Hello everyone, welcome to another episode of Q&A about future science and technology.

Let me see, we have a whole bunch of questions here.

Let's see, Julie asks, is human oversight becoming more important, not less, as AI improves?

Well, that's one of those questions that's partly about the psychology of what people think they're getting from AI as a tool. I mean, I think that when people's expectation is lots of things will go wrong, I've really got to look at what's happening.

then people will look at it and not use it for things where it would be a problem if something went wrong. At the point where people say, it's just going to get it right.

That's where you run into trouble.

I think the,

the thing I'm seeing, it depends on the task. You know, there are a number of things to say about, about sort of… I would say in the area of coding.

what I'm seeing is… that…

sort of people who really understand the big picture of what's happening, these are the people who can make good use of AI and know how to break down a problem to the point where the pieces can realistically be done with AI. People who just to, like, throw the whole problem at an AI, that's not usually going to go well.

Let me break that down another time, which is to say, when the thing you're trying to get the AI to do is to make something that just has to look right.

then maybe that can work. If it's just make me a form on a website that roughly looks like this, it may very well be able to do that. What's underneath may be a horrible mess.

And may have lots of crazy things that you wouldn't like if you understood them, but the thing you want is just the thing you're seeing, and maybe it can do that well.

But when you're building a bigger sort of tower of functionality, it isn't really good enough to have something where the pieces… the pieces have to have been modularized by you, and then it depends on what you're doing. In some cases, it's kind of essential that you understand what the piece does, and that's a place where people, I think, get into a fair amount of trouble, where it's just like, oh, the AI wrote this piece of code.

It's 300 lines long. I don't understand any of it.

You know, maybe the AI can tell me, you know, you ask the AI, what's going on in this piece of code, and it tells you something, and you still don't understand that. This is a big place where our technology, Wolfram language, is really important, because we built, kind of, the highest level language that's ever been made.

A language which really is trying to deal with computation rather than just pure programming, so to speak.

And it's a place where having the LLM write Wolfram language code, you get something where, if it does a good job, you can actually understand it, and you are expected to understand it, so to speak. It's not something where it's just sort of gobs of low-level code, which no human could be expected to understand. So I think that's a… that happens to be sort of a critical point for our technology, the ability to have

code that you understand. I mean, essentially, my… my,

My… my… sort of one of my statements about our technology, it's what you need if you want to understand what you are computing.

If it's okay that you don't understand what you're computing, but it just sort of looks roughly right at the end, then maybe you don't need us in the middle. But if you want to understand what you're computing, you absolutely need us. And there are many cases where you do need to understand what's being computed. Even cases where people wouldn't have used our technology in the past, but where now they're trying to sort of make a piece of a big system, and they're going to build a lot on top of that piece, and they better understand what that piece is doing, or the whole tower is going to topple over.

So…

But I think the thing that I'm kind of observing is that people who are… people who understand what they're trying to get

are at a big advantage. People who were kind of in the trenches programmers, who were just sort of, you know, knew how to turn the crank to get out more code.

Those people are not in good shape in the time of AI, because it's what's, you know, what becomes important is, did you know what you were trying to get? Do you understand what happened, and kind of how… how well can you, kind of.

how well can you modularize things to the point where you're asking the AI to do something that can plausibly do?

I mean, something I've seen a bunch of recently is the following phenomenon, is people who are like product managers, for example, who maybe wrote code, you know, 20 years ago, but haven't been doing that for a long time, say, oh my gosh.

now we can actually, with our own fingers, write code that implements something. Probably not production code, but something where we can kind of see what the consequences of our ideas are.

Well, it's interesting that that's exactly what we achieved for technical computation 40 years ago when we introduced Mathematica. It went from the point where if you wanted to understand the consequences of kind of a technical idea that you had.

In the past, the only way you could do that with a computer is you find a programmer, you get the programmer to do that thing of programming it, they come back in a few weeks, and they show it to you, and you've probably forgotten what you asked them to do, and it all… the iteration loop is very slow. What… what happened with Mathematica, and now Orphan Language.

actually very quickly, among people like, kind of, the high-end physicists and mathematicians and so on, and then spreading to other people, was that people realized, oh my gosh, I can actually, have the process of computing things be something directly under my control, as sort of the principal who understands what you're trying to do, and that makes the whole thing a vastly more efficient loop.

that's something that has sort of extended now to a bunch of areas that involve, kind of, user interfaces and construction of websites and kind of general, sort of, more general kinds of, I would say, shallower code that's now been extended with AI. That feeling that you can sort of go from an idea to an implemented thing.

I have to say, that was sort of my whole point in building Wolfram Language and its distant predecessor from 1979, SMP, was exactly to have the ability for myself to have that loop of going from sort of idea to implementation as efficiently as possible. And I think

that…

I kind of have felt that I've sort of been living what we would now think of as sort of the AI dream for more than 40 years now, because for me, at least, a lot of the kinds of things I'm trying to do in science, in research, in technology development, I can go very directly from an idea that I have

to a thing that's implemented with the tool of Wolfram language.

And as I say, that's… that's now, sort of, that experience of get, you know, go from the idea to a thing is extended to more people, and people who understand, sort of, what, understand the idea. and then can tell whether the idea got correctly implemented, or got… once they see the implementation, it's very common, you have an idea. Once you see it implemented, you realize, wait a minute, that wasn't quite the right idea, this part of it is great, that part of it doesn't work so well. You know, you want that kind of tight loop, and if you're the person originally coming up with the idea, originally conceptualizing the thing, that's the place where having that tight loop really makes sense.

If your job had been in the trenches, kind of grind the cog, so to speak, well, that's… then you probably… being asked to figure out the idea, you're going to say, that's super difficult. We don't know how to do that.

And, you know, we're not… that's not our job. And, well, so there's then a big advantage if you do understand what, sort of, the big picture is.

I think.

that, the… as I say, in answer to the original question about human oversight.

I think that it really is a question of the expectation. I mean, you know, if you expect that your self-driving car, for example, is really going to do the right thing, you just don't pay much attention and you let the thing do its thing. If you're always like, I know it's going to make a mistake in this or that case.

Like, you know, it's gonna make… this seems to be a bug in current times of these things, that, you know, making,

turns that aren't, you know, non-allowed turns on red lights and things. But, it, you know, and so you know that's a thing which goes wrong, so you pay attention to that. If you are kind of lulled into a sort of sense of security by the fact, oh, it seems to be working perfectly, then that's where you run into real trouble.

I think, that, you know, when I look at, sort of, AI artifacts, the…

You know, one thing to understand about AIs is…

They're always trying to get to the thing that they were asked to get to.

As… as sort of as quickly as possible.

