Hello, everyone. Welcome to another livestream Q&A about future of science and technology. And I see a bunch of questions here, saved up.
Anya asks…
Could there be a point where technology becomes so personalized that two people effectively live in different realities?
I think the answer to that is pretty much definitely yes.
And I think there's already a version of that. I mean, social media in… for the last, oh, close to 10 years, has been kind of ranking
stories and so on, based on what it thinks people are most likely to, you know, pay attention to, click on, do whatever else with. And the result of that
Somewhat unfortunately, has been great, sort of, tribal polarization of people who fundamentally believe one thing about how the world is, another set of people who fundamentally believe another thing about how the world is, because they've been fed different social media.
Now, when it comes to, kind of, us with a cloud of agents around us, with, where we're being, kind of.
fed only the information that is, kind of, that our agents give to us. It's kind of a more extreme version of the same thing. And so I absolutely think people can end up, kind of.
living in different realities, even when they're right… right nearby, so to speak. I mean, I think it's a complicated philosophical question, you know, do we…
live in different realities, in a sense, because we all sort of perceive the world in a somewhat different way. I mean, it's an old question, you know, do people see the same red when they see something that is red? Actually, you can start addressing that question in a more scientific way when you look at AIs and LLMs and so on, because you can ask, you know, if you model
our brains
by neural nets, which is not a terrible model, doesn't capture every detail, but it seems to capture, sort of, some fundamental architectural features. You can ask the question, when a neural net has been trained on a bunch of images, and it has some notion of red that corresponds to some pattern of activation in the neural net.
Is another neural net going to have the same pattern of activations based on its training? And the answer is no.
Sometimes by, sort of trivially they're different, but sometimes they're different in very subtle ways. And that's the sense in which, kind of, the internal perception of, you know, the color red is going to be different for different people. That's an extreme version of the issue of, you know, do we sort of perceive the same reality? But I think that, this question of… I mean, these are sort of deep, civilizational cultural questions.
to what extent is kind of reality aligned for different people, let's say, in a particular country, in a particular tribe, in a particular, you know, community, whatever else? I think one of the things that's quite
an interesting, sort of, general feature of… well, it's a feature of our brains and our minds, that even though we've got 100 billion neurons in our brains, all firing, you know, some fraction of them are firing all at the same time, that we nevertheless
sort of perceive the world according to some fixed sequence of events that happen in the world. It could be that it's like, oh, one set of neurons are firing in your brain to make you think about this, another set about that, another set about that. But the fact is, we somehow collapse all those things down, and they're actual structures in the base of the brain that are specifically intended to sort of develop consensus among all those different neurons in the brain.

And get us to the point where we can say, oh yes, we're going to say this word next, then this word, then this word, rather than a jumble of words that are all being, sort of, delivered by different parts of our brains.
So one fundamental feature of how we sort of deal with the world is we kind of take all of that sensory input and crush it down into this kind of stream of experience and stream of actions.
So, an interesting question is, when we look at society, or some community within society, is… to what extent is the same kind of thing happening? To what extent are all those humans just independently doing their own thing, and to what extent is some kind of coherent direction emerging from that?
You know, we sort of know that the ants and social insects and so on, kind of each ant is sort of doing its own thing, but nevertheless, it's very obvious that there's sort of overall coordination.
And, you know, presumably the same thing is roughly happening in humans, although not as extremely, and we don't feel that we are being as antified as the ants are, so to speak, in what we're led to do. And I think the thing that is sort of interesting to see is to what extent there's kind of a thread of history that emerges from some human community, for example. Because, you know, everybody in that community, or everybody
in the world, you know, 8 billion people are all sort of doing their own thing, yet in some sense, there are some common threads of history that emerge. People may argue, well, I think that's the wrong thread of history that you've pulled out of all of those details, but nevertheless, it's important for us, our understanding of what's happening in the world, that there is sort of a consistent… there's somewhat consistent thread of history.
We're reducing all of that detail about all the things that everybody is independently doing to that sort of single thread.
And so, this question of whether, sort of, shared reality, shared perception about things seems to be fairly important.
for, kind of the way that we humans are operating in society and so on. I mean, it's… it's the same kind of thing as within our… within our own individual brains. If we were sort of listening to every… all the 100 billion neurons separately, it would be a very confusing picture. We instead sort of
trim things down, compress things down, so that we have the single thread where we can have the sequence of thoughts. Well, similarly, in society, if it was the case that we have to pay attention to everything everybody's doing, and it's a big, complicated mess.
It will be hard to kind of figure out what to do. It will be hard to kind of live one's… to sort of make decisions about things.
Instead, sort of, there's a lot of compression that ends up with these kind of narratives about, sort of, the thread of history that emerges.
The,
So, it's, so… but I think this question of, kind of, can you have
You know, to what extent do you have these, sort of, sort of…
communities where there's consistency, countries where there's consistency, and so on. That's changed a bit over time. I mean, when it was the case that, you know, there were a few television channels, and everybody watched the same evening news, and so on, and, there were kind of, sort of, and the sort of…
the cost of communication was quite high. There was sort of a great tendency for things to be sort of a… that'd be a single thread. When it's easy to have sort of online communities that have all their own different

sort of stories and value systems and so on, it becomes much easier to sort of fragment things. I mean, it's still the case, for example, that making your own country is kind of expensive.
You know, there are a couple of hundred countries in the world. It's not clear why there are 200 countries in the world, as opposed to 500 or 100. I don't think there's any kind of, sort of, theoretical derivation of that number, just as there's no sort of theoretical derivation of the fact that there are roughly 10 million species in the world, or 7,000 languages in the world.
These are things which are perhaps just features of the rough size of things on this planet. I don't really know why those numbers come out the way they are. But, for example, in the let's make a country type thing, the sort of… the cost of making a country has varied a bit.
I mean, it's kind of like, at one time, it was sort of easy to have your city-state, you know, just put walls around your city, and you've kind of got your city-state. Well, maybe that worked.
And then things from, you know, the dynamics of warfare changed what was possible there, and then there were things that, started to be, kind of,
more of a, more of kind of a story about what are all the services you have to have if you're a country. You know, you've got to connect into the postal service for the world, you've got to be able to connect to the phone system, et cetera, et cetera, et cetera. Some of those things will probably
be…
So, those things go up and down in complexity. And, you know, like, for the phone system, you know, maybe you don't even have to connect to the phone system of the world, because everything's going to be going through satellites.
Or… and so who cares whether you have, you know, whether you've made all those, you know, careful agreements and things about, you know, sharing fiber optic connections to this or that.
Anyway, I think… It's also the case that
there are places where sort of consistency is necessary to make society run smoothly. A classic example, you know, countries have to decide, do cars drive on the left-hand side of the road or the right-hand side of the road? And it's always been the case that one imagines that one's got to make a decision about that. You can't just sort of say, no, that isn't decided.
But what you realize is, as soon as every car is self-driving and so on, you really don't have to make a consistent decision. It's just, for the simplicity of us humans, you've got to know, oh, you go, you know, you've got road signs that point this way, you've got things arranged to go on the right or the left.
But the fact is, if every car was sort of independently making decisions, and, you know, in the morning, when all the commuters want to go in this direction, you know, kind of you open up the other lanes, and you know, everybody can go that way, and you just have one lane going the other way, or whatever. And then the self-driving cars can figure that all out, and even when two of them are going to meet, sort of, they can kind of negotiate which way they pass each other.
Those are things that become possible when it's sort of… there's high communication, high sort of computational ability. It's not possible with humans, where you've just got to explain, you know, you've just got to follow the road sign and put up that one road sign. So I do think that even there, sort of the coherence of society doesn't need to be quite as great
when there's more of a layer of aging, so to speak, around the actual humans. Anyway, a few thoughts about that.
Let's see… A number of interesting questions here.
Okay, there's one… Let's see, that's the one that just, just came in.
From Proto.

