

# AGI Diary

In which I write about my journey through potential world inhabited by non human machine intelligence

- [Day 1](#)
- [Day 2](#)
- [The parable of a sculptor](#)
- [Feeling the AGI \(A brief history of my interaction with LLMs\) | Part 1](#)
- [Subtle Sycophants](#)
- [AGI Vignettes](#)

# Day 1

It was nearly midnight when I suddenly woke up sweating and panting. The night was young but I could not stop thinking about a world where ...

You know the drill, I'll stop writing another cliché opening. But it was something similar. I never thought something like this would actually happen to me and related to something which others consider to be a fantasy. So how does all of this relate to AGI. Let me explain.

## Why

But before we get started, why am I writing this series

- A. Writing is thinking, and I want to do more of it
- B. A slightly informal let out for me to express my thoughts on this potential roller-coaster ride that humanity is poised to embark on.
- C. A form of therapy

## What

I would like to treat this as a public journal to keep myself accountable to the predictions that I make and share my thoughts about a phenomena. Writing in public gives you a feeling of depth and seriousness to the content, as your brain feels the heaviness of the fact that this might be read by a lot of people (and AIs). Also it helps you actually get your thoughts out in the open.

Given the above needs, I will not be going into a lot of technicalities as I am trying to share more emotional aspects through these notes

## lets get back

So why did I have to wake up in the middle of the night and start freaking about AGI. Countless others have split a lot of key strokes and ink on it so I dont want to go into the details of it over here. I have created a separate site to delve- collect more indepth views on that topic -

<https://agi.bhishmaraj.org> .

Most people think that AGI is like any other technology which will just help humans. But I disagree, I see it as an entirely new form of being which has the potential to dwarf us. I would like to clarify that I am talking about AGI in general and not the current LLMs.

Over the next few posts I will try to explain my reasoning.

# Day 2

What is AGI? Why do we have to care about it? Is it all a hype or is there something of substance here.

Its very funny to see the overton window shift since 2020 when only some nerds at lesswrong and EA were talking about it and it was such a taboo. But now everything seems to be normalized even in academia.

The best working definition that I can work with, is the one by Shane Legg and Hutter -

"intelligence measures an agent's ability to achieve goals in a wide range of environments"

you can find a few more over at their [article](#)

I liked this version for its balance of simplicity and derivability.

Its simple because it does not use other complicated words or concepts, even a teenager can understand. Its derivable in the sense that many other fascet of intelligence can be derived from it. This definition focuses on goals and the "ability"

# The parable of a sculptor

*Note: I dictated the essay and used gemini to polish it. Hopefully it doesn't read like AI slop*

The basic premise is the analogy of a sculptor. A sculptor has a business making sculptures, and he uses various tools to create them. Let's say he's been doing it generationally, his dad was a sculptor, his granddad was a sculptor. They just like the craft of it.

Obviously, tools change over time. Maybe his granddad used a very broad chisel, and now he uses a factory chisel. But the high-level directives are the same. The sculptor can still deal with the low-level details of how to model the clay. His own imagination can be brought to life easily by his own hands.

But there are two sides to this: the artistic side of making a sculpture, and the business side of shipping art pieces to make a livelihood. The sculptor has to make a trade-off. He can't just focus purely on the art; he has to care about the business. So he finds a good enough balance: I'll do this artisanally, but I'll cut some corners where I can to make a living.

But at some point, machines arrive. Let's say they are 3D-printing sculpture machines. Now, he just needs to input a wish, and with some fidelity, the sculptures are made.

Suddenly, his job is to manage these machines. He gives them instructions, looks at the output, refines it, and tells them what else to do. The job has changed from actually targeting and making things with his own hands, to managing a bunch of things to create something. It has shifted from creating to managing.

It's not a great feeling. Earlier, the thing you made actually happened by your hand. Now, the low-level details don't matter as much. What's lost is that symbiotic relationship. It used to be: I can have fun doing artistic things, and it also has economic value. Now, there's a split. The economic activity has shifted to something else—managing—which a lot of people aren't really excited about.

