can grow more capable as the underlying models improve. Your agent that handles 30-minute tasks today will likely handle 2-hour tasks next year without any rebuilding.
3. **Break complex workflows into smaller chunks**. Since current models excel at tasks under an hour, design your agents to handle distinct phases of longer processes. This lets you deliver value today while preparing for tomorrow.
4. **Prepare for multi-agent collaboration**. The next frontier will be agents that can effectively work together, extending their collective time horizon beyond what any single agent can manage. In case you missed it, I shared more on my thoughts about multi-agent systems here.

Most importantly, recognize that we're still early in this exponential curve. What seems impossibly complex for an agent today might be routine in just a few years.

At agent.ai, we're building with these principles in mind, creating a platform where both builders and users can benefit from this remarkable growth curve.

If you've been waiting for the right moment to start building agents, consider this your wake-up call. The capabilities are doubling every 7 months! That means the agents you build today will be twice as capable by the end of the year.

It’s time to build!

—Dharmesh (@dharmesh)



Update your email preferences or unsubscribe here

© 2026 @dharmesh

228 Park Ave S, #29976, New York, New York 10003, United States