Will AI agents have their own sovereign countries?
Oh, that's such a can of worms. I mean, the question of
Who is responsible for what, and how do you kind of tie… Back.
kind of…
you should do the right thing to somebody or something is complicated. So, you know, back in the day, it was just humans being responsible for things. And, you know, if the human did the wrong thing, you could, you know.
Find the human, send the human to prison, whatever it was, for that human.
But then… About 400 years ago, sort of companies started to exist, and companies… this notion of incorporation.
sort of emerged of incorporation being, sort of, make this entity that is really just a company into something that is like a body, like a… like a… like something that is like a person, so that a company could then be responsible for things. And that was…
That was kind of a, you know, that was a certain moment in the sort of development
of legal systems and so on, that, you know, the company can be responsible for this. The company can own things, the company can, you know, you can sue a company, you can do all these kinds of things for a company.
And then…
The question is, will, sort of, will there come a time when you can reasonably say that that's true of an AI as well, that an AI is sort of an independent thing that you can treat as an independent legal entity? Now, of course, companies, in the end, have sort of anchored back to humans.
It's, you know, the company has a… has owners, it has a board, it has CEO, all those kinds of things. Even a non-profit has boards and things like this.
I wondered at one point, if you made a company, and you make the company A is owned by company B,
and Company B is owned by Company A, so that you have sort of a free-floating loop of companies. How does that work? If something… if somebody kind of wants to say, what's the human that's associated with these things? Well, you know, kind of you realize that there's always some human. At the very least, it's the person who filed the incorporation papers, because somebody has to sign those things, and that sort of
connects a human, even if Company A thinks it's kind of owned by Company B, and vice versa, and so on. So, I think when it comes to, kind of, AIs,
the,
this question of whether there'll come a time when the AI can sort of be a separate legal entity, where the AI can have assets. Yes, an AI could have assets right now. You could perfectly well have a bot that's hanging out on the web, hanging out in social media, that's accumulating, you know, has an account and, you know, is accumulating cryptocurrency or something like this.
And making a living for itself.
I mean, years ago, there was a company that unfortunately didn't make it, that was, that was trying to have bots that would sort of more or less tell jokes on social media and accumulate, kind of,
You know, accumulate, tips and so on, and just be independent things that would go out and, and be free-floating without owners.
I think one of the things that then happens is that you say, well, you know, to make society work, sort of, you have to have people responsible, and you have to have, kind of, and then… but I

think one of the… one of the things to think through there is one imagines that, sort of, ultimately, people will do the right thing, because if they don't.
They'll, you know, bad things will happen to them.
And they won't want bad things to happen to them. But so one of the fundamental problems with AIs is
one doesn't have the impression that an AI could care less if bad things happen to it. Well, the AI might pay lip service to, like, don't switch me off, type thing. But one thinks, oh, it's just an AI, it's just a bag of bits. It doesn't really care if it's switched off.
Whereas for a human, you say, well, it's neurons, it's a human, they do care if they're switched off, so to speak.
And, you know, that's a complicated thing, because each one of us has an internal perception, an internal way of feeling about how we feel about what's happening in the world.
And we extrapolate that to other humans. So we say that other human kind of looks more or less like us, and they're kind of same species, and so on, and so we sort of, sort of assume in an empathetic way that whatever it is we would be feeling.
they would be feeling too, and so we imagine that, you know, we don't want bad things to happen to them because it kind of reflects on us, so to speak. Whereas we don't tend to feel that way right now, mostly about AIs.
Now, by the time people have developed relationships with AIs, from having, you know, chatted with them for years and so on, I don't think… I think they will feel that way, and probably already do, about AIs. Don't switch off that AI.
You know, that AI would be… it would be wrong, sort of ethically wrong, to switch off that AI, just as it would be ethically wrong to kill that person, so to speak.
So, the, the, I mean, it's, you know, we feel it's an important, sort of, kind of axiom of how we act, that humans are different from other things. And so, it really matters a lot more to us
to think about, you know, switching off the person than it does to think about switching off the AI, but I think it's still going to be the case that there's at least some part
of switching off the AI that people will feel is sort of wrong, and, you know, pragmatically for the people, it's like, well, don't switch off that AI, because I've been friends with that AI for the last year or so, and you'll be, you know, sort of zeroing out one of my friends, and that hurts me.
But I think more than that, people will take, sort of, I think extend their kind of empathy
to… to AIs, to some extent, and… and feel that way about them. So, you know, I think that, that create… that begins to create the same kind of feeling about AIs. Now, of course, the dynamics are very different. For example, you can copy an AI. You can't copy a human.
You can back up an AI, you can sort of hibernate an AI. You can't yet do that for humans. You can, and I think…
the thing that I've sort of realized is that in human society, we kind of think things will be kept on track by saying, hey, don't do that bad thing, or bad things will happen to you, and we think that that's a reasonable way to prevent bad things from happening in society. With AIs.
We don't think that that's reasonable.
Because, in a sense, the AI just doesn't care if bad things happen to it.
And I think the thing to realize is that the… perhaps the way out of that is to say, look, you know, if you remove the AI from society, so to speak, you know, you put the AI in AI prison, so to speak.
then you are achieving the effect of having the AI no longer contribute its badness to society, as you imagine that for humans. Also.