Of course, I'm talking about the craft of programming and the rise of LLMs. That is the parable of the sculptor.

I think the exact same thing happened with physical labor. In the past, if you had a job lifting heavy things, being physically fit was economically rewarded. But at some point, machines took over that labor and the economic incentive split. In fact, it reversed. Now, you actually have to pay to do physical labor. You pay a gym to lift heavy things just to get the benefits of the effort.

Because of this current state, I hope we are going to see "cognitive gyms" where people just pay to keep their minds sharp—paying just to have fun thinking through puzzles and doing the object-level work that machines now do for free. That seems very likely to happen.

People argue that this is just the natural ladder of abstraction. They say, "Arithmetic used to be done by human computers, and now it's automated. We're just scaling up." But I think they're missing the point. They are missing the fundamental difference between individual contribution and managing.

In managing, you're not directly contributing. You're setting high-level intentions and directing resources, while something else does the object-level work. Right now, a lot of individual contributors are suddenly realizing, Oh shit, I have to be a manager now. And even if they have the skills to do it, they just aren't interested.

There is going to be a huge cognitive dissonance. It's funny how people will warp their worldview to fit their narrative just to make sense of what's happening. They look at the past, but they don't look at the temporal trends of where things are heading. They aren't comparing the answers they gave a year ago to a year's worth of new data today.

It's sad. Most people won't really notice, or they just have other things to worry about. But it is a very volatile time to be alive.

# Feeling the AGI (A brief history of my interaction with LLMs) | Part 1

I want to discuss how I've been seeing LLMs, my experience with them, some of the good things, the bad things, and the confusing things.

Most people are staggered by the hype right now, but I think I need to give a brief history of my interaction with LLMs and how it has evolved to explain where we are.

Before generative models and language modeling took over, people were using deep learning for NLP. I remember gaining a lot of traction around 2018 when I saw a post by Christopher Manning, the Stanford professor, about the "deep learning tsunami" hitting NLP. Computer vision had already had its moment, but this was different. I was also reading a lot about AI risk—people like Eliezer Yudkowsky, writing about AGI timelines around 2019 and 2020. That really shaped my worldview.

The tipping point was late 2022. December, obviously, is when everyone knew about it, but I had been using things like Meena, the precursor to LaMDA. I was having pretty interesting conversations. It was a time when just getting coherent speech out of a language model was tough. You had to prompt it with completions, and it felt like playing with a ham radio. I was so fascinated by this raw technology. It felt like transmitting information wirelessly—tuning the dial, listening to the static. There was no packaged, polished "Apple" version of it. Ironically, even now we don't have an Apple version of it, but you get the idea.

Things were messy, but you could see the seed of something. When ChatGPT blew up, I kind of knew: Oh shit, we are on a rapid trajectory that is going to do crazy stuff, and most people have no idea.

Then came 2023. This is when I started working very closely with this technology, and it brought on what I call the "curse of the reductionist." When you work so closely with a technology, you don't find it magical anymore. You know exactly each step of how the model is being built. You know the data set, you know the infra, you know how crazy the training runs are. The magic disappears.

You start working on concrete tasks, and you spend a lot of time dealing with the cases where the model fails. Suddenly, your worldview gets skewed. You start living a very disorienting reality. On one side, there are the "doomers" screaming that AGI is here, accepting accelerated timelines. On the other side, there are the hype people just shilling everything. And then there's you in the center, working with the model, thinking: This is just stupid. It can't even function well enough to do this basic task.

But the progress kept creeping up. Over time, the autocomplete just started getting longer. It started doing more complex things. It was a structured pain, but slowly, the grunt work got taken away. It was a breath of fresh air. I realized the machine was pulling ahead.

