Hello there welcome to another Q&A about the future of science and technology and I see a whole bunch of uh questions here so let me uh take a look at some of these so LC asks "What is the future of expositions like the one I'm doing right now in the presence of AI and personalized pedagogy? So 1st thing to say is where the raw material
for whatever an AI might do, and I think
there's sort of a question of when can you get
an answer that is specifically the thing you need to fill in to your sort of cognitive map versus getting something that where everybody's going to hear the same thing.
I have to say, some of my experiments on personalization
have been sort of unsatisfactory for certain kinds of reasons.
Sometimes there's value to everybody, sort of hearing the same thing, then they can talk about it.
and so on. There's also sometimes value in. I want the thing that's customized to me, and that is sort of the best possible fit to me.
I mean, I think, right right now. We're building, for example, an AI tutor
that is initially targeted to algebra one, but its goal is to kind of go through a particular curriculum in a way that is completely customized and personalized to each particular student. So if they have a particular thing, they misunderstand, for whatever reason they, the AI tutor, can sort of explain the particular thing they need. It's an interesting question. If you ask, for example, teachers of something like math.
how many different confusions have you seen in your career. Say, well, it's a few 100
and different people are confused in different ways, and it's kind of rare that they see kind of a new confusion that they've never seen before. But there's still a pretty broad range of them, and some people are just not confused by things that other people are confused by, and vice versa.
I think the so. You know the idea that you can be fed. Just the piece of information you need to understand something that's a really useful concept and personalization. I mean, I know for myself, there are things I try to understand where, for example, I'll ask people about them, and they'll be going, you know. I'll tell you this, this, this and this, and I'm always interrupting them, because, like I know all of that stuff. Tell me this one thing I don't know.
and that's the thing which I can expect. Kind of an AI to get to know me well enough that it knows. No, you don't have to explain this piece of math or something to me. But this other thing no, I don't really know anything about that other thing. You're going to have to explain that much more slowly.
I think this question of how do you take material that was intended for sort of broadcast presentation and turn it into something that can be kind of personalized. That's an interesting question. Actually, we've been working quite a bit on that. We've been building kind of a whole framework
for letting people work with kind of broadcast material and sort of
chop it up and reconstitute it in a way that can be tutorified, so to speak. I can't yet say exactly how well that's going to work. I think we're going to know more a lot more in the next couple of months, probably about how that works. I'm somewhat hopeful
that will be able to do that. Take sort of the arc of kind of a. This is what you should learn about this thing. This is what a textbook about that thing is like, and turn it into something that can

where you can kind of meaningfully have sort of a personalized educational experience. And I think that the concept of talk to the book

is is works okay with Llms. As they are today. I think we can do beyond the sort of talk to the book. There's kind of lead you through through learning a particular kind of thing which has a bit more to it than just talk to the book. Ask a question about this particular paragraph, but I'm sort of hopeful that we'll be able to make that work well, and that that will provide a good personalized experience

with the caveat, that I think there are situations in which everybody hearing the same thing, is useful, and is a good kind of I don't know social ish, experience, or a good thing, where people can talk about what they're hearing rather than well, I'm just, you know, on my own. It's me and my AI kind of in a box, learning things that are completely different from what other people are learning, even if the the final endpoint is is the same.

Let's see, Francis asks, will computer simulation eventually replace physical experimentation. Well, there are lots of situations in which people talk about the digital twin of some system. There's a lot of cases where it's much more efficient to simulate things than to actually do the physical experiment.

Because each run of the physical experiment may be expensive, you can, perhaps you only have one piece of apparatus, whereas you can run many simulations in parallel all those kinds of things. But in the end you're never going to be able to simulate everything about the universe. The universe will always surprise you in all sorts of ways.

and I think you you'll still always need to have a base of physical experiments, I mean, that's all the more so in the life sciences. Right now, we're really not in a position to simulate all that much at all in the life sciences at this point. It's, you know, the question of, can we replace kind of the clinical trial of the drug with the simulation of what happens? Well, we can get some information from simulations, but we just don't have the sort of depth of information depth of computation. I mean, there's a this biology, I think, is rife with computational irreducibility. It's rife with situations where you just have to run it and see what happens to know what's going to happen. So I think that's a case where, when it gets close to life, sciences, biology, and so on. I think that's a case where experiments are still very much the thing one has to do. Now one can automate the doing of those experiments.

you know. That's a thing like our friends at Emerald Cloud lab have done a lot of work. Sort of getting one from you. Write the Wolfram language script. You press go, and then kind of robots and some humans kind of swing into action to do your experiment in the end. You're still doing experiments, you know, pouring liquids into things and stirring things and and running things through sequencing machines and all that kind of stuff.

But the the sort of the interface to the the user looks very computational. It looks more like a simulation. But in the end there's an actual experiment going on underneath.

Sort of interesting. Maybe that's partly, you know. Maybe that's sort of an intermediate in the way the future looks is the kind of computer controlled experiment, the simulation where there's no underlying thing, the underlying thing is just bits in your computer to. There's a thick layer of bits in your computer that's setting up the experiments. But in the end there's a robot actually doing the experiment underneath

to the case where you are much more directly connected, where humans with their own fingers are pipetting things and so on, and doing the actual experiment. I think that's sort of a gradation there between sort of it's just simulation to. It's the pure sort of human doing experiment.

And probably there are sort of intermediate steps there, where you're simulating a bunch of things and only measuring one piece of it, and then deducing from that measurement a whole bunch of things. So I think it's a sort of more continuous gradation between those things.

But, as I say, the the question of Do you need to do the physical experiment to know what's going to happen. There are cases where you absolutely need to. Now, you know, in our kind of approach to fundamental physics, we kind of believe that in the end

there is a model that just gives everything in the physical universe.

But it's very difficult to run that model. In fact, it's so difficult. The amount of computation that's needed to run that model is just the same as the amount of computation that the universe has spent to get those answers for itself.

So you don't get an advantage in so far as there's computational irreducibility you don't get an advantage from oh, I'm going to run it on a computer because you might have to do as much computation as the universe had to do. And the only computer that's sort of powerful enough to run the whole universe is the universe itself. In other words, you basically just have to watch and see what happens.

So I think it's inevitable that there'll be pieces where you just have to sort of watch and see what happens. How much layering of kind of computational control of that one can get. That's a different issue. And I think lots of such control can be got. And I think that's what things will look much more like in the future is is, you know, some things are pure simulation. Many things are a thick layer of computational control on what is ultimately a physical experiment going on underneath.

