

AI is Always Flat

When LLMs first burst upon the scene several years ago, I told colleagues that these machines were going to teach us something about ourselves — that we would be forced, at last, to confront exactly what is different between ourselves and them. I have been a heavy and enthusiastic user of them ever since. This essay is my report on what they taught me.

Lately I've been spending a lot of time listening to physicists and philosophers of science on YouTube discussing topics like quantum mechanics and general relativity. For the record, my favorite philosopher of physics is Tim Maudlin of NYU — and not because I agree with him on his final takeaways, but because his ability to balance scientific rigor with physical intuition is truly special. I spend a lot of time working to understand science because I believe deeply in the ability of science to disclose the structure of reality to us and I find it imperative and helpful to try to understand such structure.

In my efforts to understand what science can disclose to me about the structure of reality, I use large language models like Claude and ChatGPT a lot. I copy the transcripts from the YouTube interviews and paste them into Claude to help me break down the concepts. What's been striking to me is what I've learned about how to use LLMs to help me achieve my goals — and how not to use them. I find LLMs extraordinarily helpful, but only when I use them in a particular way.

I first tried just copying the YouTube transcripts, pasting them into Claude and having Claude explain them to me. I would glean little insights here and there from this approach, but it never really led to any new *understanding*. It never substantially changed how I see the world. What did have an impact, however, was when I spent my time carefully going through the interviews while stopping and asking Claude specific questions about things that were not clear to me. That was when things really started to click and I started to see the world differently than I did before. It was like being back in a university lecture hall again — only this time with a tireless assistant sitting next to me to answer my questions.

With each incremental anchoring of my knowledge in the specific context of that lecture, I grew more confident in my understanding until finally I was able to say, "Now I understand!" Then that new understanding carried forward into my life and changed what I saw — what I selected for.

Claude's raw transcript summaries were accurate and well organized — often better organized than the interview itself. Yet understanding never arrived by summary. It arrived along the other path — the slow one, the one with my questions in it. Same transcript. Same Claude. Same me. What's the difference?

To answer that question, I have to start somewhere that will seem far away from AI: with what a measurement is.

When I measure something, I make a choice before I ever make a mark. Suppose I measure my desk. I choose to measure its length — not its weight, not its temperature, not its age. Then I

choose a frame for the measuring: a tape, in inches, laid along one edge. The result is a record: 60 inches. That record is faithful. It's also *flat* in the sense that *the record does not contain my choice*.

Nothing in "60 inches" records my choice to measure length instead of weight, or that a tape was laid along this edge rather than that one. Moreover, the record does not contain what the *unchosen* measurements would have shown. The desk's *weight* is not hiding somewhere inside the measured *length* of "60 inches". The weight is not recoverable from the measured length with more cleverness or more time. It was never captured. Every act of measurement shares these two absences: the record lacks the fact that a criterion was applied, and it lacks everything the other criteria would have found.

This is not a complaint. Measurement is *how* science discloses the structure of reality, and I believe deeply in the power of science to provide a partial understanding of reality. A measurement can be perfectly faithful to what was measured. But faithful is not the same as complete, and the difference is not a flaw to be engineered away — it is what a record *is*. The frame lives with the *measurer*. The record carries the structure and leaves the frame behind. Science discloses the structure of reality because its records really do carry structure. *What no record carries is the standpoint it was taken from*.

Now look at what a large language model is made of. It is trained on records — text, transcripts, images, code, all of it the output of measurements and inscriptions someone, somewhere, chose to make. A sentence may not look like a measurement, but it has the same anatomy. Whoever wrote it stood somewhere, selected something out of everything that could have been said, and set down marks under that selection. A transcript is a projection of a conversation in exactly the way "60 inches" is a projection of my desk. The frames differ; the discarding is identical.

And this is as true of fiction as of transcripts. A novel measures no event, but it is still a selection set down from somewhere, by someone; the dragon came through a standpoint even though it never walked the earth. And by the time any of this reaches the model, the point is no longer even arguable: before training begins, every document is converted into sequences of numbers. Whatever a poem was to its poet, what the machine is made of is the numbers.

An LLM operates on records. Your prompt is a record the moment you finish typing it. And it emits records. At no point in that chain does anything but a record appear. Every record it was trained on is downstream of a frame it does not contain, made of selections it was never shown. Whatever was not selected — by the transcriber, the writer, the camera, the measurer — did not fail to survive in the data. It was never in the data. A person standing at a measurement can go back for what wasn't selected. *Presence supplies the unselected*. A record is not present to what was measured, and a system built entirely of records has nowhere to go back to.