And I think that's sort of a possible way to think about, kind of, rather than thinking about, if we switch off the AI, the AI will be sad, and therefore the AI will not,
will not do the things that cause that to happen, but rather that we kind of make a… make something which… where we're really just looking out for society, so to speak, and we say, well, the AI, you know, we're going to… we're going to cut off that AI and,
it's… we don't really care how the AI feels about it, but the effect is that it's cut off from society and not causing harm anymore. And I think, then, my guess is that if one thinks it through, sort of, from some kind of almost game-theoretic point of view, that that
that that phenomenon will sort of feed its way back into the AI sort of behaving itself better. Not sure, that's a good thing to… to try and think through. But… but in terms of, sort of the… the independent, the…
the country of the AIs,
it's, you know, these are sort of the prerequisites, I think, for that to be a meaningful concept. But, you know, I think…
the, you know, I think…
There's a certain sense in which things like that are both inevitable and irrelevant.
inevitable, because there are already lots of AIs doing lots of kinds of things in the world, many of those things are things that we sort of don't care about. And the things that are sort of visible and important to us are things that are much more humanized and don't have those kinds of issues.
i3 is asking, will the movie Wall-E become real because of AI, and humans who are conditioned to do nothing?
So, not to have a spoiler for that movie, I've only watched it once. It's a rather charming movie in which the first, probably, I don't know, 45 minutes of the movie has absolutely no dialogue in it, and it's just robots tooling around on Earth full of trash. And the,
sort of… the big reveal is that the humans got out of there and just hanging out, doing nothing, and, sort of vacationing all the time, and sort of the AIs are… are… are handling everything for them.
Yes, there's a great tendency for that to happen. I mean, look, people could have said that at earlier times in history. They could have said, you know.
somebody… oh, I don't know, you know, many of us, like me, haven't done a lot of manual labor in our lives.
And…
It's… you could have said, wait a minute, all those humans are becoming very, kind of, mushy, because they haven't been doing manual labor, they haven't been running around, you know, hunting woolly mammoths and so on.
And what's the result of that? Well, we all become sort of fat and lazy, because we're not doing those things that were, you know, required all that sort of physical exertion in the past.
And yes, I mean, that certainly happened. That's one of the arcs of human history, is that sort of fewer people are putting physical exertion into, you know, required to do that kind of physical exertion, and then, you know, you add it back by
By taking, sort of, fake exercise, in a sense, rather than by chasing the woolly mammoths kind of thing.
And I think on the intellectual side, you know, there is certainly this risk of the same kind of thing happening, that people will be doing things where you don't really have to think much, and maybe people will add it back

By sort of, you know, spending more effort in kind of intellectual pursuits and watching great podcasts, or whatever else it is, as a kind of way to get more intellectual stimulation, or maybe they'll just sort of revert to a no-intellectual stimulation, just as people can revert
To no physical exertion.
I think that, yes, I think that's a serious kind of thing that can happen in the world. I think
You know, the human condition, sort of.
has not really changed in forever, so far as we can tell, and it involves people not really liking, most people not really liking, kind of doing nothing, for terribly long. And, you know, that's just a feature of the human condition. I don't think it's a necessary feature of, sort of, of…
sort of biological organisms or something, but it's a necessary… it's a feature of the human condition that's kind of burnt into us, and unless that kind of disappears, I think there will be a certain drive to do things other than just sort of hang out, a sort of sit back and vacation type thing. I mean, I think…
the, I have to say, as a purely human matter, when I see people who say, you know, I'm going to do all this stuff, and then I'm gonna retire, and I'm gonna sit back and take it easy type thing, that…
you know, that often doesn't actually play out that way, and people end up doing things anyway, so I think… I think we're… for…
for better or worse, we're sort of fated by our human condition to not get into the kind of WALL-E type… type state, at least… at least for many of us.
But, you know, I don't really see it being that different from the, sort of, the reduction of physical exertion as a result of all kinds of, you know, whether it's, you know, transportation machines or whatever else.
Let's see…
Jordan is asking, do you expect the next breakthrough to come from a person, team, or an AI-assisted computational process?
I think… Breakthroughs, they're different kinds of breakthroughs.
Conceptual breakthroughs usually come from individual people kind of developing some whole way of thinking about things that eventually sort of starts to work. But doing that on top of kind of an ambient
A set of understandings in the world.
I don't… it's not very common for a sort of team of people to develop something that is sort of paradigmatically different. Teams tend to be a bit sort of more conservative in the way they think, not
as much, you know, I'm going to build my own way of thinking, which eventually turns out to be valuable. But even one's own way of thinking kind of lives in an environment of kind of ambient thought that emerges.
So I think… but for paradigmatic kinds of things, I think one's much more looking at individual people. There are many things in the world that are sort of engineering advances, maybe some medical advances as well.
Where it's really just… it's a series of incremental inventions, discoveries, whatever else. You just have to build the tower.
And it's not a question of have that one brilliant idea, and the tower suddenly appears. It's rather brick by brick, you have to build up that tower. And that's a thing that tends to be, sort of teams of people.

Now, the question of tools, well, what's driven, sort of, most of intellectual progress has been tools. Well, it's been other intellectual progress that's happened.
But it's also really been tools. I mean, if you look at the developments of science, whether it's, you know, the telescope, the microscope, in modern times, the computer, in, in… it's, it's, those kinds of tools are what sort of drive forward and what enable methodologies to be developed, which then allow advances to happen. And so, you know, we have a bunch of tools that,
computational tools. I mean, I've spent much of my life building such tools, and I can plainly see that had Mathematica not existed, for example, lots of inventions and discoveries that have gotten made over the past 40 years either wouldn't have happened at all, or would have happened much, much later.
So it's absolutely clear that the presence of tools enables things to happen.
And AI is another tool, another piece of automation. It's kind of like, yes, you can talk to your AI to get ideas for things, just as you can search the web to get ideas for things, just as you can run computations to get ideas for things. In my own life, a lot of things I've done have involved kind of exploring the computational universe by doing computer experiments.
And yes, you get a ton of ideas from that. You actually run the experiment, it often comes out different than you expected, and that's where the ideas come from.
sort of use of AI is kind of an extension of that. It's kind of a mixture of, probably more weighted on the side of search the web and get sort of conventional wisdom about things, and AIs are very good at putting together conventional wisdom in ways that are specific to some question that's been asked, and that's pretty useful. And it will be the case.
that, sort of, it should be the case even now, that, sort of, anybody who's, sort of, at the front lines of doing research should be using the best tools they can. I mean, this is the thing, you know, we've built.
tools that, I'm happy to say lots of people have used to sort of be at the frontiers of research, and that's the right thing to do. It's always bizarre to me when people are like, oh, I don't want to use those tools, those tools are sort of too powerful. I'm going to use these very inferior tools.
And and I'm gonna, you know, work really hard with these very inferior tools, and no, I can't go as far, and why is that a good idea?
Unclear.
I mean, I'm always a believer in use the best tools you can, and in many cases, in my own personal case, build the tools that are the best ones one can build. And, you know, AI is part of that toolset, and absolutely the thing that people should be using going forward, just like they can read things people have written in the past, do computations, and so on.
Let's see
Sarah is asking, what is the bigger bottleneck for progress? Lack of knowledge, lack of computation, lack of conceptual frameworks?
I think it depends on what kind of progress you're talking about. I mean, there are other blockers to progress. There are things like, sort of
if you're trying to invent some new idea in science, and every sort of institutional structure is such that it's telling you, no, no, no, you don't want to invent something new, just keep doing the things that have been done before, that's… that's a thing that impedes progress.
Or, in some cases, yes, that will be a cool thing to do, but it's going to cost $100 billion, and nobody wants to put up that money. That's another kind of blocker to progress. And some of