Let's see, next question we have on our list is from Sammy.

What do you think is more valuable for the future of science, stopping, aging for prolonging life, or figuring out how to de-age becoming younger and living longer that way.

Well, I think none of us want to sort of be be kept in a state

where it's kind of like we're 100 years old, and lots of systems in our bodies have failed. That doesn't sound like much fun to me. I think one wants something where one is in a state where one feels sort of perfectly sort of healthy, I'm happy to say, even at my advanced age. I'm touch wood. I'm still. I still feel pretty healthy and energetic.

so I'd be perfectly happy if I could just stop aging where I am right now, I think. Will it matter if I stopped aging at age 40, instead of my current age? I don't know, for for me in particular. I think I was less healthy at age 40 than I am now, so that wouldn't be such a win, but in terms of what it takes

to

what it takes to kind of stop aging versus reverse aging. I kind of think those are somewhat the same kind of thing. I mean, if we ask, you know what what leads to aging? We don't completely know yet. I mean, there is genetic damage.

There's and the genetic damage has many, many downstream consequences, both sort of material damage and genetic damage. Material damage, as in actual cells, dying kind of sort of joints, creaking things like this. And there's genetic damage as in

the kind of the program for successive cells, gets more and more kind of inaccurate and error prone, which has the consequence that there will be sort of material damage later on. But you know, there are many different components to this. I mean, biology. Life is a very complicated, we would call it technology stack if it was built by engineering. But it's kind of a stack of kind of capabilities and systems that's pretty tall, and that's been kind of being assembled for the last. You know, more than a couple of 1 billion years, and it's sort of shocking, in a sense, how many of the systems we have

are incredibly ancient. I mean the the methods we have. I don't know for DNA repair, or something like this, probably a billion years old. There are lots of other kinds of things where we are. You know, it's been a very slow process to build up this stack of capabilities that represents what is now sort of in our current biology, so to speak. And I think the

the question of sort of what leads to aging. Well, there's just I mentioned some genetic damage, material damage, oxidative stress. That's sort of just chemical damage. There's there's lots of things that sort of just don't

degrade the immune system degrades. Now what can be done to reverse these things? Well, clearly, biology has a mechanism to reverse them, because when it makes the next generation, it's kind of rebooting and starting from scratch and everything's healthy and young.

And so the question is, can one take the existing biological organism and sort of while the plane is flying. Can you reassemble the thing? Or is it the case that much like modern computer systems? The only way to fix it is basically to replace the thing. To start again from sort of the next generation.

I think we don't know, I think, sort of the intermediate case between the we've got to sort of start a new generation and we fix it while it's still sort of ongoing is kind of all sorts of regenerative medicine plays where you're saying, well, we replace that component and so on. There are questions about, you know. Can you replace?

Can you sort of roll back the clock on everything versus, can we replace component by sort of big component like replace this organ versus, can we replace this class of cells? That's much more microscopic? And there are things of interest in all of these directions.

I mean, I think the so sort of the big piece of progress made I don't know 15 years ago, or something now was discovering, sort of more or less how to turn the clock back on cells and revert them to stem cells from which all the few 1,000 different kinds of cells in our bodies can eventually be derived.

And it's not a perfect thing. You take a cell skin cell. Whatever else use these Yamanaka factors, you kind of feed it various kinds of chemicals, and it successfully manages to go back to a State pluripotent stem cell ipsc induced pluripotent stem cell. That is a thing that sort of is roughly a cell from which all other cells in the body can be derived.

It isn't perfect.

The Ipscs have various kinds of genetic instability. They still seem to have some trace of the type of cell that they were before they were reverted, and probably the reason for that is that there's other stuff in a cell that's beyond its well. It's not the things that that cause a cell to be of a particular type

are still not fully known. There are things that have to do with the actual sort of cycle of chemical reactions that are happening in the cell. There are things that have to do with sort of annotations to the genetic material, and so on, that indicate the cell type.

And there are other features of the cell that are still not completely understood, that say, this is a skin cell versus. This is a heart muscle cell versus. This is a brain cell, or whatever else.

But any case, there's a sort of in a 1st approximation. You can revert things to the point where there's a a stem cell that could turn into anything. Then there are a certain number of procedures known to get that cell to turn into various different type cell types.

But you know, if you do that with brain cells, you get some brain organoid. That's a blob of brain cells. But it isn't really quite like a brain as a brain is usually configured because a brain, as it's usually configured is full of more detailed structure. That isn't just a blob of cells. It's something which has grown with all these different features, and so on.

And so one thing one can do for some organs is to say, let's sort of 3D. Print a scaffold, and then let's feed stem cells into that scaffold and use that scaffold that we just 3D. Printed. We just sort of mechanically printed it, and and hope that the stem cells living in that scaffold will behave the way that, or the cell differentiated cells living in that scaffold will behave the way that they would for that kind of organ.

So I think the that's sort of another play is, you know. Can you can. You make kidneys, lungs, whatever else, and just replace them?

And you know, just plug new ones in

that looks reasonably promising as a as a strategy. But that's sort of organ at a time type work.

Another thing one can do is sort of operate at a somewhat more. Well, okay, there's the molecular level. And then there's sort of a level of organelles of structures in the cell. So one of the big targets is mitochondria. They are sort of the source of energy in us.

And you know they presumably came from a completely separate, free living organism. You know, a billion or 2 years ago that got integrated into each of our cells to provide an energy source for our cells. And now they're so deeply integrated they're making use of proteins that come from from the standard nuclear genome, and so on, as well as the ones that are their own, separately replicated little circular genomes, and so on.

But in any case, mitochondria, I think, represent like 10% of our body mass. And the evidence is that they're kind of they're the things that when you need energy somewhere. They go very quickly and deliver themselves to the thing that needs energy, you know, if you're running a race or something, and you need more energy to muscles and things like this. The mitochondria will go there. Maybe, when you're doing more thinking the mitochondria are going to the brain. I don't know

but in any case they seem to respond very quickly, to sort of go to sort of be flowing freely through even the bloodstream, and so on, and then sort of be be ingested into cells to go and provide energy for those cells, and they go elsewhere, and so on.

Our mitochondria degrade over time, and as a result of genetic damage, presumably maybe chemical damage of various kinds that leads to genetic damage.

And so that's 1 of the targets of, can we just supply ourselves with new mitochondria? Can we, for example, pick out mitochondria that are nice and healthy, and sort of grow them in a bioreactor just separately grow them in kind of a in more or less than a dish. They're not really dishes. They're differently shaped things that look

kind of

The the how do I describe them? They just sort of pots

that are carefully controlled, and that allow cells to independently replicate kind of in the pot, so to speak.