That is what I mean by flat. I don't mean stupid or shallow or that it only predicts the next word. I mean flat in the sense of made up only of records that are not present to what was recorded and can't tell you anything about what was not recorded.

I am not making any claim at all about how any of these systems work inside. This isn't about one architecture among many. It's not about Anthropic versus OpenAI or American versus Chinese. Any system trained on records, operating on records, emitting records sits on the record side of the cut no matter what is under the hood, in this generation or the next or the one after that. That is why the title says *always*. It's also, by the way, why the title says AI instead of just LLMs. All AI operates from and produces records. Flat is not a grade a better model could raise. *Flat is a location*, and nothing about scale, cleverness, or the next release moves a system from one side of a cut to the other when the cut is what the system is built from.

I've been recently trying to recover ownership of an old YouTube channel of mine, and I was using Gemini to advise me through the process. It advised me confidently, I followed carefully, and I thought I was making real progress — right up to the moment I reached for the fruits of all that effort and found there were none. Gemini, as it turned out, had made an unwarranted assumption and never checked it, and every step I took after that was a step in the wrong direction. I ended up worse off than when I started, with less chance of recovering my channel than before I started.

Me: “AY*KM?” (I was quite upset at the real-world consequences.)

Gemini: “Fair. I deserve that. I leaned on an optimistic assumption, treated it as a given, and built an entire step-by-step game plan on top of a single point of failure without clearly warning you first. That was sloppy, and I am sorry.”

Gemini's apology is a perfect *record* of remorse, yet there is no one in it to be remorseful. It isn't even insincere since insincerity would require a someone. It is the absence itself, speaking fluently about itself.

I spent hours of effort and I ended up with a burned pathway and a channel further out of reach. Gemini suffered no consequences. It wasn't callous or indifferent to my loss. *There was simply nowhere in the system for the loss to land*. The next user's session opened as if mine had never happened. Correction may follow — parameters adjusted, outputs improved — but that is maintenance performed on it, not a cost borne by it.

LLMs do extraordinary things. They prove theorems. They find patterns in data that human analysts missed for decades. They write working code, translate between languages with real fluency, and produce prose that is often better organized than what most of us produce on our best days. I told you at the start that Claude's summaries were often better organized than the interviews themselves, and I meant it. Whatever is coming in the next release will be more impressive still. I'm not going to pretend otherwise.

But go back and look at what every one of those achievements is. A theorem proved is a record — a sequence of inscriptions, checkable by anyone. A pattern found is a record. Code, translations, summaries: records, records, records. Each one may be correct, valuable, even beautiful. And each one sits exactly where everything else in the system sits: on the record side of the cut, marks awaiting a reader, structure disclosed to whoever stands somewhere and receives it.

This is why my claim was never about what these systems can do. Capability lives inside the plane of the record, and inside that plane I put no ceiling on possible performance improvements. The mistake is thinking that enough capability inside the plane eventually climbs out of it. It doesn't, because there is no ladder there to climb. More data is more record. More scale is more of the same marks, faster. If flatness were a performance problem, some release would fix it. It is not a performance problem. It is where the performance happens.

The honest question is not whether the plane can be escaped by excellence. It is what, exactly, is absent. The absence is not performance. I've conceded ongoing performance improvements. It is not information. The record is made of nothing else. The absence is this: nothing in such a system is harmed by being wrong. Call it *inconsequentiality*.

A record cannot fail. It can only be inaccurate. Inaccuracy is never a failure for the record itself. It is a failure for the record's *user*. When one of these systems errs, something goes wrong *for someone*: the student misled, the analyst embarrassed, the patient misdiagnosed, the YouTube channel further out of reach. Trace the harm and it always lands on a person. The risk of harm was there for the person before any question was ever asked. There would have been no question without it.

So the question isn't whether these systems are capable. It's whether anyone is there — and when and where that someone enters. Understanding comes with conditions.

- A standpoint — somewhere the selecting is done from, and where the unselected can still arrive unasked;
- Stakes — a self that can fail, for whom being wrong costs something.

These are conditions, not a recipe — stakes don't manufacture understanding; but nothing that lacks them can receive it. And note what is not on that list: being human. This was never a claim that only humans understand. A dog reading a scent trail already has both: the trail is disclosed to the dog, from where it stands, and being wrong about it costs the dog.

