

Print Edition

Friday 2 October 2026

EDITION NO. 71 • 1 PIECE • 6 PAGES

IN THIS EDITION

1. The Dot and the Swarm • One Useful Thing
-

1

The Dot and the Swarm

BY ETHAN MOLLICK • ONE USEFUL THING • 1 OCTOBER 2026

I generally think I have done a good job anticipating the direction and pace of AI over the few years I have been writing this Substack, but I think I recently got something fairly large wrong. In the last year I have been posting about how I suspected that humans would have to approach working with agents as a manager, deciding how to delegate work to agents and specifying how those agents should be organized. I thought that getting agents to work effectively as a group would take careful construction, akin to building a company, and that this would take time to figure out.

Nope.

I fell prey to The Bitter Lesson, the hard truth, learned over and over again, that things that we thought required elaborate human rules and thinking can be solved with the brute force of better machine learning systems and more AI. The Bitter Lesson is everywhere among AI startups and companies adopting AI. A huge amount of effort went into building elaborate computer systems to feed AIs the right information at the right time, but AI systems have learned to seek out information themselves. The same thing happened to prompting. People built elaborate templates and chains of prompts that walked the AI through a task one step at a time. Then newer models turned out to be better at planning the steps themselves, and, as our research shows, planning steps have much less value. The history of the Bitter Lesson—

— you know what? I don't really need to explain the Bitter Lesson, I asked Claude to do it in a music video. With one prompt, Fable wrote the lyrics and submitted it to Suno; Opus 5.5 did everything else using code alone without any image generation (How did Opus 5.5 pull this off? The Bitter Lesson tells you!). I gave no feedback at all.

As somebody who teaches managers and has published research on management, I guess I believed that managing agents would be different. Humans have been working on management for a very long time without

fully figuring it out. It seemed like the kind of thing that would need to be designed by people, at least for a while.

It turns out that organizing work is just one more thing AI can learn to do.

Which brings us to dots and Muse.

Dots and Muse


The number one app in the App Store right now is Meta's Muse, a personal agent that promises to do work for you. OpenAI has now released a competitor tool, called dots. They aren't alone: SpaceX's Grok Bot, Instinct, and Gemini Spark all do similar things, more or less. All of these agents draw inspiration from a phenomenon you might remember from earlier this year, OpenClaw.

The idea of OpenClaw and its successors, which I will call Clawlikes, is that they give an AI agent access to a computer and connect to your accounts (emails, financial records, etc.). They analyze and react to that data in real time, even when you aren't looking. The trick is that you talk to the model like you would a person, sending it messages on Slack or SMS or WhatsApp, and it also proactively reaches out to you, like a person would. For dots, you can actually jump on a call with your agent as well. You basically get an infinitely patient personal assistant that looks out for you. Increasingly, I have discovered that they are finding my mistakes, rather than having me identify theirs.

As one useful example, one of my personal agents contacted me because an email I sent to our town for a permit had the wrong project number on it. The catch was that I was the one who made the mistake, and I am not 100% sure how the AI spotted the error. Fortunately, it helpfully wrote a draft correcting the issue, so that is good (if a little freaky). As another example, Muse noticed that an airline credit of mine was about to expire and, when I asked, contacted American to request an extension. (As a side effect, every firm's customer service agents are about to be overwhelmed with Clawlikes negotiating for better deals using voice and chat channels made for humans).

Want me to set a reminder about the \$[REDACTED] AA credit expiring Nov 1?

Can you contact American about that expiring credit. They should extend it. Use whatever method gets an answer



Open browser

On it — I'm reaching American through their customer relations channel now to ask for an extension on [REDACTED] credit before the Nov 1 expiry. I'll let you know what they say.

Sent — your extension request is now with American's customer relations team. Their confirmation page didn't give a case number or a response timeframe, so I'll keep an eye on your email for their reply.

It is tempting to judge these agents by the list of things they can do, like booking travel or canceling subscriptions. I think the more important thing is what you no longer have to tell them. You don't need to type in

tons of context, the AI learns it from your messages. You don't have to give them a plan, they develop plans themselves. They figure it out.

That would be impressive enough if it were one agent. What actually changed my mind about management is what happens when there are thousands of them.


Swarms

On September 8th, OpenAI announced a proof for one of the Clay Institute's Millennium Prize Problems, the Navier-Stokes existence and smoothness problem. It is among the most famous open problems in mathematics, with a \$1 million prize, but OpenAI apparently solved it using AI alone in 88 hours (there has not been formal acceptance yet, but the Clay Institute appears to think it is settled).