Actually, you know, it's a tricky business, because you end up with sort of mitochondria embedded in some other kind of cell. And those cells replicate, and then you get them to spit out their mitochondria. Then you inject the mitochondria. Those kinds of things. These are all things. People hope that they might be able to do. It's not yet a proven kind of thing to be able to do, but that's another level is replace the mitochondria, you know and give yourself more energy, so to speak.

Another level is things where you're dealing with, for example, immune system cells. And then other levels are more at the level of things like individual DNA, and so on. But at the level of well, the immune system is kind of our

good

kind of thing that goes around our bodies kind of surveilling for damage, surveilling for problems. And well, really, the only thing it can do is to zap the cells that are problematic. But already that's useful, because what tends to happen in many situations, whether it's you get an infection, whether it's even the beginnings of a tumor or something. It's like this is a bad thing, and the best thing to do is just sort of cut it out

and do that at the cellular level. And that's what the immune system is set up to do.

How do we tune up the immune system? Well, that's been the subject of a lot of work, and we don't have as good an understanding of how the immune system works at all. So it's hard to know kind of how to tune it up. But there's lots of things that seem like they're promising as ways to tune up the immune system. Make it better at surveilling for things that are bad if we detect something that's bad being able to sort of tell our immune system. Go surveil harder for that. That's kind of what vaccines are trying to do

and but that's something that's happening at the level of individual cells.

I think the there are.

Oh, gosh! There are just all sorts of different strategies. But the one thing you might say is, well, okay. We know these strategies for reverting cells to stem cells sort of making the cells. Young again, why not just use those factors and just sort of inject those factors until every cell get young again? Well, that profoundly doesn't work. Because, as you revert cells

and stop them, having differentiated into the type of cells that they are in your body. If you try and make them go back to their stem cell state. They basically start a bunch of tumors that's their sort of cells that aren't cells that know what they're trying to do. And so the often cells that sort of don't know what they're trying to do will kind of revert to that very primitive stage of life

from I don't know billions of years ago. Now that corresponded to organisms that were just cell cell. Just build up, you know, this big blob of cells. Which is which is sort of what's happening in in tumors and so on.

Is. It's kind of this. This collection, you know, malignant tumors, or whatever that's just the cells are just growing, growing, growing without bound. It's sort of a more primitive form of life than the one that we have where we grow for a couple of decades and then stop growing, and then just replacing cells after that.

But so to the original question about, can we kind of revert or just prolong things in their current state. I mean, I think the whatever's going on we have to kind of

get rid of the bad stuff and put in new good stuff.

and whether that sort of whether the dial of what the effective age goes backwards or stays the same, I think, is sort of a detail. I mean, I will say that another thing

that is another sort of feature of biology, as it is right now, is cells that are kind of not really making it senescent cells of various kinds. You know something that we try to do. Our immune system tries to do. Just get rid of those cells. You know that cell is not doing the right thing. Okay, zap it.

You know, we can always make more cells because we can always divide existing cells, and it's kind of like if something goes wrong, just get rid of it and expect to replace it with a newly divided cell that had a better chance to do the right thing.

So that's kind of that approach. But I tend to think that sort of the. There are multiple prongs here from things like energizing the immune system to be more helpful to and using the fact that we can sort of compute what the immune system can do on a computer outside of the body, and then go tell the immune system to do that rather than letting the immune system kind of figure it out for itself.

I mean, it's kind of like with our brains. We get to use computers to help us figure out lots of stuff that we couldn't figure out just with our brains.

Well, our immune system is also a kind of computational system that's trying to figure out what to do with all these antigens that get into our bodies and so on. And right now it doesn't have any AI assistance. It doesn't have any computational assistance. It's just doing its own thing with our brains. We get to read from our computer screens, you know, run our Wolfram language code, whatever else it is. But with our immune system. We don't really get to do that. you know. We only get to do that through the beginnings of various kinds of medical interventions where we're seeing sort of some aspects of what's happening. You know, something is not going quite right in the immune system. Okay, inject this thing to make it better.

You know there's some direct cases like for antivenoms, for example, where you can just inject antibodies into a person and expect those to sort of clump around the venom molecules and help things. But most of the time we don't know how to do that.

And so the question is, you know, for the immune system, for example, can we externalize the sort of computation that the immune system has to do? How do I make an antibody? That's the right shape to attack this Antigen. How do I make this T cell so that it goes and kills the bad cells that I don't want, and so on.

Well, you know, one can imagine a time when that process is externalized, I mean, that's already started to happen.

For example, car t therapy for sort of liquid cancers, and so on is something where you're assembling a cell. That's a type of T cell, you're assembling a cell that is specifically targeted to go and sort of attack the tumor cells and so on. And I think that's

that's a case where sort of the intelligence has been externalized and you're growing cells in test tubes, and so on, and then reinjecting them. And I suspect that's sort of part of the part of the general future of biology is to what extent can you externalize things that right now our bodies have to sort of figure out for themselves. But one can imagine a future where that's been externalized. But then the question is, how do you reinsert

whatever you figured out back into the body? And that's that's the part that's difficult. And I mean, I suppose, in a sense, with some kind of