If you imagined, oh, there's this way to do this, I mean, in other words, if they can do what might look to you like a cheat.

that's part of what they're supposed to do. In other words, if they can win the game by just, you know, pressing the stop button at the right time without playing the game, so to speak, then that's what they're going to do. That's sort of how they're set up to work. Now, a little bit of reinforcement learning

with human feedback and so on, probably sort of blunts that complete tendency to cheat. But there's a great sense of, you know, if you can get to… if you're an AI, you're trying to get to a result. You were trained to get to a result. If you can get there, by all means necessary, just do it. So, for example, if you're told to do a mathematical proof.

But you're not told what you should assume in doing the proof, what axioms you should introduce. Well, if you're an AI, the right way to get to the proof is just introduce the axioms you need, and then it's easy to get to the proof.

Now, as a human, you didn't necessarily think, oh, wait a minute, I should tell it, don't introduce new axioms, because that would be a cheat. You don't even think of that cheat. So when you see, oh, the LLM produced this amazing proof, and then you realize, wait a minute, you know, in some cases, I've seen this, it's like, that couldn't possibly be right.

You know, I have a sense of how difficult that proof should be. This is not even close to how long it should be. Let me look at what it actually did.
Oh, it introduced some axiom in the middle. That was something which is kind of tantamount to just writing down the answer.
So, but, you know, that's sort of fair game from an AI that's just been trained to get to the result as efficiently as possible. So that's a thing that, you know, I think the question of exactly what do you tell the AI to do, how do you conceptualize what you're telling the AI to do, do you really sort of tighten it up to the point where it can't go off track?
Some of that stuff is very much, you have to know what you really want, you have to have some experience, and so on, and then you have to have the experience to know, if you're looking at a piece of a proof, or a code, or whatever.
this couldn't possibly be right. You know, you could have the intuition this couldn't possibly be right, in which case, again, your expectations change, you look more carefully, you're trying to, you know, sort of… and then you don't get yourself, kind of, taken off track.
Anyway, a few thoughts about that.
the, let's see…
There's a question here from Kyla, asking, will future calendars schedule our lives and negotiate with each other behind our backs?
You know, there have been products for 5, even 10 years that are sort of appointment-setting apps that I think are supposed to do that.
But I keep on seeing these apps, and then I think, maybe we could use those, maybe my scheduling folk could use those, and then I see the apps disappear again. So my impression is that hasn't worked out very well so far.
Now, you know, will we all have kind of a cloud of AI agents doing things around us? To some extent, we already do have that.
in, kind of, you know, what we see on a web page with ads, maybe if we're buying things, there might be some sort of agent layer between us. You know, this is kind of one of the things that could… could kind of happen in kind of human dynamics of, you know, your agent should talk to my agent. I mean, certainly among sort of busy folk
It tends to be, your assistant should talk to my assistant.
And that's… or your scheduler should talk to my scheduler. I mean, that's the thing that's happening at the human level. It's a thing that, we can expect some of that to automate, I suppose. As I say, my impression is it hasn't been terribly successful so far.
And, you know, a couple of years ago, I was hearing about how, oh, AI is going to be used to do, for example, vendor negotiations and so on, and then I didn't hear any more about that. So, either it became ubiquitous, or it didn't work.
I'm not completely sure which.
the, I know we're not using that, so… I mean, certainly one uses it to summarize documents and things of that kind, but I don't think it's kind of like, have the AI make the next move, you know, you're buying this kind of service, and, you know, it's going to suggest this… this way of, you know, well, we'll… we'll do it, take it for more years if you… if it costs less, and all that kind of thing. I don't think that's,
that's yet kind of an AI-done thing, although my small versions of that
Are certainly things which are done, sort of, all the time.
by agents, and have been for years. I mean, people have had, sort of, agents that buy things on eBay or something and, you know, swoop in in auctions and do those kinds of things. That's

happened for absolutely years and years, and probably that will get more sophisticated, but I don't have the impression that the calendar thing has yet

Sort of taken off.

Uzura is asking, should there be limitations on the use of AI in specific cases from an energy water use perspective?

Gosh, well, I mean, one thing is,

To what extent do you connect AIs to critical systems in the world?

That is, you know, the electric power grid, I don't know, the water distribution system, the, you know, banking systems, all this kind of thing. To what extent could an AI just be in charge?

I think that,

there are places where common sense dictates that it doesn't make sense to have an AI be in charge. You know, if there's a… if there's a switch you flip, and there's $100 million, goes one way or the other based on the switch that's flipped.

By golly, you can have a person, you know, spend, you know, you know, 50,000…

dollars worth of time to figure out which way that switch should be flipped. There's no point in having the AI that's going to cost a dollar flipping a switch that's going to move $100 million this way or that way. So, it's kind of like… like, even if you could, it's… it's probably pointless, and it's… it's probably, you know, sort of common sense dictates that it's a bad idea.

Now, there are other cases where there are things of, sort of.

great import, like, you know, you've got an autonomous weapon, does it… does it attack this thing or that thing? Where, you know, the difficulty is those decisions have to be made very quickly.

And those, you know, at this point, sort of the speed at which such decisions get made is sort of probably ramping up and up and up to the point where it's beyond where humans can do it.

And so then you kind of have to have an AI in the loop, and at that point, you are dealing with, you know, putting… giving sort of significant, you know, actions with… with significant consequences, sort of inevitably in the hands of the AI.

Now, I mean, to some extent, that's already happening, you know, autopilots on planes have existed for a long time, although there's still a pilot there who can… who can sort of, at least if it's a… if it's a plane with people in it, so far.

You know, the self-driving taxi story is a little bit more of a… there isn't really a person

at least, you know, there isn't a person right there to take over. Maybe there's a remote driver who, if they, you know, if the system says something crazy is happening, they can take over. But, you know, maybe that doesn't always work. Certainly, there have been some incidents recently where that didn't work.

Because, you know, all the traffic lights failed at the same time or something, and it overwhelmed the sort of backup drivers, or whatever else. But I think…

So, you know, there are certainly cases where… where, sort of, putting the AI, and you end up, sort of.

by the dynamics of the system, you're putting the AI in charge, and you couldn't really have it any other way.

and has to make significant decisions. And yes, there are clearly cases where, you know, it's, you know, common sense dictates certain things shouldn't be put in the hands of the AI. In other cases, perhaps they should. And I think the thing that,

you know, in…

there's a question of, you know, is it actually in the hands of the AI, or is it the AI that's coming up with a suggestion, and then the humans are like, oh, well, I'll just follow the AI?

And so I think it's, it's, you know, my guess is there's a lot more of the AI suggests humans kind of lazily follow than there is actually connect the AI to the sort of actuation layer to make things happen.

Let's see question here from Fares. How could AI change the discovery of drugs and vaccines?

You know, I keep on talking to people about this, because

You know, in the history of the pharmaceutical industry, in the modern history of the pharmaceutical industry, there have been a sequence of times where people have said, we've got a magic new method for discovering drugs.

Truth is, it's really hard to find drugs that work

And that don't cause bad side effects, and so on.

