Hello, everyone. Welcome to another episode of Q&A about history… about future of science and technology.
And I see a number of questions that have come in. Let me see… See one here from Dazdian. What do you see as very real possibilities or plans for the future with regard to the blending of technology and human physiology?
I think I've talked about some of these kinds of things before, but let's see.
Well… Right now, you know, we've had…
Forever and ever, you know, we've had dental fillings since, what, the 4,000 years or something. So, there's some level of, you know, technology that, of that kind. We've had, you know, sort of replacement bones, we've had these kinds of things. You know, one area that is sure to
to develop as kind of replacement organs of various kinds. There certainly, is a question of kind of what
What can you replace
in the same form that it is in the body, what do you replace in a purely technological way? You know, if you're trying to replace, I don't know, insulin production by the pancreas, you might have that as a completely separate external thing that, where the signals are being measured, the glucose level or whatever is being measured, and insulin is being, sort of put in.
But that could be a, you know, it already is, a separate kind of box attached to you, so to speak. Whereas, if you're trying to replace, I don't know.
Something like kidneys or lungs or something of this kind. Well, kidneys, again, you can… you can do dialysis external to the human, so to speak, but it's not a convenient thing to kind of walk around with, so to speak.
Much better would be to have kind of a… an in situ sort of replacement for the function of a kidney. And there's some kinds of organs where the function is sort of
basically, chemical, mechanical types of functions, I think the,
And I think the, sort of, the sort of probably leading potential plan there is you… well, again, there are two possibilities. One is you make a technological replacement where it
just, you know, is pumping blood, but it works in a completely different way than the heart works. Or, number two, you're actually trying to make something that is very much like the heart, but maybe made of slightly different materials or whatever else.
And I think those are two slightly separate branches.
But the… for many of these organs, the kind of… to get the functionality that the organ has, it's convenient to kind of, sort of, copy what biology already did.
And I think the leading approach to doing that is to use very high-resolution 3D printing to kind of make scaffolds for these organs, where what's important is kind of the structure of the organ in terms of muscles or cartilage or whatever else. And then… then there's the question of sort of how do you cover
that scaffold with actual cells that can do the kinds of things that are needed in those organs. And I think the main plan for that would be you build the scaffold, then you fill in stem cells, you let them turn into the kinds of cells that line the lungs, or the kidneys, or whatever else.
And, you try and make something that is sort of a replacement part
that is a technological replacement part that sort of works the same way that the original thing worked. Each of these organs has all kinds of different issues. I mean, you know, with hearts, the main issue is preventing blood clotting and keeping, sort of, the flow of

of blood nice and continuous, and not letting anything stagnate, because as soon as things stagnate, they tend to clot, and then you have a bad situation. With lungs, I think the main thing is being able to have
Kind of the… the place where you're transferring oxygen into the,
into tissues, sort of the… be able to diffuse… it can diffuse only a small distance, maybe 2 microns or something, and you have to be able to have things, sort of your… your replacement tissue, accurate, sort of, to the point where you've got all these little tiny microscopic branching pieces that get down to that scale.
I'm not completely sure. I think kidneys, I gather, are also kind of somewhat
you know, one can imagine sort of 3D printing them. I think there one has to have kind of a filtering system that involves kind of a fine mesh filtering kind of mechanism.
I think the liver, for example, whose main function is kind of producing lots of chemistry for the body, my impression is that there's much less known, that's not as much of a sort of structural play. I mean, something like lungs
Once you have something with that many, you know, whatever it is, 23 levels of branching, you're down to these incredibly tiny channels and so on, you're going to be able to have oxygen transfer. With the liver, you're trying to produce all these different chemicals, and that's a complicated story that isn't just sort of the mechanics of the thing.
I would say, when it comes to other, sort of,
parts of us humans, again, they're things that are sort of purely mechanical, where in many cases there already are, you know, prosthetics that can replace these purely mechanical features of humans. I think there's a sort of big issue of how do you interface to that.
I mean, for example, if you have a damaged spinal cord, you know, can you interface to something which is essentially a sort of a kind of a thing where, you know, an exoskeleton or something that is controlled
by your brain, but directly neurally. And it seems like
things have been advancing along those lines. The main issue is you've got a bunch of nerve fibers, and you can potentially detect what signals are coming on each of those nerve fibers, but then the question is, well, which fiber does what? And I think it's sort of a mixture of being able to kind of learn,
being able to learn from, kind of, the person thinks, oh, I'm moving my right leg forward, or whatever, then you see what the actual, nerve firings are in that bundle of nerve fibers.
Or the person notices, if I think this, so to speak, then this happens in the kind of exoskeleton. And it becomes sort of a feedback loop where the person is learning and the system is learning. And that seems like a thing that has gotten to a certain point. I mean, it's definitely not…
all the way there yet. My impression is I don't know
My guess is that's limited by the kind of sensor technology of how accurately you can kind of get data from a large number of independent little tiny nerve fibers, which is the kind of thing that eventually gets solved. It's kind of a sort of precision engineering problem.
I guess that other, pieces of,
Well, that's… that's kind of on the… on the… that's one level of… of scale. I mean, if we look at, kind of, a lot of, sort of.
The ways that we include things in ourselves, a lot of that is at a chemical level, through drugs and so on, where essentially we're saying, send in molecules that are going to do the right thing.

if it's some kind of, oh, I don't know, slow-acting hormone or something, or, you know, where you just take a pill once a day, you are doing the same thing that, you know, gland X does in your body, you're doing by just getting the pill that came from the factory type thing.
And so, in a sense, you're including… you're integrating something from the external technological world, but it's being integrated, in that case at the molecular level, rather than at the level of a sort of a physical, mechanical thing.
My guess is there's sort of an intermediate between those things, which one's sort of waiting for, which is kind of the nanobot-type approach
Which is, you know, can you put sort of artificial red blood cells in the bloodstream? Can you put artificial white blood cells in the bloodstream? That's… so can you have a technological thing that is sort of floating around the bloodstream? Right now, the technological things we have floating around the bloodstream are molecules that come from drugs and so on.
But, you know, could you have your own technologically developed white blood cell, for example? That would be very interesting if it was possible, because the white blood cells are, you know, T cells and B cells and so on, are integral to the immune system, and one could imagine something where right now.
you're at the mercy of these still somewhat not very well-understood mechanisms of how the immune system decides, I'm going to attack this thing, I'm not going to attack that thing. Sometimes it makes too little of an attack, and the infection, you know, takes over or whatever. Sometimes it makes too much of an attack, and it attacks something that's part of your actual body, and then you get an autoimmune disease, and so on.
And so, being able to kind of technologically decide what do the, you know, what do the antibodies get generated to, what do the, what do the T cells go and attack, and so on, that will be highly interesting.
It's… it's not… we're not there yet, but that's the kind of thing one can imagine, where you're sort of integrating at the cellular level, so there's kind of integrating at the molecular level, you take some drug, you get those molecules in you. There's integrating at the more mechanical level of you've got, you know, an artificial hip or something like this.
And there's integrating at the intermediate level of, sort of, artificial cells or something like that. the technology of those things, I mean, again, there are different ways one can imagine approaching that. One is, you take existing stem cells.
And you try and reprogram them. Another is you make fully artificial cells. On the existing cells that you reprogram, that's kind of the whole idea of stem cells.
Where you're taking a cell from, well, pretty much anywhere in the body these days, and you are reprogramming it to be a cell that can turn into any kind of cell, and then you're trying to coax it through many days, weeks, months, to turn into exactly what you want, the pancreatic beta cell, the heart cell, whatever else.
There was some…
It seems to be the case that, you know, if you generate… so you can take a sort of a… you want a cell from you because you want cells that are, that have the same immune system kind of markers as your other cells, so that those cells won't get attacked by your immune system.
The sort of less sophisticated version that's more current practice is you take a cell from a generic cell line.
from, you know, cells that were derived from somebody completely different, and then you kind of have to have immunosuppressive drugs and so on to prevent your body from attacking those cells as being sort of foreign cells. But then, kind of the…