proteins, one is already doing that. I mean, if one has, you know, an insulin pump or something one has kind of instead of using the pancreas to figure out. Oh, there's a high glucose level. Let's secrete insulin. You're using some continuous glucose monitor combined with an insulin pump to figure out externally by sort of using external computational intelligence. You're figuring out what to do without relying on the body to do it for itself.
and that's surely a coming trend in terms of a lot of kinds of medical interventions.
Let's see
Debug is asking, do you think in the future machines will hack into human minds to acquire consciousness like the matrix. I didn't think that the
the concept in the matrix was that the machines wanted to use humans to get consciousness. I thought it was something much more mundane like they just wanted to use humans as good sources of energy or something. I'm not sure. But in any case the question of will machines kind of
how do machines kind of symbios with with human minds?
Well, I tend to think that consciousness is not quite such a magic thing. It's really just the story of
kind of taking in all this input from the outside world and sort of kind of grinding it down to the point where we feel that we're having a single single thread of experience, and we're deciding what to do next. And that's something that machines can absolutely also do and sort of have the experience of consciousness themselves.
The question of whether we attribute consciousness to machines is a different question. That's much more a question of how we feel about things. I mean, the fact is, we only have the experience of things that are going on in our own minds.
But we kind of see people around, and the people kind of seem more or less like us, and so we can make the extrapolation, the sort of inductive inference that well, what's going on inside them must feel about the same as what's going on inside us feels to us. So we kind of extend our idea of consciousness that we innerly experience to other people who we see around.
Now, the question is, do we extend that sort of feeling of what must be going on inside to things like computers? Do we extend it to robots? That sort of seem human like. Certainly, as a practical matter, you know, one often finds oneself saying, you know, it's trying really hard to do that of the computer. Let's say, or it doesn't like that.
Those are very human kinds of statements, even though maybe that's only at a superficial level one's imputing to to the computer the same kinds of inner feelings that we have about ourselves.
My my guess is that it's just a matter of kind of one. If one is just having a text chat with some chat bot that is doing things that are quite human, like. I think one will tend to sort of impute to that thing much of the same, much of the same kind of inner experience that we each have ourselves, so to speak.
And I think it's more on us whether we sort of attribute consciousness to the thing than some kind of intrinsic character of the thing. I think the intrinsic character is rather easy to get. The only places where it doesn't
so much show up are things in the natural world where there's not the same kind of sort of crushing down of information down to this sort of single thread of experience and decide what to do. It's like something like the weather, let's say, you know, there are lots of different pieces

of the weather, lots of different pieces of fluid that do this or that, and clouds that do this, and pieces of rain that do that, and there's no kind of coherent picture

of sort of it's all making a sort of single thread of experience for the weather, a little different from the way that the computation plays out for for us

now in terms of how one can make sort of a combination between sort of the AI and the human. It's an interesting question. One of the issues is in our brains. There are 100 billion neurons that reach sort of firing, you know. I think a billion of them are firing at any given time. And the question is, what

how does one take those neuron firings and interface them to anything else?

One way we know to do that is, we can be talking. We can take all those independent neuron firings. We can sort of package them all up into this single thread of action, which is, you know, moving our mouths and and saying things. But there's a hundred 1 billion neurons in there that are all sort of doing their own thing, but somehow we're channeling them all down into this sort of thread of what we say

or what we do, how we move our muscles and things like this. So the question is, if we're going to sort of hack into our brains and sort of deal with those 100 billion neurons all by themselves. How is that going to work?

Is that all? A very incoherent pile of stuff that's going on, and the only time when it becomes coherent and communicable is when we've turned it into. Let's say a thing we're saying. So, for example, for a computer, you could say, Well, how does the computer interface to another computer? Well, it has all kinds of network protocols and things like this, where it's taking whatever internal information it had. It's packaging that up

into this sort of thing that's been going to be sent down the network channel to another computer that's going to unpack it. And so on

for the computer, we could equally well say, Well, why aren't you just taking the bits that are in the memory of the computer and transplanting them to another computer. Well, the reason is because the arrangement of bits in each computer is different. Even if they started the same, they have the same operating system, and so on, unless they had the exact same experiences where that interrupt came in at the same. You know, Nanosecond, as on another computer. The detailed arrangement of bits inside, the computer is going to be different.

And so there's no way of sort of taking the direct memory of one computer and just transplanting it into another. And so, similarly, for for humans and brains, the detailed configuration of that, that arrangement of of neurons is going to be different

now, could we expect an AI to sort of learn enough about the detailed arrangement of our neurons that it can kind of feed things in at that level without having to have sort of a neural user interface that sort of packages, things at the same level that we might have a graphical user interface. That's an interesting question. I mean, there are certainly cases where one's doing things like repairing spinal cords, where you're taking this bundle of nerve fibers. And you're saying, well, what does each fiber in that bundle

mean? Oh, well, it means move your toe, or something like this, then you can potentially learn that. But that's a very describable thing like, Move your toe.

I think it's much more difficult when it comes to sort of inner thoughts. It's really like trying to figure out in an Llm. What does every individual neuron and every individual weight in an Llm.

Really do? They all sort of add up to make that thought? But it's very hard to dissect it and say, this piece corresponds to this component of that thought.

I mean, there's some hope that you can get sort of a mechanistic interpretability for Llms. I'm pretty dubious about that being something that is generally doable. I mean the view that I have of what's happening both in Llms and in brains is that there are these kind of lumps of irreducible computation that do what they do.

and it so happens that you can often fit them together to do things that you want to have done. But each of those lumps of computation is some kind of thing that's very hard to understand, and the way that they fit together. It's like, well, they happen to be able to be stacked in this or that way to achieve the objective. But it's not something where one's going to have sort of a narrative explanation of what's going on.

So can one sort of hack into those lumps of irreducible computation, or to the interfaces between them, and usefully do things. I'm not sure. I think that's conceivable, I think by the time one's down at the level of individual neurons, it becomes much more difficult to imagine sort of having a meaningful communication there

that is going to be sort of coherent enough to usefully be able to to to do things with. I mean, it sort of depends a bit on the hardware level

of just to what extent can you actually get into the brain and measure what every individual neuron is doing? I mean, if you have just a surface where you have a bunch of microelectrodes. And you're just sort of reading out from that surface. Then you're just going to get one sort of one collection of neurons that happen to to get to that surface. If you want to get into the whole volume of the brain, you have to have a different approach. Perhaps you get to use magnetic

kinds of recording processes rather than electrical ones. Perhaps you can get all those neurons to start producing light with some kind of optogenetic approach

that rather than being using electrical signals. I'm not sure. And perhaps, you know, light goes. You know, a certain distance infrared light, at least in the brain, probably far enough that you can kind of see infrared light. At least you know deep enough

that you could sort of sense what was happening in these neurons, particularly if they were set up to actually produce light themselves through some kind of change in the in the genetics. That was again, it's a complicated thing to achieve, but is not inconceivable. I mean, maybe there was a way to actually detect

through, let's say, infrared light, the activity of individual neurons without having to modify those neurons. I'm not sure, I think going back, actually injecting sort of input into the neurons. I think they have to be set up so that they accept that input. I don't think by just shining light on the neuron type thing, you're not going to get it to do anything. You really eventually have to have an electrochemical process to do that.

Let's see.

Oh, let's just look at some questions that have just come in to see if they're ones.

medically, is asking, what do you think the 1st mainstream disease to be cured by current technology will be?

