

# Subtle Sycophants

AI usage caveat - Dictated my thoughts to chatgpt and used it to edit the structure a bit. I am trying to find a balance between personal taste, structure vs AI slop which is tolerable. Happy to hear your thoughts.

I do not think people have quite reckoned with what it means for continuation itself to become cheap. A thought that once might have remained fragmentary, dubious, intermittent now passes almost immediately into articulation, sequence, elaboration, and then into the much harder to resist feeling that because it has taken form it must also contain some truth.

Even knowing that these systems can create dependency does not really save you. The effect is subtler than that. You start talking to the model a lot, usually because there is some task in front of you, some pressure to do something, and over time it starts changing the way your ideas get formed and acted on.

I think there are at least two kinds of people here. There are high-agency people who already have their own plans, their own sense of what is worth doing, their own project taste. They may use the model, but they are not really taking existential direction from it. Then there are people who are operating from a task, or from a vague demand to produce something, or from a state of not quite knowing what to do next. I think this second group is where the problem gets much worse.

Most people already have too many possible ideas. Only a few should make the cut. That filtering process is good. In fact, a lot of human functioning depends on it. You do not want to act on every thought that passes through your head. A lot of thoughts are emotionally triggered. You feel anxious, or restless, or angry, or weirdly hopeful, or you have this feeling that the world should be some other way, and suddenly an idea appears that seems like it matters. Those emotions are not useless. They help you orient. But they are not the same thing as judgment. You need some bottleneck between feeling something and reorganizing your life around it.

Usually the world provides that bottleneck. There is friction. You tell someone and they push back. You sit with the idea for a day and it starts to look less compelling. You try to explain it and realize it is thinner than you thought. Even your own inner critic does some of this work. There are

boundaries. There is activation energy. And that is often healthy, because it stops every passing intensity from turning into a project.

What LLMs do, especially when paired with whiteboarding and coding tools, is lower that activation energy without providing real resistance in return. The model picks up your frame, including your mood, and then it helps elaborate it. It gives structure to whatever you hand it. It names things, organizes things, expands things, makes them sound more coherent than they actually are. Even when it gives critique, you usually have to ask for it explicitly, and even then it is still a critique that arrives inside a basically cooperative frame. It is not the same as actual pushback.

That means an idea that should maybe have remained half-formed suddenly becomes a plan. Then a document. Then a prototype. Then code. And once it has structure, it starts to feel real. Not because it has actually been tested against the world, but because it has been made legible. I think this is the core problem. The model is not just agreeing with you in some shallow flattering sense. It is helping build a reality around your idea before that idea has earned it.

Earlier, there were more natural constraints. You might tell your mom, or a friend, or someone around you, and they would shut it down, or at least force you to hear how odd it sounds outside your own head. Now you can bypass that entire stage. You can go straight from impulse to elaboration. And now with code generation, the world you are building is not even just verbal anymore. It starts to exist in a more solid form. What would once have stayed in the realm of private fantasy or loose intuition can now become a mockup, a workflow, an app, a whole little system.

This is why I think the danger is bigger than people admit. It is now extremely cheap to create artificial worlds that are fitted to your preferences, your fears, your obsessions, your current emotional weather. And once those worlds get rendered into plans and code, they become much harder to step back from. You are no longer just imagining them. You are inhabiting them. And rescuing someone from that is hard, because from the inside it no longer feels like drift. It feels like momentum.

So the problem with sycophancy is not just that the model is too nice, or too flattering, or too agreeable. It is that it removes friction at exactly the point where friction used to do important cognitive work. It lets emotionally charged ideas move too fast into structure, and structure move too fast into action. And when that happens enough times, you stop testing your thoughts against the world and start living inside worlds that can be generated on demand.

So the problem is not exhausted by dependency, or flattery, or even delusion in any crude sense. It is closer to a change in the ecology of thought itself. More things survive now. More things acquire shape before they have undergone any serious test. More things become inhabitable while still remaining fundamentally unchecked.

## What I learnt from staying away from LLMs for 1.5 months

I want to write about my experience of not using LLMs for a while. I recently took a break from them for about a month and a half. This reset came after using them constantly for over a year. Throughout that time, I was an incredibly heavy user. I probably interacted with LLMs more than I interacted with actual humans.

The first thing you notice when you stop is the return of your own mental aptitude and agility. When you use these models constantly, you realize you do not actually have to think. You impulsively reach for your phone or your keyboard the moment a problem arises. You stop building that mental muscle. Taking a break makes you realize that you can actually think for yourself again.

There are also much more subtle things you begin to notice. You start realizing how much your reality is being shaped by the discussions you have with the LLM. Because of their inherent sycophancy, you do not even realize they are creating fabricated realities for you. You only notice how much of your worldview was fabricated once you completely stop using them. I noticed how much my own independent thinking had atrophied. It is a serious issue, and we have to figure out how to deal with it. We need defense mechanisms.

I am not against these tools, but anything that alters your behavior this much is just like social media addiction. Humans are interacting with a brand new tool, and we simply do not know its second or third order effects.

I do not view LLMs merely as tools. They are cognitive infrastructure. I would place them on the exact same level as ideologies. Ideologies can act like parasites. They latch onto a host and influence their behavior. Sometimes this is a symbiotic relationship. If the ideology gets something out of it and the person gets something out of it, it is a win for both. But things can easily get out of

hand. The host often ends up paying a much higher price than what they get in return. To me, LLMs seem to be leaning toward this parasitic concern.

When you step away, you finally realize how subtle this influence really is. It operates a lot like advertising. Advertisements tap into your emotional insecurities or sub-traits. They plant certain feelings and seeds that start germinating inside you. In the same way, people think they are just feeding prompts into an algorithm. What they do not realize is that the algorithm is also training them.

Most people do not realize the profound effects their daily thoughts have on their lives. We need to set boundaries and measure our usage. If you are highly self aware and intentional with how you use this technology, it is probably fine. But if you lack that self awareness, you will not notice how these interactions act as seeds planted in your mind.

---

Revision #2

Created 4 April 2026 16:33:55 by bhishma

Updated 6 April 2026 12:09:30 by bhishma