The approach there is take a cell that you reprogram, you know, let's say your immune system had attacked all your pancreatic beta cells, and you got type 1 diabetes as a result.
then can you replace those beta cells by putting in, by reprogramming cells to become beta cells? And that's something that's sort of just about possible now, and, you know, seems to be just about working now.
It was hoped at one time that you could take cells that might be, let's say, heart muscle cells. and get them from stem cells, just sort of inject them and say, okay, they're going to go to the right place in the body. I don't think that works as well as people had at one time hoped. I think it works better, for reasons I don't think are really well understood, for cells that are, like, secreting things at a molecular level, like secreting insulin, let's say.
That seems to work better, it seems to be less important where those cells are, they're just sort of sensing their local environment and deciding what to do, whereas a heart cell, heart muscle cell, doesn't really do you any good unless it's actually integrated into heart muscle.
So… That's on that branch. I mean, the other branch is, can you make an actual artificial cell? And there's a certain amount of technology that's been developed to try and make something which is sort of encapsulated like a biological cell is. Biological cells have this lipid bilayer surface that is something that you can more or less technologically reproduce. You're just making little bubbles, basically, with those kinds of surfaces.
Now the question is, what do you put into that cell
Both can you put in things that will be useful, and if you really, you know, by the time you're putting in something which can replicate, you might as well just be using a stem cell. So it's really, kind of, can you, sort of.
make artificial cells that have… that do useful things, and that's still pretty far away from the state of the art, to be able to do that. But that's sort of an interesting direction. I mean, in fact. the lipid nanoparticles that were used for the mRNA vaccines for COVID, for example, they come from sort of that technology. You're trying to enclose, so RNA is a very unstable molecule. in DNA, the thing that sort of stores our genetic material and gets passed from generation to generation, that's a pretty stable molecule. You can have DNA from, you know, 50,000 years ago that was frozen somewhere, and it'll probably still be, be reconstructible today. RNA, which is kind of the intermediate between the DNA, which encodes genetic information.
and proteins that are actually the things we're made of, so to speak. The RNA is kind of an unstable thing. Possibly it was, in the early history of life on Earth, there was… it's certainly imagined that RNA might have been sort of the original thing that was around before DNA started getting produced.
But in any case, RNA is fairly, fairly unstable, and so if you want to use RNA, and you want to get it into cells and so on, you have to enclose it in some kind of membrane, and that was what was done. It's got a very tiny micro-nano cells, so to speak, of which were essentially lipids, much like the, the,
The cell membranes of ordinary cells.
But so I think that's… that's, as I say, the artificial cells that really do things and have machinery inside, that's still kind of further away. I mean, the thing that is always the thing to remember is biology and us humans
are examples of, kind of, molecular scale computation. We have a lot of machinery that's operating at a molecular scale, and doing, sort of, amazing amount of stuff, given how small that machinery is.

We don't have a good technology stack for that yet. The only technology stack we have that does that kind of thing is basically biology.
I mean, there's no reason you couldn't make something technological with a completely different molecule that stores genetic information, completely different machinery for going from that, quote, genetic information to something structurally interesting. I mean, proteins are a particular set of chemicals that are sequences of amino acids that
turn out to make lots of different shapes and be useful. There's nothing that says, there isn't something that's based on, I don't know, silicon and boron and something else that would be a complete alternative as a way to make
Sort of molecules that are of specified shapes, which is essentially what proteins seem to serve for us.
So that, and then if one does that, if one has a completely different technology stack for doing those things, then there's a question of how do you interface it to biology with its technology stack.
And sometimes it might be easier, because, you know, most likely the immune system will just say, oh, that's not a thing I recognize, that's not a protein of some kind, you know, an alien protein.
That's just, I don't know, something like a piece of plastic, which the immune system is not going to, kind of respond to.
So, you know, it's, it's, so I think that possibility of kind of a different technology stack to make microscopic molecular-scale machinery is, is kind of another direction. And as I say, there are these different sort of integration levels of different scale sizes
for integrating things with biology. I mean, the obvious one that one can also talk about is brains. Where, you know, brains do things that are a bit like neural nets and computers.
And it's like, can you replace some pieces of brain function with some… something that is digital electronics? Can you make an interface between those things? I think this is… it's… it's a comparatively difficult thing. It's, you know, our brains were not made, really, to have internal interfaces. Our brains were made so that we interface to the world.
through, you know, eyes and ears and speaking and moving and all those kinds of things. The internal workings of the brain weren't really intended to be externalized. Just like if you have a program running on your computer, it's got some nice user interface that has menus and pull-downs and so on, but that program didn't ever expect you to go poking into the internal memory of the program
And, trying to interact with it that way. It had an interface, it had an external interface that was built for you to interact with, and so do brains. So I think it's somewhat non-trivial to go and intercept things in the sort of… internally, as it would be for a typical practical program running on a computer.
And then, sort of, the question of what does it mean to do that?
and sort of what kind of a thing do you have when you have this kind of mixed digital and human-like brain and so on? And what does it feel like to have something which is kind of connected to you? I mean, in a sense, whenever you control any machine.
You know, you've got some complicated joystick or whatever else it is, you're controlling it.
it is sort of like an extension of you, and I'm sure that the wiring in our brains, when you're, you know, flying a drone or something like this, I'm sure you're making use of the same basic wiring in your brain that you make use of for, I don't know, walking and…
gesturing and all kinds of other things. In other words, it really has been the case that

Through just holding those, you know, controllers.
you have integrated something outside of you into the same system that you're using internally.
And I suspect that if you had something that was a sort of more direct neural connection, it would feel kind of the same way. It's just like, I'm connected to, you know, I… in one case, you can say, well, yes, I'm moving my fingers to do that.
In another case, you know, one isn't usually conscious of when one's typing. It's just like one's thinking of a word, and the word is coming out through one's fingers. One's not usually, if one's a fluent typist, thinking, you know, I'm typing T-H-E, and so on.
So, anyway, a few thoughts on that.
Well, let's see, gosh.
jumping through questions here, but, I see Ivy asking, could life exist in forms we currently classify as non-biological?
That really depends on what is the definition of life, and people have had a hard time with that, because life on Earth is all connected. You know, all life on Earth has DNA and RNA. Well, it all has RNA, for sure. There are viruses that don't have DNA. It has cell membranes, it has things like that.
I mean, all life on Earth seems to be derived from
It's sort of last universal common ancestor, which was around about 3.7 billion years ago, people say, and sort of all life on Earth derives from that. Now, Luca, the last universal common ancestor, probably beat out a lot of also Rans.
as other possible, you know, seeds for life on Earth.
And we don't know what any of their physics or chemistry might have been.
I think that, it's a very interesting question whether some of those branches that Luca kind of crushed
Maybe they still exist somewhere. Maybe, you know, there's some weird microorganism-like thing that lives deep in rocks inside the Earth that has been having a grand old time for the last 3.7 billion years, and it didn't know that Luca took over, so to speak.
That's a possibility. By the way, that's not an easy experiment to do, because, you know, you go fish things up from stuff that's, you know, miles down and so on. It's hard to avoid it getting contaminated, because, you know, there are kind of ordinary microbes everywhere.
That's one problem. The other problem is, when you get the thing, how can you tell whether it's a living thing? What is the definition of a living thing?
You know, that's been something that has been very confusing. Life on Earth, as I say, very much defined by the particular details of chemistry and biochemistry that we have in the particular chain of things that our life as it has been progressively evolved on Earth.
People have sort of said, well, self-replication, that's the hallmark of life.
Of course, there are computer viruses that do that. There are, you know, the famous example of something very non-biological that does that is fire. If you have some fire, it will make more fire.
The, and then there are also, organisms like mules, for example, which are very definitely living, but nevertheless aren't capable of reproduction.
So that definition is sort of a guide… some… provides some guidance, but it's not the whole story.
And then there are stories about how, you know, biology is something that organizes the environment into some form that was, kind of, you know, structured form. But, you know, crystals do that too.