Well, I mean, there's been lots of progress in things like heart disease. Both in detecting things and in kind of just replacing the plumbing

to to improve things there. I mean, there are sort of endocrine type diseases that have

been in cases where there's sort of long
where the the lifetime of the
of the hormone is long. You just take a pill every day and it fixes it. You know, a typical case where that isn't. The case is type one diabetes where the lifetime of insulin is quite short, and that seems to be very well on its way through sort of stem cell based therapies well on its way to being able to just replace the beta cells and
and have have that have that work. So I think that may be one of the ones that's sort of early in in being able to be cured. I think also, it's looking promising to make these sort of changes to the immune system to deal with cancers that are blood type cancers where you can kind of get access to all the cells. It's more challenging when it's kind of a solid tumor, and you can't sort of get into all those cells in an easy way.
I think those are a couple of things, I think. you know, if if
there start to be artificial kidneys that work well, there's a whole class of diseases there that seem like they're they're on their way to being to being resolved, I think, when it comes to
You know, people blow hot and cold in terms of their beliefs about things like Alzheimer's disease.
and it's it's hard for me to tell what's what's possible or not. There.
I think, while there are other sort of slightly less mainstream diseases which are potentially fixable by gene editing of various kinds. Those tend to be not so mainstream and affect particular parts of the body, and so on. And perhaps I'm not thinking my diseases well enough. I'm trying to
trying to remember a few more. I think you know there are about 100,000 kinds of diseases that get classified in the international classification of diseases. Icd 10 or something.
I think my own efforts in understanding sort of the computational foundations of medicine tend to suggest that the idea of discrete disease isn't really right. There are sort of islands of disease, but then there are sort of continuous gradations from one disease to another, and I think that's the that's
that. That makes it more difficult to sort of say, will this be fixed or not? I think
you know. Sometimes there are questions like, is it really a disease? I mean, is it just, or is it just I think it's a disease if it kills you for sure. But there are other things where it's not quite so clear. What
you know, what's what's good, what's bad? And so on.
Let's see,
well, here's a difficult one. Jason is asking, what do you think will be the next big technological innovation and computer and communication tech. They're saying from the landline to pages to cell phones, to smartphones.
Well, you know, it's interesting because these different form factors.
they're really things where they attach to us humans in different ways, landlines. You hold something up to your ear, so to speak. You know, pagers, it's a thing where you don't have to be connected to a wire. The thing is sort of a thing that you can carry around anywhere.
you know, to phones to smartphones where we've got, you know, the touch, interface and so on. I think it's this question of well, what's the next interface modality? And you know, virtual reality has been a next interface modality since at least 1990. It hasn't quite made it yet.

It's sort of slowly getting there in some areas, but not in others. I tend to think the next thing that will really work is when there really are glasses that give you sort of an augmented reality display. That's that's really useful.

exactly what it will take to sort of adapt that I don't know. I mean, I always kind of think that there are moments where, you know. I think glasses will be back

as a people wear them instead of getting contact lenses and avoiding wearing glasses. Type thing. If there's an augmented reality display. Well, that's at least until you can get contact lenses that have augmented reality displays. I'm not sure how the optics works, for that. It's easier to imagine optics on a on glasses that either have, you know, that have a display. That, for example, is a laser that paints onto your retina or something to keep the thing in focus, even though it's very close to to your eye.

But in any case, the I think that's probably

the next obvious interface modality. Quite how you give input, whether it's talking. I mean, I would argue that in the sequence, from sort of landlines, to smartphones, and so on. That voice interfaces are definitely another sort of partial step in that direction.

And you know, another thing to realize is these different interfaces are good for some purposes and not others. I mean back when phones were 1st introduced. It's like nobody will want to do business on the telephone.

But then you find out there are lots of things you can do on a telephone that you can't really get around to doing in person. There are things that work well in person don't really work very well over the phone. These different modalities introduce different things you can do, but they don't necessarily subsume everything that was there before, and I tend to think.

you know, voice interfaces have added things you can do now that you couldn't really do before, but they haven't taken over. It's not like

you write code with a voice interface and so on, I think. And you know again, with augmented reality interfaces, I think that will be another thing where there's a bunch of things you can do like. I've always thought the post-it notes in your environment like you come back in the same room, and it reminds you oh, you didn't put back that thing that was on the shelf there, and it's sort of showing you that as you look at that thing, so to speak.

you know that's an example of a type of interface that we just don't have right now other than putting, you know, a sticky note on that thing to say, to say, to do that. This is a virtual such version version of that. And I think that's the type of thing that's 1 of many kind of new things that will become possible if you have this new interface modality. What other interface modalities are coming.

I don't know. There are direct neural interfaces which I was just talking about. I'm not quite as optimistic about those as being sort of any kind of mainstream story, just because, actually, we got these thick skulls that make it pretty hard

to sort of get information directly into sort of electrical style information directly into our brains. Maybe there'll be some clever engineering solution with infrared light, or some such other thing that gets around that I don't see that as a particularly near term thing.

I do think.

You know there are other.

Well, I suppose there are other types of things that are much coarser like, you know you do things to change your mood, you know, if you're if you're that kind of person, you, you know,

drink alcohol to change your mood. I don't. I don't happen to do that. I don't particularly like alcohol, but, or you, you know, take something to improve your mood. You're sort of operating at this chemical level. It's a it's a very coarse interface to the brain. But
you know, that can be thought of as another kind of thing like this. And and I suppose one can imagine a sort of science fiction future where there's sort of a smart
drug, like thing, or nanobot like thing that goes into your brain and sort of starts sort of activating particular areas. Again, I think that has the same problem that this direct neural interface has of well, the unformed thought, so to speak, that that sort of interacting with unformed thoughts is kind of hard to do.
Let's see
Well, Sneaky is asking, what do you think will be the most impactful applications of gene editing in the near future?
I'm not entirely sure. I think that various
sort of definitive, inherited diseases seem possible. But the thing you have to understand is that current, you know, Crispr-cas 9 type gene editing is is quite limited in what it can do. It can knock out sort of single
single bases and so on. And they're starting to be more technologies to do slightly more elaborate things, and and start to both edit things out and and put things in, and so on. And so it's somewhat limited in terms of what kinds of edits it can make.
So you know, if you have a change in your genome that has, I don't know something that's missing, and that causes the frame to shift, so that all the proteins downstream of that are all nonsense, because you know the way it works is every triple, every 3 bases on your DNA code for one amino acid, the Amino acid will be amino acids. Get strung together to make proteins, so if you have one missing base on your
DNA, you shift the frame. You no longer are reading the correct triple of bases, and so that totally confuses things, and the protein that you get is a mess. After that, potentially, I mean, it's a mess until you get to the end of that exon, that region of DNA that is going to make a piece of
the of that protein. It's all a rather complicated thing that goes through this this strange machine called the spliceosome, that takes these sections of DNA and splices them together to give the sequence that is going to make the whole sequence in the protein, and so on. But so bunch of details. But the main bottom line is, it's sort of an emerging technology. How can you do better editing more elaborate kinds of editing
with gene editing techniques.
And so then the question is, Well, what diseases.
for example, have are the most kind of just, oh, it's just a single snip that's gone wrong. A single, a single base that's that's out of whack. And and what can you do about it?
And there are some number of those diseases. But most diseases are pretty complicated and involve multiple different changes. And it's a more elaborate story. I mean, I was just looking recently at the one that I thought was really simple is eye color. It's a typical, you know, Mendelian trait, where it's kind of like, you know.
the the the brown-eyed parent and the blue-eyed parent
that's actually the situation in my family may come.