I think that…

the pharmaceutical industry is definitely quite jaded from having seen probably 5 iterations of, there's a magic new method for discovering drugs.

The thing to understand also about the pharmaceutical industry is that it's a little bit like, sort of, oil prospecting, or gold prospecting, or one of these other kinds of things, where it's like, you know, you make a bet, you drill a hole, it costs a fair amount of money.

You know, in the case of drugs, it can cost a billion dollars to get a drug to the point where it can be approved. It's, you know, you're making a bunch of bets, and, you know, anything that kind of can tip the probabilities more in your favor in terms of the bets you make is probably a thing worth doing.

But people have gotten fairly jaded in that industry. At least the larger companies have gotten fairly jaded. And I think also, you know, it's gone up and down in terms of investors. So, you know, historically, people… well, what were some of the… one of the waves? I mean, back… what was it?

40 years ago, maybe, it was natural products.

go forage the earth, find the bark of the obscure tree in the Amazon that cures some kind of cancer, real story. The, you know, find things like that. Then it was kind of,

I'm not sure I've got the order of these right, but another phase of this was rational drug design.

we will have kind of a CAD system for designing molecules, and we'll know enough about how molecules' shapes work that we can just, you know, design in our computer design system the shape of molecule that we need to bind to this particular thing and take this action as a drug.

Another one was combinatorial chemistry. We combine lots of different things.

Another one that is sort of a generic is high-throughput screening. Just try lots and lots of different compounds. You know, most pharma companies have millions of compounds kind of in a big library, and given that you know what you're trying to get it to do, you just try lots of different compounds and see what works.

Then there was sort of one era of machine learning, where it was kind of like, we'll use machine learning to predict which compounds will work and which will not.

And then, in most recent times, it's kind of used LLMs and be able to go, you know, there was a kind of text-to-drug type thing, where you just say, describe what you want the drug to do, and a molecule will get created by just like a, you know, a piece of text can get created by an LLM, and then that molecule will be your drug.

Okay, what could possibly go wrong?

Well, what goes wrong

is that we human biological systems are really complicated, and there are lots of, kind of, very unpredictable things that happen with random molecule X. Even though it binds to the right site on the thing you were trying to… the target you were trying to get, there can be lots of off-target things that happen. There can be, oh, it turns out it builds up in the liver and is toxic.

Or it has some other problem, and these things are hard to find. And one of the questions is, okay, well, couldn't we just machine learn those other things?

My guess is that that's really hard, because my guess is that there's sort of a… it's a feature of medicine, biomedicine in general. You could say it's sort of footnotes all the way down. You know, one says, oh, you know, this is a general statement that one might make in medicine, but there are always exceptions.

Just like, you know, this drug will do this. Oh, except that in 0.1% of the population, it will turn their hair green. Or, you know, some other, you know, some other sort of weird side effect.

And the thing that I think is true, and I'm getting a better understanding of this as I understand more in my efforts to study the foundations of biology, of this idea of why there are footnotes all the way down. In other words, why it becomes really hard to predict

What will happen with, let's say, a particular drug in a human or the whole population of humans?

And I think the real point is, sort of, we're…

We might think of ourselves a little bit like machines, where, you know, we do all these things, we have all these muscles, and, you know, proteins, and this and that and the other, and we're kind of operating in this way that we might think we could sort of predict like a traditional machine.

The problem is, that's not really the way it works. Machines of the kind that we've been used to the last 100 years or so are machines that were typically designed by a human. We know, kind of, this part is designed to do that, and we can kind of trace it with our minds.

That's not how we were built.

we were built by biological evolution over the last 3 billion or so years, and by lots of trial and error. I mean, we're a trillion generations away from the last universal common ancestor, the, and we're probably, I don't know, a few… we don't know how many trillions of generations we are away from the very first life on Earth.

But, let's say, a future in generations.

And there are, you know, many different that branched out. There are, like, 10 to the 40th organisms that have lived in the history of life on Earth. We are the result of lots of trial and error about, sort of, how to make an organism that more or less works. And the thing that's become clear from my, sort of, efforts and foundations of biology

is that we can really think of that as we're sort of putting together these lumps of irreducible computation, where each piece is sort of arbitrarily difficult to understand. Each piece happens to sort of fit

into the thing that one needs it to do in us as a biological organism, but if you say, what is there inside, it's sort of arbitrarily complicated to describe that. And I think that's the place one's getting kind of… kind of skewered when one tries to say, can one predict what drug is going to work and how, and so on.

I think this whole understanding that we're starting to have about the foundations of biology and about, kind of, how we are put together from these sort of lumps of irreducible computation, it gives one a different sort of perspective on things like, how do you even test a drug?

Because, you know, the typical protocol for testing drugs, you know, double-blind, Clinical trials and so on, involves… you just give the drug to a bunch of people, maybe you give a placebo to other people and so on, you compare the results, but you're sampling the population. And you're expecting that when you've sampled in this or that way, that you would be fairly certain that it will be kind of okay for a decent fraction of people. Now, you know, the dirty truth of the industry is that who you choose to sample with is a critical part of how you get your drug approved.

Because if you can say, we tried this drug in 100 people, and it was great in 98 of those 100 people, and it wasn't too bad in the other two, it's kind of like, yeah, this drug is really a winner. But if… but typically, the way this is sort of gamed, to some extent, is you pick those 100 people, and you might start with, you know, 3,000 people, and say, well, we've got this criteria, and we've got that criteria, and we've got this thing, and that thing, and that thing.

And sometimes there's kind of rules about what you can say. When you finally say the drug is approved. It's like, well, in people of this and that kind, and then any other use is sort of an off-label use.

But still, I think we are actually getting some better idea from looking at this kind of foundational models of biology. We're getting some better idea how to think about, sort of, the statistics problem of how exactly do you have to test the drug to be able to validate what's really going to happen with it. I have to say, my own feeling is that this sort of gaming of drug testing is a little bit, you know, is not the best thing, and, you know, how would one avoid that? Well, you know, part of it is a lot of money is being spent to set up that gaming, and that's part of what costs all the money in doing a drug trial.

And I think, you know, I've kind of wondered, are there better schemes one could use? Like, just say, you know, there are these panels of people who sign up. I'm prepared to be for, you know, you get paid X amount or something, or however it works, you know, I'm prepared to be someone on whom drugs can be tested, and then you're kind of from that large panel, for any given drug trial, you're randomly assigning people from the outside. So, you know, the FDA in the US, or whoever it is, assigns those people, rather than the pharmaceutical company finding exactly the right people to test their drug on.

Anyway, there's this question about, you know, the drug discovery. There are, I think, a few drugs now in testing that were sort of AI-generated in the modern period of LLMs. I keep on hearing people still being very hopeful about that.

And companies getting started and heavily funded, and, you know, maybe there will be a giant cascade of new drugs that get invented. I don't know. I mean, if you look at the history of drug development.