Crystals, you can start off with a tiny seed crystal, and it'll gradually add atoms in this lovely arrangement that is… that makes this perfect crystal.
So, most of these definitions don't really work.
I think, actually, in some of my recent work on what I've been calling bulk orchestration.
the way in which, at a molecular scale, sort of processes get orchestrated in biology. Kind of at the molecular biology level, sort of one molecule is actively moving around some other molecule and positioning it just right so that it can interact with the next molecule, and so on.
This sort of stack of bulk orchestration that seems to exist in life, that might actually be the main hallmark of life.
It may be that we are at the top of, essentially, a technology stack of bulk orchestration that… that has
that necessarily took a certain amount of time to evolve in biological evolution. That there's sort of a certain burnt-in, kind of, effort, essentially, irreducible computational effort that was kind of burnt into the forms of organisms that we now have.
And that may be, I think it is a serious possibility, that that really is the actual hallmark of abstract life, so to speak. And so then we go, you know, Planet X or whatever, we go and we scoop up some soil, or we gulp something in from its ocean or something, and then we say, is there something alive here? Even if the chemistry is completely different.
seeing some kind of bulk orchestration where there's a big stack of capability there would be an interesting sign. I mean, if we compare
Biochemistry with geochemistry. What's happened with rocks, so to speak?
By the way, rocks on Earth are much more diverse than rocks on the Moon or Mars, because a lot of rocks have been involved with things that are biological processes. Even the fact there's so much oxygen in the Earth's atmosphere, that's presumably a biological origin.
So, in a sense, the biology has been a thing that has had chemical effects and geochemical effects on the Earth.
But, you know, you can ask the question, left to its own devices, would kind of geochemistry lead to, sort of, an elaborate layers of organization, any kind of bulk orchestration process? It doesn't seem so.
It seems like that's something that is a special to something like life.
So that's kind of the, but, you know, I think the life as we know it today probably isn't the first form of life
that arose on the earth.
you know, we might be life 6.0 or something, and there are previous versions that provided, essentially, the templates that allowed life of the kind that we know with DNA and all that kind of thing to be generated. It could be that we needed some kind of, you know, pockets that existed in some other kind of molecule to be able to form the things that will become DNA and proteins and
things like this. That's one feature. I'm sure that's the case. I'm sure we were not…
the first form of life on Earth.
I'm sure there were things that were fairly viable as
sort of lifelike things that were different in detail. I mean, it would be completely remarkable if the ribosome, the thing that transcribes RNA, well, ultimately DNA, and makes proteins, it's a big, complicated molecular machine. That didn't just, you know, plop into existence from nothing.

That was probably, you know, that is… it's not ribosome 1.0. It's probably the result of many, sort of, little pieces that eventually assembled into the ribosome. By the way, all life on Earth uses ribosomes. It uses the same basic molecular machine
To, to, to, to create itself, so to speak.
So… I think this, you know, we're not the first form of life.
The question is, were there other branches of life, and did any of them survive?
And we don't know. I mean, so far as we can tell right now, we're the only ones, so to speak. But I'm not sure that that's… it's a very interesting experiment to try, and it's something that, you know, I don't think in any of the, environments that we're used to, that,
That one can… that there's sort of…
different life, alt… alt life, or something, is, exists there. Now, having said that, you know, I have to just explain a little bit about some history. You know, people have known about kind of cells of the kind we have, eukaryotes. People have known about prokaryotes, cells of the kind bacteria have, but it wasn't that long ago, how long was it, maybe 40, 50 years ago, that the archaeobacteria were discovered. They're a completely separate branch of life. They're a completely separate kingdom of life.
And it's not like our kaleobacteria are very hard to find. It's not like you have to go to obscure places to find them. They're all over the place. It's just that they were not, sort of, there wasn't a good assay that could pick them out from the other stuff you get in a scoop of pond water or whatever.
And the same was true, for example, I don't know, prions, for example. Prions are little pieces of protein that flop around.
And that affect the folding of other pieces of protein that can lead to nasty diseases and things like this.
But again, those are pretty, you know, they're pretty common, actually, but they weren't identified as a thing. It was just like, well, that's just stuff we don't care about, so to speak. You know, we're in the scoop of pond water, or whatever it is.
I think this is an interesting general point in science, that, you know, there often are things where there is some object or phenomenon that actually is quite obvious.
But you're looking at a thing that has that phenomenon, or has that kind of object, but the measurement devices, or the kind of way you have of thinking about things just ignore that completely.
I mean, I know in my own life and work, sort of a dramatic example of this is all the work that I've done on simple programs and how they do very complicated kinds of things.
People had seen that phenomenon for a long time before I really kind of tried to make a science out of it.
But mostly, they just said, oh, well, there's kind of this noise and messy stuff that happens. That's not what we're concentrating on. We're concentrating on something systematic that we can explain and that works with mathematics or whatever else.
And so, this effect that has been kind of, to me, is a very central one in science and nature and so on, is an effect that people just sort of, you know, that's not very interesting. That's just sort of, you know, random noise that we don't care about, so to speak.
So this is a pretty common thing in science, that until you… until, for whatever reason, you say, this is really an important thing.