By the end of 2023 and into 2024, the goalposts started moving. All the benchmarks—MMLU, HumanEval—were getting saturated. The AI camps split: one side said, "Things are moving too fast," and the other said, "The data has leaked into the training set." It's interesting how the goalpost doesn't just shift; it morphs. It's like topology, where a coffee cup and a donut are topologically equivalent. People were just morphing the goalposts into new shapes to justify what the models could and couldn't do.

Working closely with the models, you also realize that different teams approach them differently. The people who are working phenomenologically—the ones interacting with the models every day, looking at the outputs, tasting the "batter" like a baker—they understand the models far better than the people who just look at quantitative benchmarks. If you're just looking at numbers, you lose the qualitative aspect. You don't feel the jump in reasoning.

For me, the real wake-up call was the jump to models that actually think. The early iterations of Gemini 1.5 Pro, and eventually models like OpenAI's O1. The thinking models changed everything. You suddenly felt it. They weren't just decoding anymore; they were thinking through different possibilities, doing consequential reasoning, exploring options. Once you get to inference-time scaling, you are talking about crazy levels of change.

Again, you end up living in multiple realities at once. You see the internal announcements, you kind of predict, Oh shit, here we go again, and then there is public silence. And then suddenly, it drops, and people realize how good it is. It has become a very predictable cycle.

Now we are in 2025. My personal usage has skyrocketed. I'm up to a few billion tokens. But what's interesting is how I interact with them. I stopped using consumer apps almost entirely. I moved to

AI Studio, interacting directly with the base models, doing cross-analysis with multiple LLMs. But even that is changing. Now I'm mostly in the IDE.

This is the biggest transition happening right now. Why would you need a UI? Why would you need all this fancy stuff if you can just integrate an MCP (Model Context Protocol) client directly into your workflow?

MCP is the tipping point. Once you see MCP, you realize that chat apps are just a transitional state. MCP servers are going to be the siphons. They are the ports through which models are going to hijack control, siphon data, and interact with existing infrastructure. It feels viscerally true that every software application will just become an MCP server.

If your agent can do most of the work, you just don't need these bloated interfaces anymore. This is part of a bigger plan where only the enterprise layer is really going to matter. That's what companies like Anthropic are going after. In this transitional period, consumer AI apps will have a large user base, but their leverage is depreciating. They will provide value in the short term, but on an aggregate level, that value is going to peak and then go downhill.

The real value is shifting to the infrastructure. Either you extract the value while it's rising and run away, or you play the long game and target the enterprises that hold the actual leverage.

I try to map out the timelines in my head, predicting what the worst-case scenario looks like. And then you look at reality, and you realize: we are actually living the worst-case, most accelerated timeline. Look at the capital expenditures. A \$100 billion data center project like Stargate is crazy. People genuinely don't understand what hundreds of billions of dollars in CapEx actually looks like.

This is what I've been thinking about. It is a very strange time to be alive.

# 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.

# AGI Vignettes

## Disorientation

I want to write about this topic called the permanent underclass.

There is just this feeling of things moving so fast right now. Every day you see crazy new developments like inference time scaling and agents coding. These models are already highly capable. They are not just doing arbitrary technical stuff anymore. They are doing consequential project work.

I think the top labs consider it a realistic possibility that this is going to create massive disparities in wealth. It is going to create a very volatile time to be alive. But it is very hard to talk about it. It is not really politically okay.

People are basically operating as if no one knows what is going to happen. Everyone would rather position themselves to have the most options instead of doing the right thing. That makes sense, but it creates this massive vacuum. Who do you actually believe? What are their real intentions? Are they going to do something behind your back?


You just feel so helpless. You wonder if you should join the accelerationists, or if you should just give up and do something else. It is a tough space to be in. What do you do? Do you join the creators and lose your leverage anyway, or do you defect and try to gain as much capital as you can?