And now I'm going to have a quiz of the eye colors of my children, but which I think I could, I think I can tell you. But any case, the the you know brown plus blue supposedly makes brown, except doesn't quite, and so on. I thought this was a very simple trait.

I was looking it up recently, and like everything in biomedicine, it might start simple, but then has a footnote, and then has a footnote to the footnote, and then it has a whole chain of footnotes, and in the end it's a pretty complicated story, and it's not the case that you could like flip eye color by just changing one base, for example, in some particular gene.

So you know, it's complicated to know there are a few diseases. I guess sickle cell anemia is one of them where it seems like it's a very specific point change that leads to the disease. And there are some diseases of that are genetic diseases that affect

sort of kids as soon as they're born that are of that type. Sometimes those diseases are very bad, and the kid can't survive unless that's fixed. Those seem like good candidates for things where you might just be able to fix it and and move on. Type thing.

I think. But those things most of those things you don't the the you can't survive for very long with that genetic change.

So it's my my thoughts. But I don't know. I mean, I do somewhat follow this, but I don't. I don't think I know sort of what is precisely the leading sort of thing that gets fixed that way.

I mean, there are some. Actually, there's some eye diseases which I think people are talking about being able to to fix that way. But I'm not quite sure.

Let's see.

well, Elijah asks, what do you think about a future where most of our decisions are made by algorithms. Do you think that could be used in our favor to make up for our self-destructive tendencies.

Well, that's an interesting thought. I mean.

it's it's like, you know, you follow your GPS in your car. It takes you. It might take you through the better streets rather than the worst streets, the streets that have less traffic on them sort of spreading the traffic, so to speak. So in a sense. GPS

might be said to be sort of making the world a better place, because it's causing the one to sort of have less intense traffic jams. Of course, it might also be taking people through all kinds of residential areas where the people who live there would much prefer that there weren't so many cars going there.

and so on, or it might. If it isn't a well programmed GPS, so to speak. It might take you on that strange, obscure, dangerous dirt road, or something rather than on the more mainstream road. But in general one could say that

you know, GPS is sort of spreading the load by avoiding the traffic jams, and so on.

So could one imagine that type of thing for other things?

For sure, one can imagine some situation where sort of the thing that recommends. Hey? Go to this. I don't know. Well, it's already the case, you know. Go to this restaurant, go to that restaurant. You're kind of spreading the load, because you're being able to say sort of you're being able to tell people. Oh, this is too busy. Now go to this other place.

I think.

In terms of sort of are you going to do the right thing?

Well, you know, one can imagine a situation where increasingly people are sort of being suggested. You should do this. You should do that, and that's, you know, whether it's in augmented reality, or whether it's just talking to some system, some chat, bot, or whatever else. There's certainly places where it could suggest. Do this or that. The question is, are people going to be sort of so much kind of oh, the bot knows best. I'm not going to think that through for myself.

I mean, I could imagine something where you know the thing. There's sort of an augmented reality thing, and it's listening to what you're saying. And it it flashes up. You're being a jerk and you're like, Well, I don't care, or oh, I didn't notice I was.

you know. I think one can certainly imagine a situation where there's sort of prompting. Now. you know who gets to write the prompts. How does that feedback on kind of your personality, the personality, the collective personality of society?

I can imagine all kinds of of terrible dystopian kinds of things that could happen from the let's suggest to everybody how to be nicer to each other, or whatever else. I think that sort of a very minimal example of that for individual people might be, you know. Are you eating too much? Are you eating junk food too much

kind of like a reminder, you know. If if it if if one could see what one was eating, if a system could see what one was eating. It's like, don't eat that piece of junk. Food, you know. Eat, eat the apple, not the the French fry, or whatever

and or even better, the feedback loop where there's, you know, measurement of blood levels of all kinds of things. And it's like you need to eat that strawberry because you need more. Oh, what does strawberries have?

I've forgotten something obscure, like magnesium. I'm not sure what it is. Banana, you need more potassium something like this.

The

you know. I can imagine that being a feedback loop. By the way, that's another example of kind of this externalization of what would otherwise be the body's own sort of activity, like I might say, I'm craving a banana because something in my body is saying I need more potassium.

But it would be probably better. Probably one would do a better job if it was like measuring blood levels of things. And then some computer system was saying, oh, consider a banana. You might want to eat that.

So I think, you know, there's there's a it's an interesting thing. What one can expect to do is sort of advice given automatically about things to do. We've already seen a bunch of examples of that. We'll no doubt see more of that. Of course, things can go horribly wrong. I mean, you can be, you know, it's the social media feed that's telling you things that it thinks you want to see, and those things cause one to just get more and more enraged.

And that's the feedback loop that's sort of correct for the AI that's trying to get you to read more and more stuff. It's like, if the person is more and more upset, they're going to read more and more, or if the person is more and more indignant, they're going to read more and more.

I think those are things which are somewhat sort of commercially defined by the dynamics of what gets people to stay on that social media site. One can certainly imagine sort of

just taking a stand. Don't do that.

But that's a complicated thing because somebody's paying for that social media platform. It's a it's a complicated loop, you know. To me it's something of the story of what's free isn't free type thing. I've always been a believer in the idea that one should make things that the people who are getting value from them pay for rather than these much more indirect things that have evolved on the web and in the open source world, and so on, where it's like, well, we're going to give something away

to the people who are getting value from it. But then there's somebody else on the side, whether it's an advertiser or something else who is really the one who's paying for it. I've always found those things to be. I've always suspected those things are ultimately unstable. I have to say they've lasted longer than I might have expected, and you know I feel like the conservative old fogey, so to speak, saying, I'm just going to make stuff

that for the most part, you know, if it costs a bunch to make, we're going to have the people who use it pay for it, so that we have this kind of loop where the people who are getting value are providing the value, so to speak, to go back to make more value, so to speak.