it can be very easy to overlook it. Same thing happened with buckyballs, form of carbon. It would have been known that carbon atoms could be arranged in diamond, they can be arranged in graphite.
In, kind of layers, but then was realized that they could be arranged in these kind of, football-shaped, like, carbon-60 cages and so on.
And it's not like those are uncommon, you know, in random soot that you get. There will be those things. It's just that nobody really looked for them, nobody knew what to look for or how to look for them until, for independent reasons, people were thinking about synthesizing molecules like that.
So, I'm sure the same is true. If there was an alt life, On Earth.
we probably wouldn't have ever detected it, because it'll be something microscopic, and it's like, well, you know, using our immunoassay or something, nothing lit up. Well, it didn't light up, because just completely different chemistry.
But, that's, you know, maybe there's some clue that comes from, oh, there's some strange phenomenon in some place which hasn't been explained, and that's kind of the place to look for things like that.
Let's see… Gosh, many, many questions here. So if there are more biology ones.
Donald is commenting with respect to things I was saying earlier about stem cells, and saying. And things, saying maybe CRISPR is better, put the insulin gene into another tissue to avoid the autoimmune problem. Well, it's worth remembering that
the insulin that is made today is made by taking insulin genes and putting them into E. coli bacteria, and then growing those in big vats, and harvesting the insulin that comes from them. I mean, the earliest
The early, stage of genetic engineering from, 1970s.
is the thing that is used to make a drug like insulin by having bacteria produce it. I mean, this is kind of… if you can… if you have… you can sort of insert a gene that is going to make some protein, you want a bunch of that protein, this is a good kind of mechanism, is you just, you know, hitchhike on a bacterium to make that happen.
what has been done with gene editing is a different thing, something like CRISPR is a way of going along a DNA strand that's actually in you, and going along and noticing
oh, the DNA strand matches this template, you know, it's a long enough template, it's, I don't know, 12 base pairs long or something, the chance of that… the particular sequence occurring by pure chance, 4 to the minus 12, whatever, it's small enough
that you say, it's not going to occur unless it's the thing I'm looking for. And then what CRISPR does is it snips out
one of the bases, so it can… it can… it can change the genome by… by… I think it's by one base in the case of classic CRISPR. But… so that's a… that's editing a gene in situ, so to speak.
The big challenge with that is just how much editing can you do? It's not like you can say, I want to make an arbitrary substitution. People are working on that, but it's not a solved problem yet. I mean, that's the thing where if you say, you know, as happens with all of us, you know, in your 6 billion base pairs, I can say that everybody watching this has something that is like, oh, it's a mutation that might be bad for this or that.
reason.
we all have lots of, kind of, oh, that isn't the best mutation to have type thing. And, you know, if you look at a gene, you know, a typical

fairly long gene, it'll be peppered with thousands and thousands of changes to every one of your genes. You know, we all have slightly different changes. Typically, most of those changes, you know, come from our ancestors, so to speak. I think we each have
maybe a little less than a million places which were specials to us, so to speak, that nobody else had before. But, you know, so we all have lots of mutations.
And, you know, you could imagine, well, let's go fix all the mutations, let's go fix, you know, millions and millions of mutations for everybody, and have everybody have the standard human genome.
Actually, I think.
that the, when the Truman Genome Project was done, as usually is the case with these things, the person who really made that work was Craig Venter, and it was… a lot of it was Craig's genome that was the thing that was used as the reference human genome.
So, you know, if one went and edited everybody's genome to be the standard human genome, we'd have a lot of Craig Venters running around. But it wouldn't work, because the editing is sort of by one site at a time, and so on.
It's, and this question of whether, you know, whether one goes in and edits, oh, yeah, you know, it feels a lot like…
one of these things where you kind of… one day you go to, kind of, the genetic designer, reminds me of Blade Runner or something.
And they say, oh, looking at your genome, you know, we can improve it. We can make this change and that change. I suppose it's like people doing, you know, interior design, coming into your house and saying, you know, if you paint this that way, it'll be better, type thing. It's like, well, let's make a change to that gene, you know, we can change your eye color to this, or whatever else. It's sort of… that's an image of the future. I think we're pretty far away from that at this point.
Let's see…
Let's see…
Well…
It's a question from Thousand here. What do you see the future of computational biology and bioinformatics in general?
So, well, there are a bunch of different things. I mean, to explain bioinformatics.
You know, all of us humans got 6 billion base pairs in our genomes.
A lot of what people are interested in is, okay, you know, who's going to respond the best to this kind of drug?
who's going to have a risk for this kind of cancer? You can go and look at the genome and try and compare what people have in their genomes to kind of what actually happens to them, sort of, medically as humans.
The one case which is very, very clear, where you're looking at genomes and trying to figure out what's going on is in the case of cancer, because the genomes of cancer cells are different from the genomes of the original host.
that produce the cancer. That's sort of the hallmark of cancer, is the fact that it modifies the genetics. Sometimes it starts off modifying it just a little bit, and then things really spiral out of control, and you've got crazy numbers of chromosomes and all kinds of things like this going on.
But, so, you know, most of the human genomes that have ever been sequenced are sequenced for… it's the tumor genomes that have been sequenced, because depending on exactly what modifications the tumor made in its genome, that affects what kind of chemicals

can be used to attack that tumor, because it's kind of, you know, what we need to find always is, how is the tumor different from your ordinary cells? And if the thing has some particular kinds of, kinds of genetic changes, it will produce different proteins, it'll have different characteristics, you can attack it that way. So, a lot of bioinformatics has been looking at tumor genomes. Looking, to some extent, At the genomes of, sort of, non-ordinary genomes.
And, asking, how are those genomes associated with actual traits like response to drugs, or risks of diseases, things like this. The main message of that whole business, that second thing that I'm talking about, is it's really complicated.
In other words, people had imagined that, you know, something like eye color would be, oh, you know, old Gregor Mendel had sort of figured out something that was like the inheritance, more or less, of eye color. Like, you know, you have,
you know, you have a brown-eyed parent and a blue-eyed parent, then the kids will have brown eyes, because of the way that, you know, genes are dominant and recessive and things like this. I'm trying to think. I have 4 kids, and I'm trying to think, and my wife has blue eyes, and I have brown eyes.
And…
This is a bad test for the parent of knowing what… oh yeah, I do know what color all my… yes, okay, so my kids have not done the perfect Mendelian kind of thing that one would expect from one brown-eyed parent, one blue-eyed parent.
But which is typical of what one finds in, kind of, genomics like that. It's… it's more complicated than you think.
That is, people say, oh, if you have this particular change, you'll be lactose intolerant. But then you're not. And it's like, how can that be? Well, it's because there's a secondary pathway for digesting lactose, and you must be using that, and so on.
It's really complicated, and people have thought 20 years ago, maybe, that it would be kind of a very straightforward thing, like, you see this, you know this is what's going to happen. It isn't that way, it's much more complicated than that. So that's been a bit of a downer for that field. I think that, in,
When it comes to,
Well, one of the big frontiers there is the immune system, because in the immune system, you are getting
There's kind of your standard genome, but then in the immune system, you're always generating new random little tiny pieces of genome that are used to make all the different antibodies, you know, whatever it is, 10 billion kinds of antibodies.
That, go… are always kind of watching for… for antigens, for… for nasty stuff.
that they can attack, so to speak. And once they find something to attack, then more of those antibodies get produced, and you mount an immune response. But the thing that is sort of a… one of the frontiers of sort of bioinformatics is don't just, sequence
the genome, you know, your main genome, ask the question, you know, of the T… of the B cells, T cells, whatever, you know, what… what are the genetic variants in their… in their variable regions, so to speak.
You know, do you have lots of T cells that are trying to attack this? Do you have lots of B cells that are doing that? That's one of the frontiers, and it's something I would expect to advance in coming years, is kind of getting a much better map of the immune system, and kind of… because that is something that is genetically kind of visible, so to speak.