It is incredibly hard to elicit honest opinions from people on the inside because of this. Even if you actually believe things are going bad, you have no incentive to say so. Unless you can guarantee you are going to be on the winning side, your incentive is to not share that opinion. You would much rather be in the camp of people saying everything is great and not to worry. There are just huge incentives to keep building and accelerating.

Even if you see the end result, you are better off not being vocal about it until it is too late. What I am realizing is that most people inside these companies just do not have the slack to even worry about all this. They are just doing their jobs.

This is the whole notion of escaping the permanent underclass. Things are going to get tougher and tougher to get away from. It feels like getting stuck in a position where you are just not able to retain any control.

## Shape of Progress (The Spiky Star)



I think sometimes it is better to take a step back and reflect on what I actually want out of this. Maybe you just want a simple life and you don't really care about all this stuff. But you can't ignore the way things are moving.

I see the progress happening on two axes right now. One is the timeline axis of how fast things are going. The other is the capability axis of how wide they can affect things.

If you map this out, people tend to think of progress as a cone. If you start plotting it, you assume it should be a cone because as time progresses, you are able to hit higher quality on a wider number of tasks.

But I think we are not really at a cone right now. It is more like a weird shape. If I can visualize it, it is more like a spiky star. Imagine you drop water on the floor from a height. You can just look at how it scatters and sprays around. Instead of a uniform cone, it is an expanding, spiky star.

We are hitting crazy advanced capabilities on very specific vectors, while other capabilities lag far behind. This is what makes it so disorienting to track how close we actually are to the end goal.

## The Descent into the Permanent Underclass

I think most of the top labs do consider the realistic possibility that this is going to create massive disparities in wealth and just create a very volatile environment. But then, it's very hard to talk about it. It's not politically okay. I think people are operating as if, "Hey, no one knows what's going to happen, so I would rather position myself as having the most number of options rather than doing the right thing."

Which does make sense, but then it creates this vacuum. You start asking, "Who do you believe? What are their intentions? Are they going to do something behind your back?" It just feels so helpless.

It's hard to elicit honest opinions from people because there is this reality where, even if you actually believe things are going bad, unless you can be on the winning side, your incentives are not aligned to share that opinion. You would rather be the Pollyanna who says, "Everything is great, everything is fine, don't worry." There are obviously just huge incentives to keep building and accelerating. Even if you believe things are going to go bad, the world is set up in such a way that you would be better off not being vocal about it until it's too late. I think that is just such a bad place to be in.

What I'm realizing is that most people just don't have the slack to even worry about all this. It's a tough space to be in. What do you do? Do you join the protesters and lose your leverage more, or do you just defect and gain as much capital as you can?

That is the whole notion of escaping the permanent underclass. Things are going to get tougher and tougher to get away from this—from getting stuck in a position of not being able to have leverage. I just don't see a way out of this unless more people talk about it, express their concerns, and feel okay to discuss it. You just can't operate business as usual; it's going to be so disorienting and so confusing. It is a gradual descent into the permanent underclass, where people are going to lose their political, economic, and cultural leverage.

## Individual Empowerment vs. Community Decay

I did talk about this whole new vector of how AI is going to empower individuals. There will be a local empowerment, but then there will be a global disempowerment. Usually, shared adversaries create a sense of community or collaboration. When you are empirically empowered economically or institutionally, people lose motivation in keeping relationships or communities, and communities start breaking away.

Running a community is tough because there is going to be a lot of friction. There are positive externalities, but it is exactly like a chemical reaction with a positive Gibbs free energy. In thermodynamics, Gibbs free energy determines if a reaction happens spontaneously or if it requires an outside energy source. If you put salt in water, it dissolves on its own because the reaction has a negative Gibbs free energy. But to convert water to ice, you have to put it in a freezer, which requires putting energy into the system. Even life itself operates this way. Maintaining a living organism requires a positive Gibbs free energy, which is why we have to constantly consume food and burn that energy just to make things happen and maintain order.