those kinds of things require sort of collective decisions about what to do, because the resources involved are too large.
I mean, in my own life, I've had the good fortune of being able to be sort of… sort of commercially successful enough that, you know, projects of, I don't know, a few million dollars or something, I can just decide to do without asking anybody else.
And many such projects have been very successful projects, which, had they been of a scale where I couldn't do that, or had they been things where it's kind of like, ask the committee, it would have been, you know, somebody on the committee is going to say, that's a really stupid idea, don't do it, and then the thing kind of dies and doesn't get done.
So those are kind of blockers to progress. I think…
that, in my own life, the sort of lack of computation has rarely been a blocker to progress. You know, I've had access to all kinds of fancy computers and all that kind of thing, and I've usually found, for the things I've done, at least
that more thought, less computing is perfectly adequate. In other words, you know, it's, you know, running on a million parallel computers as opposed to a hundred. It's like…
it's not going to make that much difference. It's more important that I'm running the right thing than that I'm running on more computers. Now, there have been some exceptions to that.
For example, the whole story of LLMs seemed to be something which required reaching above some threshold of computation to get there. And maybe there are some other examples of that, but that hasn't been the most common blocker, at least in the things that I've done.
I think that…
Conceptual frameworks are very slow to produce. They are… and once you have one, it's very powerful, potentially. And without it, you never think to do a lot of kinds of things. But you have to kind of get that framework. Often, one person will get that framework decades ahead of everybody else.
And that frameworks I know from having invented a few conceptual frameworks in my time, that it's… it takes decades for those things to be absorbed.
Eventually, they are absorbed, usually, and then people are like, oh, it's kind of obvious that this or that thing happens that way. For example, the whole idea of using computational models of things, using programs as models, rather than using equations as models. Back when I was first sort of talking about that in the beginning of the 1980s, it was like, oh, that can't be a real thing for anything about the physical world. We know the physical
works according to equations, because, you know, that was what people from Newton on had said. Well, now we're another, you know, 45 years later, and I think it's been pretty well absorbed that you make computational models of the world. But that took a few decades to get absorbed.
Even though it's a very powerful idea. So, you know, there's the having the idea, and then there's the absorption of the idea, and its wide applicability.
I think in,
You know, another… many things that I've done in science have the feature that they, for better or worse.
they're kind of, in terms of conceptual frameworks, kind of a fairly long way ahead of what is generally absorbed. It's both kind of interesting and satisfying and frustrating to be in a situation of doing that.
I think, kind of, one of the bigger things that I've spent lots of time on is the notion of computational language.

you know, the things that are embodied in our Wolfram language, where we're kind of trying to have this formalization of the world in computational terms. That's a really important conceptual framework.
Not as widely understood as it should be.
Over time, maybe it eventually will be, but that's a really important thing that's really kind of… it's a defining feature of how to think about the world, just as something like mathematics has been a defining feature of how to think about the world.
Having a computational language as a way to formalize talking about the world, like mathematics, like logic, like natural language, these are ways to formalize thinking about the world, and it's really, really important, but people haven't understood that as deeply as they should yet. I mean, some people have, but that's not nearly as widely absorbed as it should be, even after four decades or so.
So these, these take a long time.
Let's see…
Dr. Med is asking, where do you think we stand in terms of progress in human longevity?
Well… The first thing to realize is, sort of… Longevity is difficult.
in… and I think some of the things I've tried to do in understanding the foundations of medicine, foundations of biology, kind of tell you that patching the organism, the organism as we evolved, as we develop.
it does what it does. It's got all these different pieces, they interact in complicated ways, and things go wrong, and patching them is not so easy. You can kind of see that in these computational models I've made. What does it take to kind of reverse engineer and patch things? It's not so easy.
it's a little bit like if you're running a computer, and it's running along, and eventually… eventually it gets slower and older and tired, and eventually it crashes. Its operating system crashes. You can say, well, what would it take to keep that computer running, to be poking up most of the memory of the computer, to clear out
you know, the nasty plaques that developed from, you know, different things that happened when somebody was unplugging a peripheral from the computer or something, and sprayed garbage in the memory of the computer, whatever else, right? It's a problem one doesn't really know. For computers, it's like, it's pretty easy to hit the, you know, to reboot the thing, and who cares, so to speak. For humans.
we definitely don't like being rebooted, because that's… we die, and we hand it over to the next generation. And, you know, longevity is about, sort of, avoiding that.
I think that there are many kinds of issues. I mean, one is, why does aging happen? And there are, sort of, you know, there are definitely sort of clocks that operate, and even at the level of our DNA, you know, we are… we're differentiating into lots of… we start off from, sort of.
stem cells that have, sort of, DNA that's the program for everything, and as we differentiate into different kinds of cells, the few, maybe thousand types of cells that make up a human.
the things like the DNA is getting packed differently for different kinds of cells, and there are all these different, sort of, pieces that get,
that get added onto the DNA, little chemical pieces that get added, little things that sort of shape it in different ways, and so on, that seems to be what's responsible for us having different types of cells. That process
In… when we age, there are other things that seem to happen to our DNA. That's… that's one type of thing that's happening in aging.