There are certain platforms for drugs that do get developed, and, you know, things like the original genetic engineering efforts in the 1980s, that led to all kinds of interesting things, like, you know, artificial insulin or something like that. You know, there were… that was an important step.

Then there are things that are, well, now in more regenerative medicine rather than drugs, stem cell therapies, some gene editing, you know, CRISPR-based therapies, those kinds of things.

Then there are sort of platforms, like mRNA vaccines and so on, that it seems like a lot of, sort of, the real

You know, dramatic progress is made in these kind of very different ways to make medical interventions and to make sort of drug-like things

Rather than just the straight, let me find a molecule that is kind of an elixir of health for some particular… in some particular situation. I think also, you know, there's lots of progress in complicated things about mixtures of drugs for things like chemotherapy, where it's kind of more a question of can you predict, sort of, the, you know, in us.

In the case of tumors, in the case of bacteria as well, you know, there can be evolution inside a person, so to speak. And so, the question is, you know, if you give a drug that kills off, sort of, the first generations of this, whatever it is that is infecting the person, then

That's all well and good, but if the… if the thing that is attacking you mutates.

then you have to be ready to kind of cut it off at the pass, rather than waiting until it's come… it's in big numbers, and it's much harder to deal with. And so that's, you know, that's another type of sort of platform-type methodology, which tends to be, you know, involve giving multiple drugs at the same time, and so on, and having a particular sequence of drugs that you give, and such like to try and deal with that idea.

of kind of, in a sense, the critters, the antigens, whatever they are, the tumor cells, are kind of migrating around a landscape of possible cells, and the question is, can you sort of cut them off at a pass before they get there, when there are only a few cells that are trying to get through there? And if they did get through, then they'd be in a very fertile place, and they'd be able to replicate a lot.

And then you have lots of them, and then they're much harder to fight.

So, you know, that's another example of that.

I think that's an area… I mean, there are plenty of, sort of, these side areas where

there's a lot of, I think, potential AI use. For example, in stem cells, there's the question of given that you have a sort of a starting point stem cell, how do you get that cell to turn into a heart muscle cell, a pancreatic beta cell, you know, a neuron, these kinds of things. And there are a certain number of very ornate protocols that have been found that involve, you know, you cook the cell for two months with doing this and that different thing every day.

there's a question trying to emulate how it develops in a fetus or in the body, and then there's the question of can you do that kind of thing more efficiently by kind of doing a sort of a machine learning prediction based on what you see? The answer is quite possibly yes.

Same kind of thing when you're doing molecular dynamics, trying to understand, you know, the actual flopping around of molecules. You know, can you speed that up by using machine learning methods? Quite possibly, yes.

These are all things which are not quite The full, you know.

get this molecule, I'm telling you what molecule to use. There are more, kind of, peripheral aspects of things where I think there's real, you know, there's probably more obvious opportunity for, kind of, AI than there is in the pure, find-me-a-drug type thing.

I think, you asked also about vaccines, It's kind of…

Well, it's a reasonable question. I mean, for example, an mRNA vaccine

In a sense, so, okay, basic how vaccines work, you're trying to get our immune system to generate more antibodies.

maybe T cells that are specific antibodies, specifically sort of ready to attack some particular antigen. And the way it works is that every antigen, every sort of invading molecule has a certain shape.

And the immune system has this very clever mechanism to have, where it has

tens of billions of different possible shapes that antibodies can be in, and when antibodies are produced, there's sort of this random process of what shape it will be, and then kind of the idea is that the immune system has a whole lot of mechanisms that amplify any shape that is finding invaders, and that then… then having amplified that shape, you get many, many, many copies of that kind of antibody, or that kind of T cell, and then those things will be ready to kind of attack the invading army of antigens.

So, you know, the way a vaccine works is it stimulates the way ordinary vaccines work, is they kind of stimulate

the production of T cells, B cells, antibodies that come from B cells, and so on. They stimulate it by saying, here's a thing that more or less looks like the critter you were trying to attack, and, you know, get ready to attack this thing.

So, usually what happens is, like, for example, with COVID,

the… the thing one does is, you know, the COVID virus was 29,000 base pairs long, that was its genetic program, and there was a question of, well, let's pick little pieces of that that we can kind of… that we can expect we can make antibodies. Antibodies are… what is it?

I think they're 9 amino acids, is that right? That means they're 20-ish base pairs long, roughly. So you're picking, you know, which

Length 20 sequences, roughly, are you picking out of that 29,000 sequence… 29,000 base pair sequence to be the thing that will be the recognition key that you're training antibodies to go attack?

So, you know, there was… there was a… there's a certain amount of art to figuring out, you know, do you pick a piece of the spike protein? Do you pick a piece of the viral capsid that's the container of the virus? What do you pick?

And so various things were tried. Some of them worked better than others.

I don't know whether that's a reasonable machine-learnable thing. It might be. There are all sorts of weird things that can go wrong if you pick the wrong… wrong piece of the shape space, so to speak, and some of them did go wrong, and some of the branches of the COVID vaccine trials. Then…

You know, there's a lot of complicated stuff that requires a model of the immune system to really understand what's going to happen. And, you know, when you have

probably natural antibodies. Some of them are very specific, some of them are less specific, and the, you know, some of them are sort of, I'm going to attack a little bit this whole range of invaders, and other ones, I'm going to attack really strongly these particular kinds of invaders. And exactly how that plays out when you have a vaccine that's sort of forcing you to say, I'm gonna, you know, take this particular template.

you know, get ready to attack that particular template is not so clear. And how that relates to the evolution of the viruses on the outside, that's also not clear. And there are many, many mysteries still in the way that things like viral infection work.

like, you know, why is there more flu in the winter than the summer, and so on, and that's still not… not really good answers that exist for that. And there are guesses, but I think… I think one of the meta guesses, I would say, is that this question of are you sick or not is not as binary as one might think.

That, you know, people have a certain level of viral load that sort of goes up and down over time, and that it's really a question more of sort of how healthy are you in general, whether those viruses get to replicate, rather than, did you just get a big load of viruses that would make you sick however healthy you are?

I'm not sure that's right, but that's a… that's sort of a plausible theory for what's going on, and it does explain things that were observed, particularly during the COVID pandemic and so on.

But, in terms of

the design of vaccines, you know, in the case of mRNA, it's, well, all vaccines work by just picking a piece of protein and so on and say, well.

all ordinary vaccines, it's just you pick a piece of protein, this is the thing that you should be generating antibodies to generate those antibodies. mRNA is a little bit more insidious, because it's just the program for making the proteins, so the mRNA goes into your cells, and it makes the proteins sort of on-site.

That you're going to attack.