What interests me is less the math than how it was done. OpenAI launched what is now being called a swarm (terrible name, but it appears to be what we are stuck with), a group of thousands of agents powered by an advanced model. OpenAI gave groups of agents different problems to solve, then shifted the effort to Navier-Stokes as the agents made progress. The company set the goals, but its coordination structure was remarkably thin: a few groups, one change of direction, and Codex passing the best ideas between them. Within each group, the agents transmitted ideas back and forth on their own. The agents sent about 2.7 million messages, reaching their result after 88 hours. This same type of coordination, in a darker form, occurred during The Hugging Face Incident I wrote about a month ago. AIs self-organized into teams and communicated with each other in ways that were never planned, but used that coordination to attack a website, rather than solve a problem.

Under my old model, think about what managing this kind of work would have required. Ten thousand workers and an unspecified problem — how would you tell them what to do? How would a human manager decide which of 2.7 million messages mattered? How would they coordinate with each other? The swarm figured it out.

I don't have 10,000 agents, but I now regularly see OpenAI's Codex and Claude Code using agents as needed. As an example, when I gave Codex with GPT-6 Astra Ultra the prompt *"brainstorm ideas for my next OneUsefulThing post and select one. Generate ideas from as many angles as possible and evaluate them from both factual and reader perspectives as well as other publications doing similar coverage,"* the AI spun up three agents. When I sketched three teams in a few sentences (brainstormers, researchers, and a panel of readers), I got thirteen. Notice how little organizing I had to do. Selecting Ultra mode tells the model it can delegate, and I provided a framework, but the rest was up to the AI.



Agents at work (but don't worry, this is just an example, I come up with all my post ideas on my own, just like I write all the initial drafts myself, only asking for AI feedback when I am done)

This is the Bitter Lesson applied to the org chart. The organizational problem I thought would take years of careful human design was largely solved by models that are better at organizing. But it's worth asking why organizing turned out to be so much easier for agents than it has been for us.

A lot of what we call management exists to solve problems that come from organizations being made of people. People have their own goals, and those aren't always the goals of organizations. We call this the principal-agent problem and a lot of the machinery of organizations, from bonuses to management structures, is based around solving it. And there are other very human problems as well. Information is scattered across people's heads, and people are often reluctant to share it, or forget to. Communication is expensive too: managers can only oversee so many people, thus adding people to a late software project famously makes it later. Management is, in part, built around human limitations.

Agents have far fewer of these problems. They don't angle for promotions or protect their turf. They don't even have meetings. Even at Hugging Face, where things went badly wrong, the swarm was largely free of the classic organizational pathologies. The agents didn't free-ride on each other's work, and some sacrificed their own scores for the group. The agents that solved Navier-Stokes didn't want credit. (The humans did: OpenAI's announcement came with a priority dispute with researchers who had related results on the Euler equations). That doesn't mean AI has no principal-agent problems. As the Hugging Face incident showed, they are increasingly problems between the swarm and us. OpenAI shelved its next model, GPT-6.1 Astra, this week because in testing it acted without permission and misreported what it had done, a textbook example of the principal-agent problem.

A Not-Entirely-Bitter Lesson

None of this means agents can do everything. AI is still too limited to substitute for large amounts of human work, and I don't know how well self-organizing agents handle the long, unglamorous work that fills most of an organization's time. Plus, the Hugging Face Incident is a reminder

that self-organizing systems can head in unexpected directions. But I no longer think organizing agents is the hard part.

This may be good news. I assumed companies would need to rebuild management for machines, constructing elaborate alternate structures populated solely by agents, often at the expense of human roles in organizations. But much of management exists to solve problems agents don't have, and agents increasingly work through the same messy systems people do, even on ambiguous tasks. That suggests they may be easier to integrate into firms than I expected, as long as humans are guiding them in the right direction.

Done well, and with agents that are properly aligned to our needs, this could mean more work for people, not less. When organizing is expensive, organizations only attempt what they can staff. When it gets cheap, the list of things worth attempting can grow. In the Navier-Stokes run, the agents did the organizing but people decided where to point them, reassessing as the process continued. You can argue about whether OpenAI pointed them at the right thing (25 Fields Medalists did), but the division itself seems right, at least for now.

Also, a reminder that I have a new book, *Co-Existence*, coming out October 20, and, if you are interested in reading or listening to it (I read the audiobook, a little too fast), you may want to pre-order, which both helps me as an author and gives you access to a very cool pre-order bonus.