There's also the… another frontier is kind of… there are maybe 10 billion… 10 million species on the Earth.
Give or take. I mean, if you include viruses, it's a different story, but of sort of reasonably substantial critters, there's perhaps 10 million. There's maybe a couple of million that have been given names in any serious way. And the question is, well, what stuff was discovered by biological evolution on Earth?
And a lot of useful drugs have been found by, you know, looking for… I know a famous one was the bark of some tree in the Amazon, made an important chemotherapy agent. People look… for things, you know, scooping up ocean water and finding what kinds of things are in the, in the plankton there, or looking at, fungi and looking what kinds of things are there. It's kind of a what did life already discover?
I mean, in the history of life on Earth, maybe 10 to the 40th organisms have lived.
That's a lot of biological evolution, that's a lot of computation. It's not that much computation. I mean, in the sense that
in, let's see, in the Earth right now, there might be, let's say there are 10 billion computers, so that's 10 to the 10 computers, and those 10 to the 10 computers, let's say they've got multi-threading and things like that, let's say they're running 10 to the 10 instructions per second, that's 10 to the 20th
instructions that are a bit… being run per second in the computers of the Earth.
And so every year, that's 10 to the 27, kinds of operations. And, well, okay, so it's still, you know, 10 to the 40th organisms. Each organism effectively did a whole bunch of computation. before it got to be, you know, before it got to the next generation. So there's still a certain gap between those things, but we can… computationally, we can probably make things a lot more efficient than going through the whole machinery of life
With metabolism, and it has to eat things, and all that kind of stuff.
And so, you know, so there's still a big gap between the 10 to the 40th organisms that have lived and the 10 to the 27 operations per year done by the computers of the Earth.
But by the made, by the engineered technological computers of the Earth.
So, I mean, there's, that's… that's sort of another…
frontier, I suppose, of bioinformatics is looking at all the different possible critters, that have… that biological evolution has produced.
So then when it comes to computational biology in general, I mean, I've been very interested in the kind of foundational questions there, and sort of what the underlying way to think about biology computationally is. You know, what is happening in biology? How do we characterize, at a bulk scale, what's happening in biology? You know, when… if you, you know, any organism
It doesn't, you know, when we, if we sort of look at what's in
sort of living tissue, it's not a random liquid, where there are just a bunch of, you know, molecules bouncing around randomly. It's a very complicated thing, not well understood yet. Where there are all these sort of complicated chains of molecules and complicated interactions between molecules. I think people think of it as being a story of chemistry. Chemistry is really, you know, your molecules are bumping around… bouncing around randomly every so often, some of them interact.
The rate of interactions in chemistry is just… depends on the total concentration of molecules. I think there's probably a whole theory of subchemistry in which you're really paying attention to what happens to every individual molecule, and how does it… it interacted with another

molecule. Maybe it's going to see that molecule again, rather than ordinary chemistry, which just says you meet strangers all the time as molecules, so to speak.
So, I think there's some… well, there's a lot to be figured out about, sort of, the extent to which one can think about biological… biology computationally.
There's sort of larger scale thinking about biology computationally, where you're just saying, you're just thinking about, you know, what's the… let's say in the endocrine system, you're saying, oh, such and such a hormone was secreted that produces the secretion of this thing, which decreases the level of that thing, and pretty soon you're writing down differential equations that describe how the level of this goes up and down, and the other thing.
That's a very coarse description of what happens in biology. It may be adequate to figure out some kinds of things, but there's a, you know, huge depth below that going down to what happens at, sort of, the individual cellular and subcellular level. I mean, I think an example of this might be circadian rhythm.
Where, sort of, there's some part of that that is kind of a differential equation, and you can kind of describe it on a macroscopic scale, and involves pituitary glands and things like this.
But then, in the end, every cell in your body has some entrainment with the circadian rhythm of, you know, day and night and so on, and its mechanisms are something much more… much different and much more sort of in the… in the bulk orchestration subchemistry kind of area.
Let's see…
Killer is commenting, I think no one knows the long-term consequences of gene editing. I think that's true. I think gene editing, I'm a little less…
I would be a little less concerned. Something like sickle cell anemia, which is, I believe, a single mutation, and I think that's one where there have been efforts to do gene editing there. I think I mean, maybe I'm wrong, but my feeling is that's likely to be fairly simple.
When it comes to editing something which is sort of an upstream effect of some complicated downstream collection of things,
That's…
you know, that's much more difficult. I mean, when you just say, hey, I want everybody to have blue eyes because they're… because they're…
Because they're cool, and then you discover, well, that has the side effect that…
You know, people get this or that thing, and, you know, people who actually had blue eyes in the lineage of their ancestors
They, you know, compensated for this by having more of that, and, you know, somebody who has less good, you know, DNA repair, their lineage compensated by having better immune response, and, you know, all sorts of things. It's complicated, and I think going in and just saying at the end, make this edit to something which is part of one of those more complicated systems. Seems, indeed, fraught with difficulty.
I mean, when it comes to stem cells, which is a different case where you've got a cell with… you know the genetics, and you're just trying to reprogram the cell to say, I want you to be like, you know, a fertilized egg cell that can turn into every possible kind of cell in the body.
There is a kind of weird sequence of chemical treatments that you give to a cell, so-called the Amanaka factors, that make the cell revert to its pluripotent stem cell state, so that it can then be differentiated into other kinds of cells.
There is a nasty effect

where the cell that has been reprogrammed, that has been deprogrammed, somehow still remembers that it was a skin cell once, even though you think you've wiped it clean with all these chemical factors. That's one of the frontiers of that kind of thing these days.
Okay, Joel is asking, isn't photosynthesis the best power generator on Earth? Can't we genetically engineer stuff to make solar cells that are based on photosynthesis? I have wondered the same thing.
You know, there is an electron transport chain in, in plants that is basically,
It's… it's using light from the sun, you know, absorbed in the green, particularly, because green is kind of the peak of the solar spectrum.
And, you know, there's a weird sort of quantum mechanical thing where there's a magnesium atom sitting in the middle of chlorophyll, the main active ingredient of photosynthesis, and it does some very quantum thing
That leads to, eventually, this whole chain, that, produces kind of power from photosynthesis. I just don't know what you need. I mean, I've wondered the exact same thing. In fact, I asked somebody recently who
said it doesn't work because… and I think the main because is that you end up,
The claim was, which may not be correct, that you end up with sort of a support system for your chloroplasts
that is kind of a whole plant, and it's then, what do you gain? I mean, you might as well just use the, you might as well have the trees grow and then burn the wood, type thing. But, I agree, it's a thing I wonder about, and it seems,
Seems worth pursuing.
Katie is asking, do you think we'll discover more significant species on Earth, whether new or just hidden from humans, bigger than bacteria?
Yeah, I mean, there's a… there's a steady rate of people discovering beetles in the Amazon jungle and so on. That's the most… that's the… that's the place where there are many little micro-niches and sort of different beetles emerged in this area versus that area.
Then the question is, is it really a different species? And, you know, can you really tell there was no, and depending on how you define a species, you can define a species by its DNA is fairly different, or you can define a species by, sort of, the very classic, it doesn't mate with the other thing.
you know, there is increasing evidence that, for example, humans, modern humans, mated with Neanderthals, who were genetically different.
And so, you know, there's… that's not… that's no longer perfectly the test, but then all of us humans are different, you know, one from another, unless we're identical twins and so on.
The, so… but we don't think of ourselves, or even, you know, we don't think of ourselves
Usually, as different species, even though, you know, different ancestral groups will have different characteristic forms of genetics that we can tell from so-called haplotypes and so on, that,
That one can identify. And so, you know, this question of, well, what really counts as a different species, it's a little bit of a fraught question.
But, you know, there will be things, and even when you say, well, that thing looks really different, it's got, you know, yellow frills, whereas all the other things that were sort of close to it were kind of
brown, unfrilled creatures. But you never know, because it doesn't take much of a change to go from being, you know, your average