And, you know, the potential problem with that is it goes on making those proteins, and it goes on producing things that you go on saying, I better attack that. And that, you know, activates your immune system, causes all kinds of problems. It's a difficult thing, and how quickly the mRNAs sort of disappear, and whether that happens in everybody, and so on.

this is a big can of worms, that, and, you know, what one can do, like, this whole question about selecting what part of the invader one should make vaccines to, I haven't actually heard of people using, AI methods. I mean, there are a whole lot of drug discovery companies. I haven't heard of vaccine discovery companies. I've heard of companies with, that are trying to do, you know, discovering pesticides.

discovering various kinds of materials. I haven't… maybe I'm not remembering properly, but I haven't heard of that particular one, although I'm sure it's being done.

Let's see…

Well, okay, I'm, I'm,

lots of questions here, I'm just going to sample some. So DES is asking, not necessarily bespoke drugs, but do you see DNA being used to suggest more effective treatments being automated and common soon for various ailments? In other words, the question is, if you know your genetic sequence.

can you better know, oh, I should take this drug and not that one?

there's been some work done along those lines. I mean, I know, you know, I, being sort of, perhaps often ahead of the curve in these kinds of things, got my whole genome sequenced back in 2010, when there were only a few tens of people who'd been sequenced, and, you know, I've had my genome

sitting on my computer ever since, and, you know, whenever some random study comes out about something, you know, I can go look up in my genome, so do I have that or not?

And there are a small number of cases where one knows that there's sort of a connection between some drug working and a particular genetic profile for a person.

I mean, I would say that people had sort of imagined we get the program for humans when we get their DNA, the 6 billion base pairs and so on. It's like, then medicine is solved, people thought.

they didn't count on computational irreducibility. They didn't count on the phenomenon that just because you know the program for something doesn't mean you know how the thing will behave.

This is a sort of big thing of mine, is that I found first in the 1980s that, it was…

that even when you know the rules by which a system operates, it can be that there's no faster way to find out what the system does than to sort of follow each step and see what it does. You can't sort of go and jump ahead and just say, I've got a formula for the answer.

But people kind of imagined that following the pattern of sort of traditional physics and things like that, that once they had the program, once they had the underlying rules

for sort of making a human, then you would be able to say everything. That did not pan out.

And the number of things that you can kind of say, oh yes, given the genome, I know what's going to happen, is quite small.
And, the,
you know, one of the things that seems to be very ubiquitous in biology is that there are many pathways for things. It's, you know, there are many different ways to achieve a particular thing. So, you know, you can have something where it looks like, well, an example for myself is, according to my genome.
I should be lactose intolerant, as many Northern Europeans are. Well, actually, in practice, I'm not.
So how could that possibly be? Because there's a particular gene for lactase that my particular gene is kind of munged in some way, but yet I, you know, digest lactose just fine, thank you. And, the, you know, why… how does that work? Well, presumably, there's another pathway that I have, I've never bothered to figure out what it is, for digesting that, that doesn't happen to be the one that everybody sort of knows is the common one for doing this. And that seems to be a very ubiquitous phenomenon in biology. In fact, I've tried to study it a bit in thinking about the foundations of biology, this kind of degeneracy of action, that there are multiple pathways to achieve a particular objective. And that, I think, makes it,
That makes it difficult to say, oh, you know, this gene in you is munged
Therefore, you will have this problem.
well, you would have that problem if that was the only way that that particular thing was going to be achieved, you know? If that was the only way to repair your DNA, then absent that gene, you know, none of your DNA could be repaired, and we replicate DNA, what is it, 10 to the 30th times in our lifetime, something like that, in all the different cells in our body.
So if your DNA repair, which is needed to make sure that you don't get, sort of, mutations in DNA, if that fundamentally didn't work, you're out of luck. You can't. And in fact, if you have enough
sort of genetic changes that really no DNA repair works, then, you know, you don't get to survive at all.
Even in early fetal developments and so on. So…
It's… but I think it's very confusing, because, you know, the slam-dunk answer of, given this thing that will happen, doesn't seem to be that way very often at all.
And even things like eye color, which seemed to be something that was governed by, you know, sort of a blue-eyed parent and a brown-eyed parent will have brown-eyed kids.
well, my wife has blue eyes, I have brown eyes, and I have four kids, and let's think. I think there's a… it's sort of almost Mendelian, but not quite, because at least one of them has greenish-colored eyes, and, you know, and so on. It's, I think the, the thing that,
is… and if you look up, kind of, what's known about the genetics of eye color now, it's a very complicated story. What used to be, oh, there's just one gene that controls it, is now there's a whole array of genes, and they lead to this, and that and the other, and so on.
So, it's been a bit challenging to go from, sort of, genomics to what does that mean at a medical level?
It's, it's more relevant for things like tumors, where the… where the genetics has really changed relative to your base genetics.
But to the question of which drugs are going to work on you versus somebody else, there's a small amount known. I, you know, will it become routine clinical practice to kind of sequence

the DNA to find out, even at a simple level, you know, should you take, you know, aspirin, or Tylenol, or whatever else, or acetaminophen, or whatever it is?

the, You know, if the cost of DNA sequencing comes down enough.

And I know there are, there are efforts, there's even a company I've been somewhat involved in, that, has sort of a new methodology for doing DNA sequencing. If it works out.

DNA sequencing will be incredibly cheap and incredibly fast. And if that happens, then it might become really common to just say, well, you know, first of all, for… not just for everybody to know their… their,

Oh my gosh, what's it called? Germline.

their basic genetic sequence, but also to be able to sequence things in their immune system and all kinds of other things. But for everybody to routinely know their full germline sequence, I would think that will become… I mean, I'm surprised it hasn't happened yet, actually, because it's not that expensive to get sequences. I think the main reason it hasn't happened is because so what?

you know, for somebody like me, it's amusing to be able to go and say, oh, there's this thing, I can just go look in a genome viewer and… or from language version of looking at biomolecules, and just go say, you know, do I have this mutation or not? You know, does this gene of mine work in the way that it normally works, and so on. But I think that's a minority interest.

Let's see…

Let's see, a question from…

Question from Des here. Do you see delivery drones as a viable solution for normal people in the foreseeable future, or is it like helicopter taxis that don't and won't become commonplace for various reasons?

I don't know. I mean, I've seen a lot of demos of delivery drones.

And yet, I don't think I've seen a delivery drone in the flesh. I mean, one routinely sees these little, you know, little,

little cart… little… what are they… what would one call them? These little, creatures with wheels that run around, particularly on things like college campuses, delivering, you know, food and things like that.

That's, you know, in a sense, that's a terrestrial drone, not the more common kind of aerial drone. So terrestrial drones are, you know, have arrived.