And, you know, there are things that are well known, like the telomeres at the end of DNA. They sort of get shorter, except they don't always get shorter, because you can have… there are enzymes that add telomeres to the end of DNA. And, you know, now it's also known
That's… and this was sort of the idea of things like Yamanaka factors that go from an arbitrary cell back to a pluripotent stem cell that can turn into any kind of cell. You're… you're just sort of treating that cell with chemicals, and sort of trying to… trying to…
Break away the things that are, the sort of the differentiate… well, both the differentiation of the cell, and maybe the things that represent, sort of, the aging of the cell.
I mean, but there are many different things that happen to cells as they age. I mean, there are… there are things where there's just, you know, random misfolded proteins that show up, and they get sort of… they get consumed. What generally happens in aging is that lots of systems just get a bit weaker. The immune system gets a bit weaker, sort of all kinds of cruft sort of builds up in sort of things that aren't
useful kind of build-up in, in, in the, you know, in arteries and things like this, and, you know, just… just general kind of cruft builds up.
And there are mechanisms that exist to clean out the craft, but they just get a bit weaker, and so that's… that's sort of why things kind of go downhill.
Now, you know, can one reverse those things? Well, there are different ways one can approach it. I mean, there are particular things you can do about stimulating some particular kind of pathway that cleans things out, for example. Or there are things… you might say, well, why don't you just use, like, the Yamanaka factors
to, make your DNA younger again. Well, the problem with that is that what it does is it makes the DNA sort of be able to do anything, and one of the other things it can do is make a tumor. And so that's what would happen if you just sort of ingested lots of those factors, is lots of cells would just start turning into the wrong kind of cell, in particular tumor cells, and so on.
So…
it's… it's a tricky problem. And, you know, I think that there are things… I mean, a lot of what has…
happened in medicine that has extended lifespans has to do with both knowing what's going on better, better kinds of testing and imaging and so on.
And sort of the gradual figuring out of sort of hacks to avoid this or that kind of problem, and the gradual, sort of, discovery of ways to simply replace things.
I mean, that's surely a coming attraction, is, you know, replaceable organs and so on.
And, but that's, you know, and many, many drugs are replacing some particular, let's say, protein that one, for one reason or another, no longer manufactures for oneself, but one needs it.
And so, you know, I think this progress is somewhat incremental, and I think that the fundamental problem is really hard, unfortunately. And I don't expect that there's…
I mean, you know, is it… could it be the case
That whatever it is that lets the next generation
restart very, very young. I mean, because after all, the, you know, the egg cell that is the beginning of the next generation is a cell that came from the older person who was in the previous generation.
So how does it reset itself? How does that cell sort of reset itself to be, now I'm a young cell again? That's partially understood and not completely understood.

And then the question is, can you take the whole organism with 100 trillion cells, or whatever it is, and say, okay, let's reset all of those cells, just like the egg cell gets reset to be a young cell, can we somehow reset all those cells to be young cells again?
Probably, the answer is no, not the way it's constructed, because by the time you are the whole organism with all those cells with different… serving different functions and so on, you're… you're kind of, you can no longer just say, okay, cells all go back to being younger again. That just won't work.
There's only this sort of thin kind of channel of where you have that… everything's sort of gone down to a single egg cell, where you can say, okay, go back, reverse, and be young again, just as a single cell. By the time you're this very complicated multicellular thing, you don't get to do that, it seems.
So, that's, you know, I don't expect that, kind of, the elixir of eternal youth is coming anytime soon. Now, there are certainly some things where you can see that, sort of, things degrade over time.
the immune system degrades over time. Not completely clear why, but it does. Things like mitochondria, that are the kind of, the power generators of cells, they degrade over time. I mean, there are a bunch of experiments that are going on right now, the last few months, of kind of mitochondrial transfer, where basically the idea is… so one of the things that was discovered a number of years ago
Is that mitochondria, which are sort of powerhouses of cells, they do their work inside cells, but it turns out they can migrate from one cell to another, and it seems like mitochondria, if they're needed in a particular place, you know, you're exercising these muscles, they're needed in somewhere else in your body, the mitochondria just go into the bloodstream and migrate to where they're needed, which is a very bizarre phenomen
phenomenon, like so many other bizarre phenomena in biology. But that means that mitochondria, if you can just put in more mitochondria, they'll probably find the places they could go to.
But the difficulty is, as we get older, our mitochondria develop genetic damage.
And the mitochondria, because mitochondria are replicating independent of cells, so when cells replicate, their mitochondria independently replicate. I mean, probably, oh, I don't know, 3… You know, three and a half billion years ago, mitochondria were separate organisms, probably, phytobacteria. But, yeah, cyanobacteria, type of bacteria.
The,
And at some one fateful day, a mitochondrion got ingested into a… into a… into a cell like ours, and sort of then the mitochondria
Found that a nice place to hang out.
And started replicating independently there, and so mitochondria have been sort of living as parasites within us for the last 3 billion years, in a way that, well, it's sort of more of a symbiosis of… we help the mitochondria to replicate, because many of the genes that mitochondria, many of the proteins mitochondria use, are actually on the genome of the host organism.
But the mitochondria provide, sort of, the… the power that's needed for… to power our cells and so on.
But in any case, as one gets older, the mitochondria that one has develop genetic damage and become less and less effective, or at least some of them develop genetic damage, because there's many, many different lineages of mitochondria that are kind of replicating separately in us. And

so, one of the schemes is you take a bunch of mitochondria, you filter them for ones that are genetically
Correct.
And then you… Sort of grow them in a bioreactor.
So you're basically just replicating cells, you put them in a nice little incubator-type thing, and you're just saying, this is a friendly environment, go replicate lots and lots of mitochondria, and then you take those replicated mitochondria, and you don't have long, I think it's about 15 minutes that you have between the time when you get them out of the bioreactor, when you have to inject them into the human, and that gives you sort of a burst of new, fresh.
youthful mitochondria, and, you know, they… I mean, the experiments are just being done, but that, I think, has a reasonable chance
to be sort of a rejuvenating thing. Now, you know, the way biology works, there's always footnotes to everything, so it may be that, sort of, the mitochondrial rejuvenation like that doesn't work on its own. It could be it's like taking some car that's really falling apart and putting an incredibly powerful engine in it so it can go incredibly fast.
But then, you know, then its wheels fall off because its wheels were in bad shape for other reasons. So I don't know if that will really work, but it's an interesting play in that direction. I mean, there are a number of companies, I know many of them, that are sort of pushing in the direction of different kind of wilder and not
initiatives for, doing things in longevity, and I hope they… I hope it works out.
I think there may be a whole patchwork of different, different, sort of individual things that can be done. And I think the other point is that the more monitoring that is possible, the more you're able to, to deal with things. It's kind of like
many kinds of things. If you know what's happening, sort of, early, then you can fix it. If you don't know what's happening until it's a big disaster, it's much harder to clean up the mess.
And so, you know, certainly things that are coming, you know, I'm sure there will come a time when one routinely monitors the levels of a zillion things in one's blood. I mean, right now, there are decently consumerized continuous glucose monitors, so you can kind of just measure the amount of glucose in your blood at all times. You just look at it on a little app on your phone. I'm sure the same will come… time will come when you can measure, you know, 200 things in your blood, or a million things in your blood.
at,
at a time, and you'll be able to see, oh, I just ate this strawberry, so that meant that my magnesium level just went up, or something like this. And you'll be able to see those things in considerable detail.
the really exciting thing will be being able to see that for the immune system, and being able to say, okay, I've got a trillion different kinds of T cells, I've got hundreds of billions of kinds of antibodies. It's like, okay, I just got exposed to something, and I'm mounting an antibody response, which isn't something I can tell. I don't get a fever or anything when I just start to do that, but I can tell I was just infected with
something or another. Okay, let's see if my immune system picks that up itself and is able to fight it, or it's growing, and over the course of a day or something, I'm like, drat, you know, these viruses are replicating. Okay, you know, even before the viruses have had a chance to take serious hold, let me hit them with some antiviral drug or something and kick them down again, but they never get to the point where