I don't know… Fox or something, to an albino fox that's completely white.
And where you might say, gosh, that's really different, that's sort of a different species, it looks so different. But actually, that might be a pretty small genetic change. So it's a little bit hard to know what the borders of species are.
But in terms of a species where it's like, gosh, we never knew that that existed, I suspect there are still some to find. I mean, there are species which exist only in one particular pond, basically, somewhere on the Earth, particularly when there are fairly closed ecosystems and caves and so on.
There are… there are places where one can certainly imagine that things could have been hanging out for millions of years in their own little ecosystem. It's interesting to study those things, because it's actually pretty hard to make a self-contained ecosystem.
So I would guess that there's still something to find there. I mean, in terms of things that are microscopic, I'm sure there's a bunch to find.
I'm sure there'll be a whole sequence of things that get found of, you know, these proteins can sort of cooperate with these other things to effectively reproduce that protein.
Even though it's just a protein, it doesn't have all the reproduction apparatus, it can kind of piggyback on other things. Just like mitochondria, which are the sort of power sources for our cells, for
for cells of, of, critters, animals like us. The,
those were presumably, at one time, a separate, freestanding organism, a cyanobacterium, I think. And at some point a couple billion years ago, or more, you know, mitochondria started just hanging out in our cells.
But by now, mitochondrial replication relies on a bunch of genes that are actually part of our standard genomic apparatus, not just genes that are part of the independently replicating mitochondrion. So, very quickly, when things sort of get entangled
Between even things that might have been separate species, or separate kinds of forms of life.
one would think they… you know, it'll kind of… prions don't really do anything on their own, but when they're in this situation, they do this. Even viruses on their own don't really get very far. Viruses try and, kind of hijack
the, the protein generation apparatus of cells to, say, make virus proteins instead of making the proteins you were supposed to be make… making for the original organism. That's how they work. But as freestanding things, they can't replicate themselves, they're kind of, they're not really… not really…
separable, so to speak. So I think it's,
It's… it's likely there'll be more discoveries like that.
Jamie is asking, would you say biology has some…
Had its computational revolution, or is it still coming?
I think that… I think it's still coming.
absolutely still coming. I mean, there's mountains and mountains of data.
automated data-taking instruments that can sequence DNA very quickly and get, you know, billions of DNA base pairs generated. That's all, you know, robotic labs, things like this. There's just oodles and oodles of data in biology.
What does it mean? What are, kind of, the big theories of biology? I think we don't really know that yet. I mean, I've been working a bunch on that, and I think we have been making progress at understanding, sort of, what big things can you say in biology?

You know, one big thing you can say is natural selection works. Another big thing you can say is we're digital, in the sense that we, you know, DNA is a digital kind of thing. You know, back in the day, you would talk about, you know, bloodlines mixing, because people imagined that what carried stuff from one generation to another was something about blood.
Where you were just sort of mixing liquids together in whatever proportions.
it's not like that. It's, you know, there's a definite sequence of base pairs that you can kind of write down as digital information on a computer, but there may be some other pieces as well that get passed down in the structure of the egg cell and things like this that aren't part of pure genetics. But in a good approximation, biology is sort of born
born digital in that sense. But the question of what sort of big things one can say about the organization of biological
Systems and cells and so on.
there's… there's a lot missing there. I've been trying to do some… some work to fill that in. I think that the stories of what… how that works are deeply computational.
They're things that are about computational irreducibility, about all sorts of phenomena that are essentially fundamentally computational phenomena, phenomena that haven't really been explored or understood in mathematics.
phenomena that are really something new in the computational world, and biology is leveraging those things, and I think there's a big sort of computational revolution to come in biology once we understand that. I mean, the revolution in biology that came with
with theory of natural selection from 1859 was pretty big. The revolution in biology that came from understanding the digital nature of DNA in 1953 was pretty big. We're going to have more of that from, kind of, a large-scale computational understanding of biology.
Let's see, Killer is asking, you know what causes lactose intolerance? Can we cure this?
I think, I mean, the lac operon is… was one of the first identified, kind of, genetic control mechanisms for producing lactase, and I don't know how many,
how many base pairs the typical mutation corresponds to. So I don't know whether it's a CRISPR-accessible thing. I think it might be easier to take lactase tablets than to go start editing your genome. But, I think the,
My impression is that it's… I think it's a fairly small number of edits that would be needed, but I'm not sure.
Dazdian asks, could figuring out the blending of technologies and human physiology,
Lead to clones being more plausible.
Well, I mean, I don't think there's any fundamental technological problem with human cloning. I mean, it might have… you know, people…
People are always very…
ethically squiggly about these things, and perhaps for good reason. I mean, it's similar to, kind of, the inbreeding story. I mean, the idea that, kind of, inbreeding of, sort of.
Small family groups, sort of, breeding with each other and so on, that… that leads to bad stuff.
I think the claim that that leads to bad stuff is more embedded in human society than the science of it necessarily would confirm. I mean, it's just something we as humans think isn't really the right thing.
It's not, you know, people had always said back in the day, you know, a couple of centuries ago, people were saying, oh, inbreeding leads to, you know, weak, sickly children, and so on.
I don't think that's necessarily true. I mean, there were important cases, like, I think the British royal family had porphyria, which is a genetic disease, and there was a whole line of, of,

of, Of,
British, you know, monarchy folk who had that, and sort of inbreeding didn't help them in that case. And certainly, you know, inbreeding within a small population will always accentuate traits that, if there's some trait that's bad, well, you'll get more of it.
But my impression, I might be wrong, is that the things you see anthropologically, in terms of the sort of prohibitions of incest and the, you know, who can you marry, how close to, you know, can you marry a first cousin, can you marry…
this and that and the other. The, you know, those prohibitions, come very much ahead of it seems, ahead of, kind of the… and we can scientifically prove this is… this is wrong. It's just something that we as a species and society just don't feel is the right thing to do. I think that, when it comes to cloning, I think that people have
kind of… it doesn't seem like the right thing to do. I'm not,
I mean, I would say that ahead of cloning, where it's like a mini-me type thing, the, you know, the less…
well, I don't know whether it's less or more savory, but the different version of cloning is, hey, I want a replacement organ X, just make me a clone.
that will grow that organ, just like my organ X, and I'll be able to replace it. Now, sort of the theory of that, which gets very dystopian, is you make kind of a decerebrated clone that doesn't have a brain, or something, or doesn't have enough of a brain that you worry about
Destroying that organism, and you're then in a complicated, slippery ethical slope of, well, you know, you've got this thing that's just like you.
And I think people, as soon as the clone comes out, and it looks just like them.
even though its brain is scrambled in some way, and so, you know, you could argue it doesn't have any conscious, you know, experience, et cetera, et cetera, et cetera, my guess is people are going to be very, you know, they're going to say, wait a minute, you can't just, you know, cut out the heart of the clone and give it to me instead. It's, you know, there's something wrong with that.
Even if you had decided ethically that, well, the clone doesn't have a brain that could lead it to conscious experience, therefore it's just meat, so to speak, and we could do what we want with it.
My guess is that that kind of medical useful cloning would come ahead of
The mini-me type cloning.
I think that, the question of how you do kind of mini-me type cloning, you know, Dolly the Sheep back from the 1990s, was a case where people had said for years, oh, mammalian cloning is impossible, even though it was known that it was possible for things like amphibians and so on, which aren't really that far away.
But then, you know, sort of tricks were discovered that allow one to make these clones, and, you know, it has become a slightly routine thing for pets and, and I think horses as well.
that, to be able to make these clones. You know, one feels like there's something wrong with it, and so people will say, oh, there's going to be something very bad about the clone.
But I'm not sure if that's really supported at this point.
And so, you know, and…
you know, in parts of the world where, sort of, things are somehow quite government regulated, one suspects there aren't. There's not lots of, hey, make a clone type thing going on.
In other parts of the world, I'm not sure. And sort of the… as more biotechnology becomes more widely available, it becomes certainly possible to do that, I would think.
You know…