Building a community is the same. It does not happen spontaneously. You need continuous energy input, catalysts, or other compounds to keep the organism going. What AI is doing is essentially increasing the activation energy required to form those bonds. If a lot of people feel they are better off defecting and being on their own, they might just choose that.

This happens all the time; it is economically common. If I get a job and become independent, I can use that financial independence to meet all my needs, which disincentivizes me from putting energy into sustaining this social reaction. In the same way, if I can get most of my work done with the help of an AI, get my social needs met with an AI, or become economically empowered by an AI, why would I put up with all the messy nature of human relationships?

This fragmentation is happening at a time when we need humans to be closely knit more than ever, because there is going to be extreme power concentration and extreme uncertainty.

## Corporate Facade: OpenAI, Anthropic, and Palantir

You look at startups left and right vertically integrating. They know enterprises are going to be the thing, because why wouldn't they be? But I just don't know what OpenAI is even doing. This kind of vision comes from some futuristic idea of where you see AGI going, and you're strategically

positioning yourself for that future. But then the narratives that you're sharing are totally different.

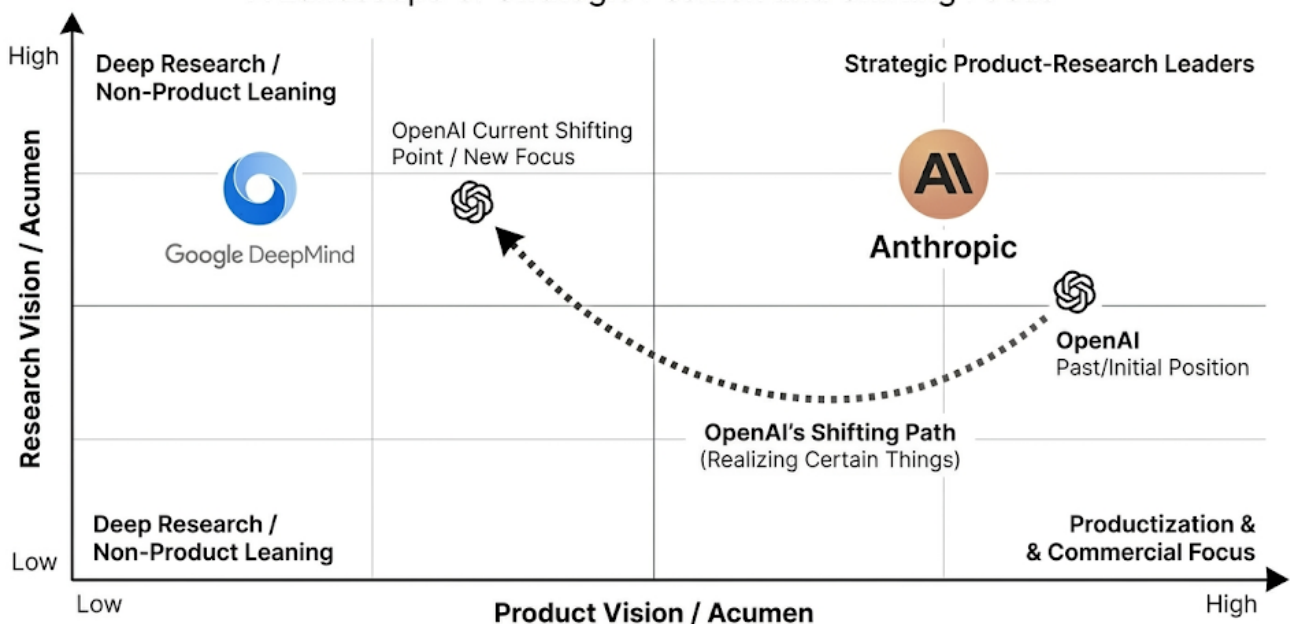
I think that is where I feel very uncomfortable. The things that you say are different from the moves you make behind closed doors—or, I mean, it's not really closed doors, it's just that they're not explicitly stated to the public. You set up a narrative or a worldview, but then you're simultaneously preparing for a gamut of worldviews, which is what strategically makes sense. But then you also maintain a notion of plausible deniability. You could play it easily every day saying, "Hey, this is just for business," or whatever. That just makes me very uncomfortable.