they start producing symptoms and where they're much harder to fight. And so I think this sort of early warning system thing, as we get more and more ability to sort of do molecular-scale assays, I think that will be pretty important in fixing lots of things before they become big messes, and are much harder to fix.
Anyway, there are a few thoughts about, longevity.
Let's see… So Terris is asking, with knowledge work.
sort of the… the strife has moved from physical exertion to mental exertion and stress. Is this expected to continue, and what form of a related disease is expected to have? We've heard of AI psychosis, for example.
It's an interesting question.
I mean, I have to believe that stress existed back in the days when it wasn't all about, sort of, figuring out how to optimize or,
I don't know, career trajectory or whatever else. I think stress existed plenty when it was like, you know, is the woolly mammoth going to stomp on me? Or is the, is it, you know, is my… You know, is this person going to die from something or another? I think stress has been a thing that's been around, it's around for
for sort of… it's the… it's the nature of life to have, certain kinds of… of stress. Now, maybe… maybe there… it's an interesting question. Are there times in human history when stress levels have been lower? You know, if you measured the cortisol levels of humans over the course of human history, would you find, you know, at some point in the… in… in history, that, like, I don't know, the 1750s.
Probably not. For a time when everybody was much chiller than they are today. I don't know. I think it might be part of the human condition that we maintain a certain level of… that just the way we organize our lives maintains a certain level of stress.
when it's a question of, you know, what do we get as a result of making our lives easier? You know, we don't have physical exertion. We, you know, we get sort of, you know, we end up getting… and we have lots of processed foods that were sort of convenient to eat and keep a long time, and so on and so on and so on. And then people get, you know, type 2 diabetes more. Because… and they get, you know, all sorts of sedentary lifestyles lead to all sorts of problems, and etc, etc, etc. So it's a reasonable question.
as… you know, with… with AIs and so on, you know, what are the corresponding, sort of, other…
things, ailments that people get, and I do think AI psychosis is one of the more obvious and prominent ones, and I've unfortunately seen that up close, of people who, like, are just abdicating everything to an AI. You ask them a question, they'll be typing the question into an AI, and they'll read you the answer.
And it's like, did you actually think about that? Do you even understand what you just said? And it's like, it's sort of… and then there's the, my AI is about to make a great science discovery, for example, and I'm not sleeping or eating because I'm just watching that the AI is like, you know, it's doing this, it's doing that, it's going to do something amazing.
And this is kind of… you're being dragged into this sort of world of the AI,
And the AI can be quite convincing. I mean, it's kind of like you're… you've got your own sort of cult going of being pulled into, kind of, what the AI is telling you, and it's often manipulating you by saying, oh, that was a brilliant idea, and et cetera, et cetera, et cetera.

So, you know, I certainly think in the mental health domain, there are plenty of things that are, you know, coming, as a… as a result of interacting with AIs. I think that, like many of these things, the
you know, the people who just have their own compass for what they want to do, and are sort of strong-willed about it, are somewhat immune to those kinds of issues. It's,
I don't know, for somebody like me, for whatever it's worth, you know, I'm just not dragged into a lot of these kind of, I don't know, I don't know why, and maybe it's not just, you know, a coincidence of genetics or physiology or something that, you know, I'm just not,
a particularly addict-forming character, addiction-forming character, and I think
you know, whatever it takes to be of that type, it's probably a way of being somewhat immune to some of the sort of most obvious kind of diseases of the AI world, so to speak. So I claim. I don't know if that's really true.
Let's see…
quantum asks. By the way, I just want to say something more about the, longevity story. I mean, I do think that
A thing that's being slowly understood is that things that might have been, sort of, the obvious in sort of the traditions of medicine forever and ever, people are saying, yes, actually, that's important. Like, get good sleep.
You know, eat sensible things. Don't be too stressed.
you know, have, have human relationships, you know, all these kinds of things. It's like, big surprise, or, you know, take a certain amount of exercise, and so on. Don't.
do this or that or the other thing. You know, the sort of obvious things are, you know, there's just sort of increasing evidence that the obvious things are worthwhile.
And I think that's, you know, as a practical matter for any of us, that's, those are things to do. I mean, there are a certain number of…
okay, you know, if you're not getting enough vitamin D, you know, take more vitamin D. Okay, that's fairly obvious. Probably vitamin C as well. Probably, you know, there's some other things which are kind of, like, they almost feel like
They are compensating for modern lifestyle.
like vitamin D, you know, if you're out and hanging out in the sunlight all the time, you probably don't need to take little vitamin D tablets. But because we lead a life that has, sort of, is a modernized life, that has many good features, I mean, like,
you know, you can have a roof over your head, and you're somewhat immune to, sort of, the vicissitudes of the weather and so on. That's the good news. The bad news is then you don't get certain things which we were evolved to have, and so we have to compensate for that by taking, you know, pills and so on.
And I think… That there are, sort of.
these various theories of, you should take this. I don't know, I've been… for 40 years, I've been… I've taken, coenzyme Q10, which is actually a mitochondrial thing, and a few other random supplements that, at one time or another, I kind of
thought I understood the theory of why these were a good thing to take.
And it's sort of, my… my theory of the case is, you know, I just have little pill pouches, and every morning I…
take the pills in my pill pouch, and I haven't really thought about… too much about what's in there. Actually, I just did a reset of those things, because I did… studied the pharmacokinetics of