what's going to be wrong with the clones? I don't know. I mean, it could be, like, the stem cell story, that there's something weird about the clones. Certainly, Dolly the sheep seem to live a perfectly happy life, and my impression is that cloned cats and dogs do too, although I'm not sure about that.
I guess… let's see, I think some people I know have cloned pets, but I'm…
I… I don't know, what their story is, quite.
Let's see, Jonathan's asking, speaking of lactase tablets, could the future include putting tiny robots in tablets to set loose on the body to rewire things? Well, I mean, one technology that has existed for decades now is all kinds of
Timed release, sort of chemically modifying things in tablets.
So that's a very simple level of that. The question of whether you can sort of set loose something which is going to go
do things in the body, you know, the immune system is our best example of that. And it's like, can you make an artificial T cell that will go and be a killer T cell and go and attack certain kinds of cells that are marked in certain ways? We don't know yet. I mean, things like CAR-T therapy that's,
Chimeric antigen receptor, is that right? I forget what the CAR stands for, T-cell. This is a recent, very successful therapy, particularly for various kinds of blood cancers.
Where one is making, essentially, artificial… one's… one's making T cells that are specifically engineered, that they're kind of… they're T cells that you would make.
But they have been specifically set up to attack
the cancer cells, for example. So that's an example of kind of a… like a specially built, kind of a custom-built immune cell that was set up to attack a particular kind of cell.
Let's see…
Okay, so Doc, Doc is asking, thinking about replacement thing, but with brain cells, do you merge with the intelligence? That is a good question. I mean, it used to be said that after your, you know.
one, two years old, and you're a human, it's like there are no new brain cells. You're gonna have, you know, your 100 billion brain cells, that's all you've got.
you know, be careful of them. I think that's a good principle in general. You know, don't feed them weird drugs or chemicals, because you've only got one set, and once those ones die, you're done, so to speak.
It was discovered about 20 years ago that actually there is new growth of brain cells in the human brain in adulthood. Quite how much of it there is, I think, is still not known.
And it could be that those of us who try and exercise our brains a lot, just like, you know, you exercise your muscles a lot, you get more, more muscle tissue. It could be that those of us who try and exercise our brains a lot, you know.
doing live streams trying to answer, all kinds of random questions coming in. Maybe that's the brain exercise one needs.
to make more brain cells, maybe. But in any case, so new brain cells are made in human brains, and somehow those brain cells presumably integrate themselves into the brain as it's operating.
And it's an interesting question. I mean, one could think about artificial neural nets, and one could say, I've trained this neural net. Let me now add some more artificial neurons, and let me then… I don't know what will happen if you… if you do that, and I don't think anybody does that. I'm trying to think, is that an experiment one knows

how it would come out, to sort of… I mean, now, remember that artificial neural nets, at least in the current state of technology, have one big difference from brains, which is brains both learn and do at the same time.
Whereas artificial neural nets, they do their big training adventure, and then they get run. Actually, until recently, it was the case that more effort was spent training up
a chat GPT-like thing than was ever spent by all the people using it. I think that's now turned over. I think it's now the case that the training has gotten easier and more streamlined, and the usage has gone up.
And so it's now, you know, all those data centers burning power, they're burning it on people asking it questions, rather than it being trained to answer the questions, so to speak.
But in any case, so it's a little different.
But, so the experiment wouldn't quite be a doable thing yet.
I think, with artificial neural nets. But this question of, can you just feed yourself more neurons, and,
and sort of have that work well, I don't think anybody knows. I mean, I think that the idea… I'm trying to think. I think in Parkinson's disease, for example, where dopamine-producing neurons die, and you just want to get… make more dopamine.
I'm thinking that I've heard of stem cell-based efforts in that direction to make… to make more neurons. And since that's just sort of a chemical being produced.
that's an easier thing to make work, probably, than, maybe. I don't know, I don't know how much the electrical activity of the neurons is… affects the production of dopamine. Clearly, it affects it a bit.
Because that's, you know, it's… whatever you're thinking about is what determines, kind of, emotional responses and other things that are associated with the production of those neurotransmitters. But,
Or the kind of diffusion of those things. But in any case, I… I mean, this question of could you just clone your brain cells, make more of them, inject them in, and, you know, have a bigger, happier brain, I don't think we know at all at this point.
the… you can make brain organoids. People have made them, I don't know how the biggest ones are now, maybe a few million nerve cells,
you kind of have to have a life support system for your brain organoid that's kind of a copy of, you know, what the actual… what an actual human with kidneys and blood and glucose and all that kind of thing is doing, but you can make these organoids. How you communicate with the organoids is a whole different story. If you're using these organoids as a thing to do medical research, where you're trying to find sort of effects of certain kinds of drugs.
In the bulk on brain tissue, that's one thing, but if you're trying to say, can you have a conversation with a cloned brain organoid of yourself.
We're very far away from that right now. I know one company where, needless to say, it's the CEO's
brain, stem cells, well, stem cells that were used to make some brain organoids. I think you're… you're… I think… I don't know what's happened to these guys, but… but they were… they were talking about selling, kind of, things that were sort of
things like a server you put in a rack that has a brain organoid inside. It's a very science fiction kind of thing. But how you communicate with that organoid, that's very challenging. I mean, the, the,

The thing you might do is you kind of give it some stimulus, and then you feed it dopamine or not, and you kind of make it… it's like the reward system of the artificial neural net, but nobody knows how well that might or might not work.
Maybe one more.
question. We've done. I've had a very biological discussion today.
Brady asks, are we like microbes in a living galaxy?
Yeah, I mean, you know, there's 8 billion of us, humans, and
There's a lot more, you know, farm animals and things, even. But, there's… and there's vastly, vastly more microbes and… and, and so on.
I think that, you know, any one of us, despite our feelings of independence, so to speak, is part of a very big ecosystem where, kind of, in a sense.
We're part of the single, sort of, superorganism that is all life on Earth.
And you could say, if we look at us.
We have suborganisms, you know, every one of our cells has a mitochondria in it. We have, you know, all of our gut bacteria that are, you know, helping us digest things and so on. They are sort of other suborganisms living inside us. We are a suborganism to the kind of superorganism, which is, well, first human society, and then
the ecosystem of the world. So I think, yes, I mean, we're… we're…
We are… each one of us is just sort of one piece of that… of that big, kind of, hierarchy of… of things.
I mean, I think, you know, the level of it that we're at, we care the most about, I suppose, typically. And because that's, you know, that's who we are, so to speak.
Let's see, maybe one more question.
Alright, I'll try this one.
This is from Carol.
I'd like to ask a question about the future of science. Most natural sciences already have well-established methods and a solid, well-proven, and developed scientific theory.
But what do you think about psychology as a science? In which direction does it need to develop to be on a par with physics or chemistry in terms of the level of evidence-based methods and the level of a verified scientific theory?
Good question. I mean, psychology…
as a reasonably serious science, is only 100 and… I don't know, 130 years old or something, of that order. 150 years old, maybe.
And…
There are, you know, there are certain laws in psychology about, you know, ex… logarithmic responses to things and so on, which are pretty well-verified laws.
And there are… then, when it comes into, kind of, the biggest stack of questions in psychology, I think it gets kind of fuzzy. I mean, in other words, there are,
You can ask, sort of, what is quantifiably knowable.
But quantification may not be the story. I mean, you could say, what is quantifiably knowable about the shapes of animals? And the answer is, not a heck of a lot. Yet, there's clearly a lot that could be determined about the shapes of different kinds of animals and so on. It's just something about which there isn't…
a traditional scientific theory. I think that particular thing is very much a piece of, kind of, frozen computational irreducibility. There were things that happened in the course of biological evolution that managed to work, and now we have a frozen version of those.