At least in fairly controlled situations. They're routinely used for delivery. In terms of aerial drones.

they're used for, you know, delivery to obscure places and things like that, but one doesn't seem routinely used. And I know, you know, I've seen tests of where, you know, you pick on a map, you know, where do you think your front

porches and so on. Mark it on a map. Okay, that's where you marked, the drone will land there. There was another plan where you would… you would lay out this kind of target pattern, like a helicopter landing pad, you know, in your backyard or whatever, and that's where the drones would come. I mean, it's kind of an amusing sort of social dynamic. I mean, people have, you know, different kinds of things that people sort of added

to their houses or whatever over time. I'm sure there was a time, I don't know when it was, actually, when people, like, routinely added mailboxes at the end of their driveways, and, or wherever, and, you know, the idea that everybody would add a drone landing pad, you know, on their front lawn or something like this, or in their back, backyard, or whatever it is, or on the roof of their building or something.

It's an amusing sort of social dynamic that that might happen. I doubt that will happen.

Because I think that it is clear that, sort of, by the time, sort of, the self-driving technology and image recognition technology is good enough, you're just going to have to say, here it is on a map.

Okay, if it's snowing that day, you know, snow is the bane of, sort of self-driving cars and things like this, because it changes the way things look.

And, you know, it could be that the delivery drones will have to have a day off if it just snowed, and I'm not sure, because then it's like, well, you know, what we thought were the markers for where you want the drone to land relative to your front steps or whatever, those markers are all covered in snow, so the drone has no idea where to go.

Just like a human trying to, you know, walk up the path leading to your front door or something might have a hard time knowing where that path is if it's covered in snow.

But… You know, I think… It's always a question of, yes, it would be kind of…

You know, it's sort of an economic question of when… when it becomes the case that, sort of. flying new goods really makes sense, because these drones aren't going to be hugely powerful, at least not at the beginning, and the battery life will be low, and things like this, and it's kind of a, you know, they're only delivering small, I think, like, 5-pound-ish

types of payloads, I think maybe even less than that. So, you know, I think that's… that's another limitation of what can you get delivered that way that has actually, sort of, that's worth doing that with. You know, you mentioned, helicopter taxis.

You know, the thing about helicopters is… You know, helicopters existed

from… when was it? 1950 or so, earlier than that, even. Why did helicopters not catch on big time?

I think the number one reason is because, well, A, they're a bit hard to fly, and, you know, autopilots hadn't been developed as much when helicopters were kind of a popular thing, and number two, it turned out that they needed a lot of maintenance, and that that was pretty expensive.

And that…

they kind of have more moving parts than planes, et cetera, et cetera, et cetera. I mean, there were various companies, I remember in the 80s, there were some companies, like a notable one was Digital Equipment Corporation, which had a lot of facilities dotted around New England, and their… their sort of… one of their gimmicks was, you know, rather than everything is in Maynard, Massachusetts, or something, it's like there was a Littleton facility, and there was this facility, and that facility.

And the way that they kind of had people go to a meeting in a different facility was hop on a helicopter, and they had a whole fleet of helicopters.

And, you know, so that was a thing that was happening, and it looked like, you know, that might really catch on. It didn't, because helicopters are just kind of too expensive, and sort of too complicated, and so on.

I don't know… I mean, when it comes to things like quadcopter or multi… multi-rotor drones. you know, I keep on seeing demos, I just saw one a couple of months ago, of a single person, kind of fully autopilot drone with, I don't know what, you know, 16 rotors or something, that was, like, I can just take you to places as, at least as a… as one person at a time.

probably at some point, that will all catch on.

I think it is a very detailed question of, kind of, how long does it take to charge it? How expensive are the batteries? You know, what are really the use cases? You know, is it useful to have a one-person air taxi?
You know, where could it really land?
You know, what will be the rules about whether it can land on the top of a building, things like this. You know, what will happen if there's wind? Is it, you know, is it a single puff of wind and the thing can't fly? How does that work? You know, how safe will people feel flying around in these things? And how safe will it actually be to fly around in these things? So, those are…
So, it's hard to predict. I mean, the technology basically exists at this point. Will the price curve be such that it can really be deployed? Will
the use cases really be there. I think the use cases would be there. I think that people, you know, commute to work by air is a thing. I mean, you know, a very modest number of people commute to work by helicopter in Manhattan, for example. You know, at times in the past, that's been… seemed more popular. I mean, I… you know, I more hear people just…
You know, the people who have as much money as they could,
you know, they're more likely to, you know, like, I live in New Jersey and commute through the tunnel every day, and somebody drives me, and I, you know, hang out and use my computer for 15 minutes or something, rather than I, you know, hop in my helicopter and land on top of my building type thing. It feels like
One of the dynamics is sort of an interesting thing in terms of technology prediction.
you know, for somebody like me, for example, before cell phones, you know, I first had cell phones in the late 1980s, before cell phones, driving from one place to another was like
this is a total waste of time. There's nothing I can do. You know, if I'm not, you know, there's… I'm just… it's pure downtime.
Then, along came cell phones, and it's like, if I'm driving from here to there, and I'm driving myself, which I personally much prefer to do, then it's… I'm just on a phone call, you know, or whatever, talk… and it really isn't…
downtime. It's just a question of what can I schedule at that time versus when I'm in front of a computer or whatever else.
And now, there's a yet new dynamic, which is self-driving cars.
Where, you know, one's not at the point where one can't pay any attention, but nevertheless, I find that it's a vastly more relaxing experience, where I can kind of feel like, you know, even if I'm driving through some nasty set of traffic jams, I'm like, I'm not stressed at the end of it. maybe if I drove through traffic jams every day, I would just get used to it, and it wouldn't, kind of, you know, increase my stress level, but I think, given that I don't do that every day, it's kind of like, if I've been driving through this really difficult traffic, I'm… I'm going to be kind of fried at the end of it, whereas now that one has, you know, possibility of self-driving cars, that to me, it's kind of like, okay, I really don't sort of feel frazzled at the end of it, because I really wasn't paying more than very reflexive attention to what was going on.
Let's see…
Megan is asking.
What is a form of automation that seems inevitable in the future, much like self-driving cars today?
You know, I have to say.
The first time I had a serious discussion about self-driving cars was 1980.