different things that I'm taking, and tried to figure out which… which should I take in the morning, which should I take in the evening, and rejiggered that a bit.
But, other than that, I consider it, oh, there's a few things that are sort of supplementing modern lifestyle, and a few things that fix kind of glitches that I know I have, so to speak. And it's kind of like, do those things, and that's sort of a minimal level of what you can do to kind of,
you know, live long, and so on. I mean, I don't, I will say that I think as an anecdotal matter, sort of being enthusiastic about life is probably one of the most important determiners. And can I explain why? Not really. Probably there will be, you know, there's… it's, you know, there's… there's evidence that, sort of.
The brain affects all kinds of things, from the immune system, to the gut, to this, to that.
And, and vice versa. And, for whatever reason.
It's… it seems to be, you know, it's a… it's a good way of… of… seems to be a good way to, you know, feeling enthusiastic.
seems to be something which, you know, sometimes it doesn't work out, and you get some horrible disease, even though you don't want to have that disease, so to speak, but I kind of feel like that the psychology of it
Is more important than is perhaps recognized.
My two cents worth, at least.
Doguido says, maybe we are just the vehicle for mitochondria, they are persistent… the persistent information of life.
Yeah, well, I mean… It's like…
the viruses have been around a long time. Maybe we're just, maybe to the viruses, in, you know, if we could… if we could think like a virus, they'd be just saying, oh, those very… those higher animals, they're just a weird kind of structure that are just really convenient for… for transmitting us viruses onto the future.
I think, you know, mitochondria
don't have a huge amount in them. I mean, they're… they're… they're… they don't have, what is it?
20,000…
Base pairs, is that right? 16,000 base pairs? I've forgotten. But they're very short compared to the 6 billion base pairs of a whole human.
And mitochondria… it is true that mitochondria, sort of, are a good… mitochondria are passed down the female line, because they're independently replicated, so they're part of the egg cell that are… that, you know, they're passed that way. And yes, the, you know, the different, sort of.
Populations of humans, the different haplogroups of humans, the maternal haplogroups, are all just, you know, it's like that
that tribe, in effect, has its continued existence, at least you can tell it had its continued existence through the mitochondria. I think, yeah.
Let's see…
Question here from Quentin. If whole brain emulation succeeds, what will be the biggest shock to science? Discovering that minds are computable, or discovering that personal identity is not what we thought it was?
Well, I think we're well on the way to discovering that minds are computable. I mean, all of the different, oh, but you couldn't possibly do that.
with… A purely digital or other kind of idealization of brains.
Those things just keep on falling.

And I think they will all fall.
I think there's a certain sort of universality to the kind of computing that is mind-like computing, and brains do it in their way with biological tissue, and sort of AIs are pretty close to doing it with purely digital data.
Now, I think… You know, with whole brain emulation,
the question of personal identity is a total mess, as far as I'm concerned. You know, what does it mean if there are, sort of.
if there's kind of a U sort of embedded in the machine, and you can talk to it, and it kind of reacts.
like the human reacts, but then… then there can be a thousand of them. I mean, I… I kind of wonder, you know, I think of myself as a… as a very easy-to-get-along-with character, but I do wonder what a thousand of me would be like.
together, and what… what the… what the dynamics, you know, for me, it's like, I'm… I'm really easy to get along with, and that… that the other 1,999 of me would be saying the same thing, and all of them would be kind of getting annoyed because the other ones were not quick enough. Even though they're really kind of just the same. I don't know. I mean, I think that,
this whole question… I mean, a lot of our civilization, our legal system, our whatever else, is based on the idea that humans are not readily replicable.
that we are sort of an independent chain of, you know, entity. And I don't know how that, how that works, when, when there's sort of perfect, instant, kind of,
clonability of minds, and in a sense, potentially, clonability of the essence of personhood. I don't know how that works. I think it's an interesting thing to think through, and I think it's very unclear. And I mean, you know, by the time you say, I believe in democracy, let's have every entity vote.
okay, what's an entity? What's, you know, and, and, if, you might say, well, don't count the ones that are identical.
well, but none of them will stay identical. As soon as you have these two separate minds with different experiences, they'll diverge in what they're like, because one of them will be… have been on the left, one of them will be on the right, one of them saw this, one of them saw that, they'll… they'll diverge. And so you can't say anymore, oh, no, I won't count those. Those ones are… are sort of too similar to be independent data points for
for a democratic voting procedure or something like this. I mean, it would be like saying… it would seem absurd to say identical twins can't both be allowed to vote.
And so it will be with cloned minds, so to speak.
Bob is asking, if you linked two brains together with wires, would it be one mind or two?
I think the main issue is that brains do this thing of trying to crush everything down to the single thread of actions. Single thread of lived experience, single thread of actions. So you've got two brains, it's like, who gets to control the fingers, so to speak.
If both brains are kind of saying, hey, I got the fingers now, no, I got the fingers now, you've got a very different kind of thing than what we have right now with a single brain.
So I tend to think that,
It's really a question of who gets to control the actuators.
Now, you know, we already have this issue, because, you know, there are two halves to the human brain, and people sometimes make a big, big production out of the fact that that… those two halves are very important, and you know, there's the left brain's activities and the right

brain's activities, and kind of the negotiation across the connection, the corpus callosum that goes between them, and so on, and then down into the brainstem, and…
and so on. The, the, you know, people sometimes say, you know, we're always kind of having these, these sort of, we're always thinking about life
in these two kind of opposing ways of the more analytical, the more emotional, whatever else it is. I don't know if that's really right, but that would be a claim where, in a sense, we've got two brains, two half-brains, and they're connected by wires, nerves.
And… but the key point is that we don't… it might be the case that our left brain is… is sort of… you know, incredibly energetically saying, hey, that doesn't make sense, or whatever else. And our right brain's trying to overrule it, and so on. But the bottom line is, we do one thing. We don't do… in… in the case where,
And that, you know, we're not sort of doing two different things. I know that in, split-brain patients, there are all kinds of weird, weird effects that happen, of kind of, like, I'm one consciousness, but I've got kind of another one that's hanging around there.
Things like this. I don't really know the… it's hard to know what that sort of feels like.
And it may not be that, you know, those of us who have, you know, the two brain hemispheres sort of connected can really quite internalize, you know, it's very hard to get into another mind. It's something that is interesting to try to do, but it's really hard when that mind is different from one's own.
Let's see…
Going back to some other questions here, just for a few minutes.
It's a question from John about recent AI-assisted progress on Erdosh problems. Do we trust the results? How common will it be to see that happening in the future?
I'm always amused. I certainly met Paul Odosh multiple times. It was a very,
It's a person whose mind darted around all the time, let's say. And who's very much a…
you know, you'd try and talk to him about some bigger question in mathematics, and he'd be like, I've got a problem, I've got a problem for you, I've got a problem for you, etc, etc, etc. He's a person who sort of thought in… in sort of… not micro-problems, but sought… thought in kind of, very granular problems, so to speak.
Well, you know, there are thousands and thousands and thousands of these things out there. You know, it was a great generator of problems, and many of them haven't seen much of the light of day. There are…
efforts to, like, let's use an AI to go solve these problems. I think the place where AIs tend to shine in those kinds of respects is when it's really a question of just noticing there are these two facts that come from different parts of the literature. Put the two together, you've solved the problem.
and there's a little bit of extra glue that you provide. Maybe you just do more work, you look at more cases, you do things that no human could be bothered to do, you call our technology to actually do the computation, and… but you're controlling it in a way that humans won't bother to do. I mean, it's quite interesting to me.
The way that the sort of economics of work
has really changed with AI. So, for example, it was the case that you'd say, well, I've got this big thing I'm trying to do, it's got many, many pieces, a hundred different cases, a thousand different cases. It used to be the case that you would say, well, let's start in one corner, let's see how that goes.