I mean, we've done plenty of experiments and other directions, like Wolfram Alpha being free, or the free Wolfram engine for developers, for Wolfram language. Those kinds of things, but I still think that the tightness of the loop of people who are getting value are the ones who are paying for things is really a healthy thing to have, and other things are ultimately on timescales. I don't know exactly what they are are ultimately unhealthy.

And I think that that's something where you know this question of well, what

you know. How do you set things up? How do you set up that? AI? Are you going to tell people that. Are you quickly going to segue to telling them things that are in the sort of that you're giving it away for free to the person. So you have to make money somewhere, and then you're making money by, for example, trying to get the person to eat another hamburger, or whatever else it is. If the person themselves.

who was sort of getting the suggestions, was paying for the suggestions, then the then the goal is to get the best suggestions. So you go on getting paid for the subscription for getting suggestions, and that seems to me like a much healthier and more kind of more robust situation. But you know people have, you know, if you offer people, there's this free thing.

it might have some problems, but it's free sometimes. That's so much more kind of people don't think about the future so much. They don't think about the implications. They just say, Hey, it's free. I'm going to get that now without realizing that that has sort of downstream costs. But you know, I think that's the main problem in that direction.

follow up question from sneaky, do you think developing better gene editing methods comes from trying to engineer more versatile Crispr proteins?

I'm not enough of an expert on all the details of this. But yes, I think that

the the whole Crispr-cas combination. My impression is there's been all kinds of work that's been done on trying to find or create sort of more versatile such proteins. And that seems to be the direction that makes the most sense.

you know. It's always shockingly complicated when you look at what actually happens. I mean, the the Crispr proteins come from the sort of primitive, immune system of bacteria.

and it is rather embarrassing how often things that are used in biotechnology, so to speak, were just like, Oh, we found this obscure thing that's used by fungi or something like that, and we managed to recruit it to be useful for human medicine. It's the number of things that were sort of ab initio created is really tiny compared to the number of things that were found where some mechanism was found in nature.

And so so it was with with Crispr in general. It's now you know the question when you look at the mechanisms that get used, for I don't know protein synthesis. DNA repair all these kinds of things. You look at all these very complicated machines and things that are happening, and little you know, pieces of molecules that are being inserted and removed, and so on. It's a very complicated story. It's a very complicated, very elaborately adapted machine.

It's not the only way you could make such a machine, but it is a way that works.

although it may be hard to describe. It may be sort of an irreducible computation that this thing is flipped around that thing and that thing, and and so on, and in the end you replace that that you know. In the end you repair the break in the DNA Strand, or whatever else it is.

I think, in imagining inventing something that sort of does that from scratch. As an engineering matter that's proved very difficult.

I don't know to what extent one can sort of move a little tiny bit away from what's already been done. One could use actual artificial selection, actual biological evolution in test tubes, so to speak, to evolve to to a better something or other. That was the kind of somewhat dubious, you know, somewhat dangerous to do, you know, gain of function research on viruses to try and make a more potent virus, so to speak, by

progressively kind of evolving it, and just sort of saying, well, we get a whole generation of viruses. Pick the ones that are the most potent. Okay, you know, they will produce the next generation of viruses and go on. One might imagine. I don't know whether people have done that. I'm sure they have

for some of these things for Gene editing, or you can go and try and do sort of the simulation version of that, and say, how could you make something that would work better again? I know people are trying to do that. It's not something where there's been a giant cascade of success yet.

although maybe that maybe that will happen. I mean this general problem of sort of make the protein that will do this particular thing really well. The kind of idea that you can design a drug with a computer design system more or less, or with a machine learning analog of that, and then just deploy it and have it work.

It's sort of hoped that there will be a giant cascade of things that get invented that way. But that's been hoped at least 5 times in the past, with 5 different methodologies for sort of making drugs by combinatorial designing, combining things, or by, you know, doing computational chemistry, and so on. A bunch of different approaches, none of which have ended up working out great. Maybe the one the current one will work out better, I'm not sure yet.

Let's see.

oh, boy, another another question about AI from Max. When do you think we will get Agi and self-awareness in AI? How big is the danger of Agi and aware? AI, you know all these terms like artificial general intelligence, self-awareness. I don't think anybody really knows what these mean.

and I think that they're good buzzwords. They're good for kind of painting a sort of Science fiction picture, but I don't think that that you know the the goalpost of what it means to have AI has moved ridiculously over the last 50 years that I've been sort of involved with this this business. I think that

sort of if you want something to behave like a human make a human, so to speak. There really isn't a way to get sort of this, this computational thing that in all aspects is like a human. There

are lots of things we've done incredibly well at making things that are superhuman in many respects. I mean, I think Wolfram language is incredibly superhuman at doing all kinds of computations.
You know, Llms are kind of a bit superhuman in being able to do the linguistic things that they do.
It's these various components work that way, you know. I think there'll be things that are a little bit more shocking when we have humanoid robots where we solve the problem of sort of robotics, which is, I think, getting quite close to being solved. And they also have kind of, you know, brains that seem sensible. That will be a little bit more kind of oh, it's getting very much like humans or superhumans, and so on.
I think. Come the the question of
so what does it mean when our ais can do more? Well, it partly is a question of well, when the ais can do more, it presents more opportunities for humans to do even more. And there will, I think, always be this frontier of humans get to decide what to do, and the ais enable the automation that's provided by the ais enables what humans want to do.
The question of, you know, will the Ais turn on the humans? I think I've talked about this at some length before. I think the the answer to that, I think.
is that yes, you can do stupid things where you set something up to be? Oh, I'm going to put the AI in charge of this or that thing, and then you're going to be unhappy about that because well, actually, something happened that nobody expected, and the AI did the wrong thing. I think the human might do the wrong thing. You might put a human in charge of something, and they might do the wrong thing. But putting an AI in charge feels like
it's more somehow when you put a human in charge and the human messes up. You're like, Oh, that's just humans. When you put an AI in charge and the AI messes up, you somehow feel. Oh, we should have been able to know that better. I'm not sure that's really the case.
and if it comes to sort of the ais are kind of out thinking the humans and figuring out how to do terrible things to the humans. I don't really think that's a i don't think that's a thing that is sort of for the future. The truth is, there are plenty of ways in which sort of a computational system can absolutely, dramatically out think humans as they are today, and can do things where
it can sort of do bad things to humans. You know, it can be the thing that is
figuring out. I don't know something with some autonomous weapon or some such other thing that's absolutely possible. Today it's a question more of, do you choose to deploy it, and so on. And then the question of whether it in principle exists. And if you, I think this sort of question of if you put the AI in charge of this sort of whole chain of things in the world. Well, then, that might include something which you say. That's a really bad thing that that happened.
but lots more to say about this. I think I've talked about this at at some length elsewhere. Perhaps. Somebody can find the the previous yakking about that subject.
and I think I'm almost have to go back to my day job here. But let me just see if there's some any particular question that term? I think I can answer quickly.
That's always a dangerous concept.
let's see.
Hmm.
Oh.
well, okay, one more medical one from incog.