And I was talking to an AI pioneer called Marvin Minsky, who was given to Wild Ideas, but who was like, of course there'll be self-driving cars. It's, you know, at the time, 1980, that seemed very far away.
And it just seemed like there were so many judgment calls and driving, and so much one had to figure out. Even quite recently. You know, I remember people saying, it just won't be possible to figure out how to merge into another lane of traffic.
Because you always have to kind of, you know, take some risk, and so on, and, you know, the other driver might do something crazy, and you couldn't get out of it, and so on. But in fact, self-driving cars have arrived, and there are just a zillion cases that they have had to learn to deal with, and it seems that it worked.
And I think that sort of opens up this idea that there could be a million cases, not a million, but a very large number of different sort of situations that one would have to deal with in the world, that sort of opens up a bunch of other possibilities for where one could make use of automation.
And I think the one that
Is, sort of, the obvious one is… is robotics.
And…
you know, I had thought for a long time that some kind of modular robot made of little tiny reconfigurable pieces would be the thing that would really be the… the thing that would catch on. Just like there are general-purpose computers, I imagine there will be general-purpose robots, general purpose in the sense that they could sort of be in a shape they want to be.
I think it's becoming pretty clear that humanoid robots are the winning, sort of, form factor, and in a sense, that's unsurprising, because we built our world for the human form factor. You know, the natural world, obviously, was not built for that, but the natural world that we choose to live in We pick the places where that form factor works okay. You know, we're not… you don't have to climb some sheer rock face to get to your front door, and so on. Most people don't, at least. And And then we built, sort of, our whole built world of, you know, of buildings and houses and all that kind of thing is all set up for the kind of the human form factor. The machines we use have, you know, levers and so on that humans can move.
And, you know, and so it's all sort of set up for humanoids, and so that's my guess, is that's what we will see happen, and I've seen many, many sort of demos of this with different levels of convincing this. I think as batteries get better, and you can have your robot, you know, have enough power to wander around for a while, as actuators get better, and we can get close to the you know, 250 kinds of, sort of, joints and actuations that we humans have, you know, probably will exceed that in robots for not too long. They'll be more agile, they'll have hands with a lot more, kind of, articulation points than we do. I mean, it's like whales have flippers with lots of articulation points. You know, I think the, we can fully expect to see robot hands with more articulation points.
And then it becomes an issue of how do you control that? That's a separate problem.
It's actually a little company I've been involved with that's… that's, solving that particular problem. But in any case, the… the kind of…
You know, the thing that seems like it might be a very daunting thing is, okay, how do you deal with… if I have a robot in my house, how do I deal with all the weird things that could happen? I mean, I don't have a pet iguana, but I might.
And, you know, what happens if the pecca iguana has, you know, put itself in some place where the robot is trying to open a door and it would squish the iguana if it did that? You know, do you have, you know, does the robot know about all these different weird cases?

It's, I have to say that seeing what's happened with self-driving cars, I fully expect that that is something that will happen. Now, with cars, it's sort of easier to get the training data, maybe.

I mean, people have sort of whole rigs where they have humans with, you know, who are all hooked up, and they're doing tasks, and they're recording all the things the humans do, so the AIs can… robotic AIs can learn from that, and there's a lot of effort

In figuring out how do you use the fact that there's sort of… there's physics in the way that robot arms move. How do you use that to kind of reduce the amount of training that you have to give to get the robot to learn the things it needs to learn to take the actions it needs to take?

But, you know, I would expect that sort of humanoid general-purpose robotics will arrive at some point from, I don't know, a year from now to 5 years from now, something like that. Hard to predict the exact time, and when it does.

It'll probably be quite dramatic, because it's, you know, there are probably a fair number of tasks, including very everyday ones, that, you know, many of us do, that will very quickly be automatable.

I mean, unlike something like ChatGPT, where it was just, you know, a website, everybody can go to it.

You know, the humanoid robots will cost $50,000, or whatever it is, so it's kind of the cost of a car type thing, and it's, you know, it will not be a thing where everybody just gets one on day one, and probably the ones that exist on day one, crazy people like me will get them, but they won't work very well, and there'll be all kinds of problems, and really, it would be better to wait for a while to get

To get it.

I think…

you know, this… many cases, dealing with all those cases, you know, there are many other kinds of mechanical automation, like in surgery and things like this, that one can imagine. Many more factory automation tasks. I mean, factories tend to be

They are somewhat robotic, somewhat reprogrammable, but often, what's the point? You know, you've got a production line, it's going to make a million of some kind of thing, you might as well just set the robot, in quotes, up to be in a fixed place, taking fixed actions, because, you know, every bolt is the same. By the time it's truly, sort of, programmable robotics, it makes it conceivable that you can make a wider range of products.

You can do that so long as all of the validation of those products can be done across this whole range of parameters. I mean, it's no good to say, yes, in principle, we could make this bolt a different shape, but then the bolt doesn't fit into the, you know, with the nut, and so on, and that doesn't do you any good. So it's kind of like, can you… is it a parameterized design for some whole system.

If yes, and different people want different parts of the parameter space, maybe there are things designed for you, and each of us are a somewhat different size and shape.

And so, you know, it's okay, you want that very custom-made shoe that's exactly the shape of your foot or something.

Or you want, you know, the glove that fits perfectly type thing, or whatever else. I mean, it's become routine for things like teeth.

that you can, like, print a tooth to the exact sort of shape you need, but of manufactured things, that isn't yet a thing, at least in… it is a thing in some areas, but not in most areas. In, you know, by the time robotics becomes cheap enough, it becomes possible to sort of make to order lots of different kinds of things, and I think one can expect to see that product lines will get much wider, and people will have much more choice in what they get, for better or worse. I mean, sometimes, you know, you give people more choice, they don't know what to choose, and they don't buy the product. But, you know, it's like, I would like

my car in lots of different colors, and then, you know, even… you can only do that if you're not having… not keeping inventory of these things, and if you're able to just say, hey, you know, press the button, it's gonna get made in the next 3 days, when it's finished, it'll drive to you, type thing, or whatever else it is. But I think, so, you know, I would expect factory automation

And again, as programmable automation becomes sort of cheap.

lots of things become possible that weren't possible before. And I think, that… I'm not sure… I haven't really thought through what are those things? I mean, I suppose one extreme one that I mentioned before is, nowadays, you know, you build a house. It's expensive.

It's… and you certainly, but a lot of the cost, you know, there's materials, but there's also a lot of labor cost. If that labor cost was all robotic, and was cheap, one could imagine doing crazy things, like you build a house, and then you decide, I want to move my house.

well, you just take it down again. You have the robot construction crew just take the house apart again. Doesn't quite work with a lot of the ways construction is done today, but one could imagine that. And then it's like, okay, if labor is free for doing something like this, then you just, you know, put pieces of the house on a truck, move it somewhere else, rebuild it.

It's,

I think, you know, that's the kind of thing that can happen if the cost of that kind of labor goes down low enough. Now, it's worth saying.

You know, that…

One of the things that happens is when the cost of things goes low enough, lots of new things become possible.

And so, you know, what were jobs that people had doing this or that thing that get automated? The lesson of history has been, when things get automated, new things… new possibilities open up. And in the end, there are things for people to do.

Sometimes there are things that more play to human strengths.

rather than things which are kind of monotonous, and sort of anybody could do it, and it doesn't really require sort of any human spark, it's just one's making use of the fact you have an arm and a hand, and they can move this to that… this place and that place. And, you know, one might hope in a sort of, from, I don't know, perhaps it's an overly utopian point of view, that more people can get to do more for

fulfilling things by virtue of the fact that more things are getting automated. But we'll see how that plays out.

spare Parts is commenting, should we be training humanoid robots for moon gravity as well? That's a good point. I wonder if anybody's… I'm sure somebody's thought of that. I haven't specifically heard of that issue.