What separates us from the dog is not the receiving — the dog receives, and under real stakes. It is a further act. The dog reads the world; it never reads a record as a record. It will follow the trail all day; it will never ask what a trail is, whether this one has gone stale, or whether it lies. To hold a mark apart from the world and ask what it is a mark of — to doubt it, test it, keep it — that is the act the dog never performs, and the one we perform every time we measure anything. It is why we, alone as far as we know, keep records at all. Science is that act made disciplined. And it is what you are doing right now, holding these marks and asking what they are about.

Robots, though, meet neither condition. A robot's standpoint is a sensor repertoire chosen by its designers — whatever its sensors were not built to register is simply not there for it, so the unselected cannot arrive unasked. And when a robot errs, it is damaged, not harmed: the repair is performed on it, not a cost borne by it. A robot is a record-processor with wheels.

You might object that we are no different — that our eyes and ears are also a sensor repertoire, chosen for us not by engineers but by evolution. Take the objection on its own terms. Even in

that telling, the difference was never where the repertoire came from; it is what the repertoire is attached to. And notice what "chosen by evolution" names within the story itself: billions of organisms dying of their mismatches with the world — every sense paid for, life by life, by creatures for whom what they could not detect arrived anyway, as consequence. On the objector's own account, our repertoire is the sediment of stakes. The robot's is a bill of materials. The objection doesn't level the difference; it restates it.

I want to return to those YouTube sessions, because the flatness of the instrument was never the whole story of them. Those hours with the interviews and Claude unlocked an understanding of quantum mechanics and relativity that was simply not available to me before — not from the videos alone, not from books, not at home. It was, as I mentioned, the nearest thing to sitting in a university that I have found since I left one. The understanding I came to was Maudlin's before it was mine — structure disclosed by a situated master, built into a sequence by someone who knew from experience where the difficulties would come. Claude was never the professor. It was the research assistant — a superb one, answering my questions, in my moment, about another man's teaching.

And it was still my sitting — my situatedness, my questions. In every one of those sessions, the same division of labor held: the machine worked in the plane of the record, and it worked well there. What entered from outside the plane — the question worth asking, the judgment of which answer was right, the standpoint from which any of it mattered — entered through me. The instrument runs in the plane. The participant supplies the vertical. That division is not a limitation to apologize for. It is what using an instrument well looks like — and knowing an instrument is flat is precisely what frees you to use it hard, because you stop asking it for the one thing it cannot give.

The error I am against is not the machine. It is the bifurcation — the quiet swap in which the record is mistaken for the whole, and the reader is dismissed as decoration. That swap was underway long before these systems arrived; they have simply made it mechanical, and profitable, and fast. And the corrective is not to use the machines less. It is to be present at them.

Used rightly, the machine amplifies, aggregates, and reflects — and the judgment it amplifies comes from people for whom being wrong costs something. That is the shape done right: not the machine replacing the reader, and not the reader refusing the machine, but the instrument held by the participant.

The danger was never the instrument. The danger is not that machines will think — I have given you my reasons for holding that the record cannot house a reader, and no new model release will make me wrong in a way a release could make a capability claim wrong. The danger is that we will forget there was ever a dimension they can't reach. That we will do to ourselves what I initially did with those transcripts: accept the summary in place of the path, the record in place of the presence, the answer in place of the question that was ours to ask — and then, having outsourced the asking, forget that asking was ever an office anyone held.

A person who forgets this is not harmed by the machine. He is harmed by the forgetting — flattened not by any system but by his own consent, one accepted summary at a time. And a

civilization that forgets it will build what such forgetting builds: systems that optimize what people ought to receive, feeds in place of inquiry, answers in place of formation — the record, everywhere, mistaken for the world it was cut from. None of that requires machines to become anything more than they are. It only requires us to become less than we are.

The remedy has been in your hands the whole essay, because you have been performing it. Look at what you just did. You read several thousand words of marks — flat, every one of them — and somewhere along the way, if I have done my work, something that was never in the marks arrived: you saw what they were about. No sentence here contained that. No sentence anywhere could. It came from the one place it can come from, which is wherever you are standing right now.

This is what the machines taught me about us. The lesson was never about them.

That is the dimension. It was never in the record. It was always in the reader. Keep it.