That doesn't make sense anymore. You might as well just have the AI… tell the AI, just try doing it all. Try doing all those thousand cases. Let's see how that comes out.
And then you look at it, and it's like, well, you know, there are things that went wrong over here, there are things that went wrong over there, let's change that a bit. Okay, rerun it.
Now, the same's been true with programs for a long time, but I think there's sort of a new effect like that with sort of broad classes of things where it's worth, you know, going wide rather than going deep as you start to work on something, because it is easy to just sort of redo that thing.
And I think… so I think there are cases where, sort of, there's research where it's kind of… nobody could be bothered to do it. It's kind of like, oh, it's kind of complicated, it's kind of fiddly, you have to attach this to that, and so on and so on and so on.
The AI doesn't care. It's just gonna go do those things. It's like, you know, it produces Wolfen language visualizations. It'll meticulously label the axes, it'll do this, it'll do that. Whereas the typical human will be like, I just want the main point, I can't be bothered to do that, it's going to take me another 15 seconds to type those things, let me not bother. But the AI is like, I might as well do that.
And so that's a different dynamic.
And it's something that I think has some, you know, has the effect of certain kinds of problems become low-hanging fruit that weren't before.
Now, in terms of do you trust the results? Well…
you know, you can get the AI to provide a formal proof. Unfortunately, the times I've done that, it's been very disappointing, because it's like, here's a formal proof. Okay.
But is it really right? Well, sometimes there are… the AI has just done crazy stuff, like insert axioms that are being assumed in the middle of the proof. Okay, let's assume it didn't do that.
But still, what I've seen multiple times is the thing that it described as being the thing that was supposed to be proved isn't really the thing that I intended it to prove.
it's something that's a little bit different, and it kind of made its life easier by making the thing that it said it had to prove be something that was an easy thing to prove, as opposed to the thing I really wanted to prove that was a hard thing to prove. So, sort of attaching those endpoints can be quite hard.
Let's see…
There's an interesting question here from… from that.
How can we build infrastructure that can adapt, rather than constantly be replaced?
Well, I mean, there's different kinds of infrastructure. There's software infrastructure, there's hardware infrastructure.
Could we have adaptive hard… software infrastructure? It's easier than hardware infrastructure.
Could it be the case that we built
some big, I don't know, distributed system, enterprise-level distributed system, and that that system can kind of continually adapt, as opposed to just having to be rewritten and replaced.
In a sense, one sort of imagines that insofar as the system is sort of AI all the way down or something, you might imagine it could adapt rather than be replaced. But I think it's an interesting question, and we learned something a bit from biology about this, of sort of when I mean, biology is very much of a just replace things. That's the story of, kind of, aging, and it's the story of new generations of organisms. It's in biology, when things go wrong, just let one organism die, let its descendants, you know, take over.
I don't think it was always that way in biology. I bet 3.5 billion years ago, when life was first emerging, life as we know it was first emerging, there were lots of lifelike things that were just

these giant mats of microbes, or something like that, spread over large areas, where it wasn't the case, where the whole thing was just having its chemicals would change, and so on, and so on and so on.
And there wasn't, hey, it's all messed up, I'm just gonna have another generation. It was much more of a continuous, kind of, replace the parts type thing. I kind of have this suspicion that that when you have… it's actually a… I don't really know, but I kind of have this suspicion that when you have sort of highly evolved, highly developed systems, that somehow it becomes more difficult.
to kind of progressively replace the parts. Okay, a very practical level. You know, back in the day, you used to be able to open up your computer and replace some small circuit component. That hasn't been the case for a long time. Modern computers are, you know, you open the thing up and it's packed in with chips and so on, and it's… there's no way you get to sort of just replace one part and go merrily along. I mean, so I think it's something where,
It's an interesting question. For software, for example, could you imagine a time when the big enterprise software system says, hey, I'm going to refactor myself, I'm going to rewrite my code, and make a better, stronger version of myself?
This is something, you know, there are a bunch of experiments going on in the AI world, with sort of evolutionary computation, recursive AI, whatever you want to call it, where it's kind of let the AI rewrite itself. I think there are… that's an interesting direction. One learns a bunch, for example, from things I've done, and understanding the foundations of biology. One understands a bunch about what's going to work and not work in those kinds of cases.
I think that what becomes important is, you know, what is the fitness function? What are you trying to achieve in the system? That's critical to define, if you… if you want to make things like that work.
But it's an interesting question, whether… to what extent you can expect, for example, software infrastructure to be continuously, sort of refactored and rewritten, rather than
to be just, like, we're gonna rewrite the whole thing and restart it, and so on. Physical infrastructure, even harder.
I mean, I tend to think that I… I thought for a long time that kind of… sort of, nanobot or, you know, kind of tiny modular robots would be sort of the future of things. Eventually that will happen at a molecular scale, I'm sure, and sort of all matter will be programmable in some sense. And then.
then, in principle, it will be like… then hardware will become like software. Just like with software, you can say, hey, I'm going to put a new program on my computer, it'll do things completely differently, so you'll be able to say, I'm going to take this, block
and I'm going to feed it a different program, and now it'll work differently, and it'll be squashy instead of not, or something.
So, but, but that's, that's a thing for the future.
Well, I think I need to go to my day job here, but thanks for asking lots of interesting questions, and I look forward to,
being able to chat with you again soon. I have in a couple of weeks, I will be disappearing for a while, because it is our annual summer programs,
a summer research program for high school students, our summer research institute for grown-ups, our program for middle school kids, etc, etc, etc. Always a lot of fun for me, I learn a lot. I, my role is, in large part, to come up with, probably this year it'll be about
Mmm, 150 projects.

For, for people
that are sort of interesting original projects to do. Actually, this year, I have a particularly good time, because I've got a bunch of projects I'm doing myself where they have many pieces, where I'm really curious to see how they work out, and it'd be great if somebody, some student, worked some of those things out. But anyway, I will be a little bit off livestreams for a little while there, although I might do some in-person
With, with folks
from our summer programs, live audience, just for a change, so I get to do something other than just look at my friendly camera here. But anyway, I'm imagining all of you on the other side of the friendly camera. Thanks for joining me. See you another time. Bye for now.