So, in the case of psychology, I think the big break that psychology has in these years is LLMs and other kinds of neural net systems, because
they are, I think, close enough to brains that many of the same phenomena that happen in brains are going to happen there, but they are, in principle, much more easier to study. I mean, it is, you know, in a brain, if you want to record from every neuron in the human brain, that's completely undoable right now.
If you wanted to tell what every artificial neuron in a neural network was doing, well, you can absolutely read out that data.
Now, can you tell what it means? Well, that's where things have really fallen down, because in the case of the brain.
there are little circuits that we can tell, particularly in lower organisms, we can tell, you know, if you, I don't know, for a fruit fly, where the brain connectome is known, you can say, if the thing's proboscis touches a piece of sugar, then this will happen.
And you can kind of see the little circuit inside the brain, which is completely determined where, you know, this thing being stimulated causes that thing to happen.
I mean, in us humans, I think there are many… much less well-understood circuits. I mean, there are circuits to do with, I don't know, eye motion, saccades of eyes, and there are… I think there are circuits to do with walking and so on.
But, some of these very low-level, kind of functions, there are circuits known, but beyond those, it's like… and then it just goes into the brain. I mean, we kind of know, you know, typical human reaction time is about a third of a second.
you know, you hear a sound, you press a button, about a third of a second is what it takes. If you work out, in terms of how far the nerve impulse would have gone in that period of time.
If you were a, you know… well, not a giant squid, because they have very special accident nerves, but to be able to
transmit more quickly, but if you're… if you're a brontosaurus with a very long neck and very, very, you know, going a long way, you can work out how much time does the nerve signal take to go along your neck, or in the case of humans, you know, down your arm to your finger, and the answer is it's tiny.
And the effective distance gone
in the brain is certainly hundreds of meters. So it's as if the thing that you, you know, the signal you heard, and eventually the button you press, it's wiggling, you know, it's going around in your brain.
over, you know, lots and lots of times. It's not just, oh, you heard the sound, and then you sent the impulse to your finger to press the button. So we know a bunch of stuff is going on inside. I mean, we've got other reasons to know that lots of stuff is going on inside. What kind of a theory would one have about what's going on inside?
I mean, one tries to classify things at a very broad level, you know, personality types.
you know, various kinds of psychiatric conditions and so on. Well, it tries to bucket, you know, this is something that is characteristic of, you know, psychosis or something like this. They're very coarse kinds of classifications.
That, are… I mean, you can see in artificial neural nets, for example, if you modify the neural net, you can make it do all kinds of crazy things. And some of those crazy things look a lot like what one imagines, I don't know, psychedelic drugs do to human brains.
or things like that. It's, I think that the,

the thing that, so one can kind of see, one can, to some extent, buckets of the large-scale behavior of brains and neural nets. But then the question… part of the question is, what's the question? In other words, what are you trying to achieve by your sort of psychological analysis? Is it that you try to predict what the human is going to do? You know, if you do this and this and this, will the human buy the car or whatever else is going on? Is that what you're trying to do? Are you trying to say things about, sort of, what humans will do collectively in society? That might not even be a piece of,
kind of individual human behavior. You know, I think
I think one of the things that's always difficult in developing a science is knowing what the right question to ask is. I mean, you know, back in the development of physics.
in the big advance of physics in the late 1600s with Isaac Newton and people like that. Somebody figured out.
to some extent, Galileo, that looking at mechanics, an object, what does it do when you bounce two objects off each other? What does it do when you drop an object in the Earth's gravitational field?
Those are things that, it turned out, there were mathematical ways to address. If what Galileo had done was, let me look at the clouds in the sky and try and figure out
the mechanics of how the clouds and the sky work. Well, we still don't know that.
It's, you know, that's not something that would have been amenable to the kind of, sort of, just-give-me-a-few-numbers type approaches that worked out for Galileo, and then Newton, and then the whole span of physics now. So, you know, a big part of developing the science is to figure out what the question should be. And I think, in the case of psychology, or even for that matter, LLMs.
We don't really completely know what the question should be.
And I think one of the things, I'll give you an example of a question I was interested in recently, which was the following question. If you've got an artificial neural net, and it's just a random neural net, you know.
it's got random weights in that neural net. And then you say, I'm going to train this neural net, I'm going to make it do this amazing thing where it classifies cats from dogs or whatever else. And then you look at the resulting neural net. If I just hand you the neural net, and I don't tell you what it was trained to do, I just say, is this neural net trained or not?
Well, the question is, can you tell? The answer seems to be, yes, you probably can tell, but if you can tell, that means there are certain kinds of regularities in the neural nets
That might give you, sort of, some global law of neural nets that wasn't obvious from anything outside of that. You know, it makes me wonder about if you're given a seal or something, or a dog, and you're asked, is this a trained dog? How can you tell?
I don't know. I mean, you know, insofar as people train dogs in the same way, you know, sit, fetch, whatever, you could try those things on the dog. Maybe it's a dog from country X, where it knows those words in Kazakh or something.
It's, so you might not have luck saying that in English. But I think, you know, it's not obvious if you don't know, you know, what the thing was trained to do. It's not obvious. And, you know, you could ask the question if you have some rat that was running a maze, and then you
You kill the poor rat, and you look at its brain.
You know, can you tell that that rat was an expert maze runner versus another rat? I believe people have been doing those experiments for a long time. I believe they're an abject failure. I believe there's nothing… no way you can see that.

I think that there are… there are some things where you know that, if… if the environment was sufficiently impoverished.
Certain development doesn't happen in the brain. And, for example, if you have a famous experiment from the 1960s, you have a cat reared in an environment of only vertical stripes. You put it in a cage with only vertical stripes, you put a collar on it so it can't see itself, and so on, then the pieces of the brain that detect horizontal stripes just don't develop.
And I think those kinds of things are the case, but I don't think you can tell. At a final level, this was a finely trained dog versus a generic wild dog that just lived its life without anybody telling it what to do, so to speak.
Alright, I think that's,
That's probably all I have time for today, but thank you for a whole bunch of interesting questions.
And, I'm kind of… this is,
addictive. I'm looking at other questions you all sent in, and there's a lot of interesting questions here. I look forward to addressing them another time.
But, until then, thanks for joining me. Talk to you another time. Bye.