Do you think there will be something akin to gene editing technologies like Crispr that can be used to develop patient-specific drugs or treatments?

I mean, this is already happening, and it's it's a story of

Well, I mentioned kind of growing mitochondria something still a little bit in the future. Things like some of these immune therapies for cancer, where you are sort of growing. In that case.

in the immune system.

the immune system works by having sort of variable regions in particular proteins in the immune system. That's how antibodies work. That's how T cells work and so on. They have a region where

the protein has sort of an arbitrary choice of amino acids, and through processes we don't completely understand. There is when you have those different sequences of amino acids, those proteins fold themselves up.

and by choosing different sequences of amino acids we seem to be able to get sort of a universal collection of possible shapes for those proteins, a set of shapes that are sort of specific to all the different possible antigens that one can imagine. They're sort of a shape space, universal collection of sequence of proteins.

And it's

starting to be a thing that's just sort of on the edge of being something that's becoming mainstream to create kind of custom.

immune system cells that are targeted to a particular patient's story. Now, of course, it's worth remembering in general, the

in general, we all have different detailed.

Genome sequences. We have. What is it?

3 million or so sort of variations

from the, you know, the reference sequences, the the overall sequence of 6 billion base pairs. I think we all have about 3 million

that are sort of variable, and that, you know 3 million changes from the average of the average, so to speak, and those make us the particular individuals we are, and those changes regularly get exposed on the surface of every cell they get broken. The base pair sequence gets broken into little tiny sequences, or actually, and they get exposed on the surface of a cell through this thing the

Mhc major histocompatibility complex, otherwise known as Hla. It's a structure on the surface of cells that is kind of a clever thing, because it's kind of like you've got the cell. And it's got all this stuff going on inside. And it's like, let me tell you what's happening inside, says the whole biological system. Let me put these little barcodes on the surface of the cell to sort of tell you what's happening inside the cell. And those barcodes might

reveal that the cell was infected with a virus because it's got virus DNA, or something that is determining that barcode, if the cell's apparatus got got taken over by a virus, or it's just here's here's the sort of the barcode of what the cell is, and if that cell isn't a u cell. It's an antigen of some kind. It's a bacteria or something like this. Then the things that get exposed will be, not self

sequences. And and as you sort of grow up in your in one's early years, the immune system kind of learns what's self and what's not self when it makes mistakes. You get various kinds of autoimmune diseases and and allergies, and so on. The

Your! So that every cell is kind of exposing what kind of a thing it is, and it's kind of labeling. Is it a self cell or not a self cell?

And so some kind of cellular type therapies. The way they've been done so far they start being done is you take a cell line. You take sort of a standardized human cell line.

One of the most common is the Hela cell line that came from a tumor somebody had in the 19 fifties that has just been growing ever since, and labs have those particular kinds of cells that are used as sort of the base cells to do all kinds of things with them, and they can be turned into stem cells, and so on. But the kind of the idea

is those cells that are made with stem cells and so on from a cell line don't have. They have the the sort of identification from the original donor of those cells. They don't have your sort of barcodes to say they're cells of your own.

And so the question is, can one can one somehow make those cells be cells that look like they're cells of your own? There are really 2 approaches. One is sort of a chemical approach to just sort of say, get rid of all the self non self labeling on that cell so that you can just sort of. So it sort of slides into the body without being noticed, as, Oh, that's a cell.

It's got proteins on its surface that are non self proteins, and so on. That's approach number one. Approach number 2 is just make the proteins on the surface be self proteins, and that's a thing where the most common way to think about doing that is to make stem cells from the actual person who's for whom you're trying to build new cells. That's been the main approach. I don't know of a way to go and take

an existing thing that comes from a cell line and sort of edit those aspects of the cell. I think that will be difficult, I think.

because it's kind of there. Those those markers are all over the surface of the cell. And so it's an easier thing to just say, make cells make stem cells that come from the original. The person you're trying to give the cells to, and then do whatever procedure you're doing, to turn those cells into whatever a pancreatic beta cell, or whatever it is, rather than try to edit the surface proteins on a cell that has come from a cell line.

All right. That was very biological.

today. But thank you for asking all those questions. I'm

I think I'm as I realize I happen to have been working quite a bit on foundations of biology recently, and I sort of had this self image that I I didn't really

hadn't really thought much about biology. But actually, it's I've been reminded recently that in the 19 eighties. Particularly, I went to lots of theoretical biology conferences and did lots of things in that area. I just never got to the point where I felt like I could write much about it. So there's not much sort of written record of my efforts in the 19 eighties, but already from that time I was learning all kinds of things about biology. Although biology is a field that has advanced a lot in the last 50 years, I would say that

one of the things that never ceases to amuse me

is that when I was a kid learning biology, you would learn the function of the Organelles in a cell. You know the the mitochondria do this. I don't know the nucleus does this, but then there was the Golgi apparatus, and when I was a kid.

Now! Oh, I know this must have been well

more than 50 years ago now much significantly more. 50 something years ago. The the thing that one memorized was the Golgi apparatus is an organelle of unknown function.

Well, now, there's a whole section in the biology textbook about what the Golgi apparatus does. And you know what the endoplasmic reticulum does. It sort of amuses me that go back half a century. And you were literally memorizing. It's an organelle of unknown function. Well, its function is not unknown now. So some of what I've learned about biology. It is a fast moving field, and at least at the level of facts, if not at the level of principles. There's sort of new stuff to learn, and I've been been steadily upgrading my updating. My my knowledge of this stuff.

So it's thank you for asking questions about it. It gives me a chance to try and sort of piece together the things that I know and have, and have learned and figured out about biology.
all right, should wrap up for now. So look forward to chatting with you again another time.
But bye for now.