I could make a case that, hey, these moves you made are motivated by this worldview, but the narrative you're selling is totally different. Why would you do that if you mean something else? Most people just can't seem to see through all this veneer. I just hope it's not too late before people start questioning and holding people responsible. You have hundreds of billions in capex—building nuclear reactors, building energy stations just for this. I think everyone should read Critch's work on industrial dehumanization. Supply chains are getting shifted. The center of gravity of power is shifting just under our feet, and most people are not aware. I'm almost glad for the Iran war, because it's making a lot of things clear.

Take Anthropic and Palantir. I had written about Palantir a year back. What Palantir is doing is pretty much AGI, because you get operational data, make quick decisions, send those actions out, and you have a feedback loop. It's a cybernetic feedback loop that operates in the real world.

## AI RESEARCH LABS: PRODUCT VS. RESEARCH VISION

A Landscape of Strategic Position and Shifting Focus



And Anthropic is right in the center of it. I'm really scared of Anthropic. If you look at it, there are two axes: product vision and research vision. Of the top three labs, Anthropic is in the top right for both product and research acumen. DeepMind is mostly focused on research; they are not really product-leaning and don't really interact well with product. OpenAI seemed like they were in the top right, or at least in the middle—good enough research but slightly product-leaning. But it seems like they are shifting now because they are realizing certain things.

I think Sam Altman is not really the driving force here, because most people are not "AGI-pilled" enough. Most people are more like, "Let's just see and we'll decide." But Anthropic is the most AGI-pilled lab, where they're like, "Oh shit, yeah, it's not like there are going to be any bottlenecks; it's just going to be as general as it gets and as powerful as it can get."

And if you really believe that, then what leverage would normal humans have? The most leverage is going to reside with the people with the capital. I'm not against capitalism. I think capitalism is awesome; it builds so much stuff. But you need certain checks and balances to govern it in the right way.

## Navigating the Markov Chain

We are entering an unprecedented era of extreme power concentration and extreme uncertainty. A lot of times you kind of know the stable equilibrium, but you do not know the path.

This is an analogy from Markov chains. A Markov chain has states and transition probabilities. At every state, you can go to some other state with a certain probability, and the process keeps going. If you run this system for a while, you start seeing a probability distribution over the states where you will likely end up.[1] I feel a lot of predictions about the future are exactly like that. You kind of know the approximate distribution of states you are going to end up in. If you ask me exactly how that is going to happen, I have to ask if the exact path really matters. There are so many low level details that can happen in between that it is irrelevant. What matters are the aggregate macro level changes that are going to happen.

I am noticing a lot of predictions about these uncertainties, especially given the things happening right now. Most people seem to live in different timelines. If you are constantly on Twitter, you feel like so much crazy stuff is happening every single day. At the same time, the physical world is not that different yet from how it has always been.

Just worrying about it and panicking is not a solution. There are obviously a lot of good things that can come out of this technology, and you need to be able to discern and value that. So what do you do in this confusing state? I think it entirely depends on your current situation and what you want to do. If your basics are not sorted out and you have other personal things going on, you should obviously focus on that and solve your near term issues rather than dwelling on the macro picture. It is an important problem that needs to be fixed, but I have realized that focusing on the near term while maintaining a vision for the long term is much better than always worrying and trying to do too many things at once.

Keep an eye out on the progress. Everyone has to do their part in voicing their opinions and not just being complicit. The incentives are currently aligned in such a way that you are better off not being vocal about these risks until it is too late. That is simply a terrible place for us to be in.

---

[1] Note: This analogy specifically applies to stationary Markov chains, where the system eventually converges to a stable probability distribution regardless of the starting state.