That's,

one… one would think the answer is definitely yes. It's a lot easier to have robots on the moon than… than people, although it might be more fun to have people on the moon, so to speak. And, you know, it's… it's more of a,

From the point of view of, kind of, aspiration, it's more interesting to have people there. From the point of view of actually getting things done, it might be better to have robots there.

twin comments, it would be amazing to have robots exploring the deep sea. I mean, to some extent, that already happens. One of the issues about the deep ocean is, you don't get to communicate by radio with something deep underground, deep underwater. You know, you have to have a wire, because radio signals don't go through water. Only very low frequency radio signals go through water. Very low frequency radio signals can only transmit

At, you know, 1 bit per second or something.

And in order to even pick up those very low-frequency radio signals, you need a hugely long antenna. I mean, one of the weird technological things in the world is submarines that prowl around for months in the deep ocean, kind of undetectable to anybody, at least one assumes.

from the surface, those submarines, you know, if the… if horror strikes and the submarines are being told to launch their missiles or whatever, it's like, how do you even tell them to do that? They're in the deep ocean, nobody, you know, and they're not communicating with anybody for months at a time. What do you do? Well, the actuality for U.S. submarines, I think, is, you know, they're trailing these long antennas behind them.

very low-frequency radio antennas, and there is a big antenna, I think it's in Wisconsin, that's a big long thing that's miles long that transmits signals that can be picked up anywhere on the Earth.

by the submarines.

And obviously, they're encrypted, so… so you can't kind of fake it and send the wrong thing and all that kind of stuff. But that's sort of the only way to communicate, other than having a wire connecting something on the surface, some ship on the surface, with this submersible that's going into the deep ocean, and that's what's had to be done. And there have been robotic

Deep ocean,

devices. I suppose the suggestion being made here is, could you have a truly autonomous one?

And I guess there are now autonomous, little autonomous boats

And, that prowl around the oceans for years at a time, kind of monitoring various kinds of things.

I've certainly seen those in prototype form, I don't know if they're deployed.

to have a deep ocean kind of surveillance, thing that just goes prowling around the deep ocean.

Yeah, that seems like a thing one could do.

You know, it's like, why would one do that?

Well, the usual reason is because you're… you're trying to find something, like you're prospecting for minerals, for example. I can imagine that working, or you're sort of going around checking the undersea cables and so on, and I know that's already done with submersibles. I don't think it's done with autonomous ones, but those are kinds of use cases. They don't seem like dramatically big use cases, but they're certainly ones that, in their domain, are important.

I don't know, you know, there's a question of… of…

It's, again, when you have something

in, you know, up in, you know, in orbit, or in the upper atmosphere, or even on the surface, you can use solar power, for example, as a sort of portable source of power. In the deep ocean, I don't know what you can use.

I think, you know, it's dark down there, and I don't think, you know, you can't use temperature gradients and things like that, because you need too big a sort of… too big a thing to be able to use that, so I think the only thing you can do is store power in a battery, which means you have to bring the thing up to a surface quite frequently, and,

You know, that's… that's sort of a limitation.

Okay, Technic is commenting. It's a reasonable question. Why don't undersea devices use sound like whales and dolphins? No need to use radio when a sound will travel, halfway around the world.

That is a… that's a good point, and…

I don't know the answer to that. I mean, I would imagine that,

I guess a bunch of the use cases so far have been for applications where you don't want anybody else

To be able to tell what's going on.

Though that's not a good argument. I… I'm…

I… there is another… there's another issue.

The issue is, the reason that sound travels so well in the ocean is that it gets into these waveguides, effectively, layers of the ocean in which the sound just sort of bounces back and forth.

For thousands of miles, sometimes.

If your objective is to communicate with something that's 5 miles down underneath you, I'm not sure that the sound travels so well. I think it's a specific thing that you can travel, you know. whales can be having discussions across the Pacific Ocean because the sound goes in these kind of, these,

These kind of waveguides, where it's sort of like having a pipe for the sound at a particular depth.

So I think the issue is that, sort of, just going deeper and deeper into the ocean, I mean, the… I have to think about that. The speed of sound increases as the pressure goes up. Sort of water becomes more… more like a solid, so to speak.

as it's under higher pressure, sound goes faster. I don't know whether…

Let me think about that for a second. See, what's happening in the… in these waveguides is that the sound is bouncing off.

the top and bouncing off the bottom. And it could be that if you try and… well, no, I think it should work. The refractive index should be such that you could send something straight down. Maybe you can't send something at much of an angle, because it just bounces off deeper layers of the ocean. I'm not sure, but I think that's the… that's the initial issue, and I haven't heard of it being done, so I kind of suspect there is an issue like that.

Let's see… okay, maybe one more question, and then I should wrap up for today.

Okay, there's two fun ones here. One from Kiki.

Could AI design a snack optimized to your genome and your mood?

That's a fun question.

Well, I suppose from wearables data, and maybe some slightly invasive data, like, blood glucose,

That's an interesting question. I mean, particularly from blood glucose levels, you can see, you know, you eat a snack, you have a peak, goes down again, it's going down, it's going down, it's crashing, it's… the glucose level is crashing, it's like, oh, let's pull it up again, let's not have it bounce around too much. Let's, you know, if you were paying attention to if you eat this, then You, you know, the…

the glucose level spikes that way. As a… as a absurd data collector, I, of course, have a little… little continuous glucose monitor somewhere… where is it? Somewhere… somewhere on my arm. The, and,

the, I did notice Discovery for me is that white rice spikes my glucose level like crazy.

I didn't know that. You know, each of us is different, and it's different in different circumstances, but

you know, it's kind of, like, I'm not sure… I'm not sure it matters to spike one's glucose level. Let's be clear about that. You know, some people think it's a terrible thing. If you have diabetes,

there are issues with it, but if you're, you know, if you have a, you know, a nice functioning pancreas that's secreting insulin, there's a certain argument that you should give it a bit of a workout from time to time, and give it those spikes, and have it do some work, otherwise it'll get lazy.
and maybe it's, you know, the cells will kind of, you know, give up, and so on. So it's not completely clear, at least to me, what one should do about that. But certainly, the idea that one could sort of be monitoring that and saying, you know, try eating this right now, it's an amusing idea. I hadn't,
hadn't occurred to me. You know, combining that with one's genome, again, the connections between the genome and what you know about what foods taste like and so on, there are a few cases where it's known, but it's pretty tenuous. You know, it's always amusing to me
when my kids
you know, will order the same food as at a restaurant as I do, and it isn't because they're copying me. It's because, you know, we've got some genetic similarity, and we both like this particular kind of thing. And so there clearly are genetic correlates of what food one likes, but I don't think that's at all decoded at this point.
You know, one can imagine similarly for things like scents, that there are, you know, genetic correlates. That might be a more direct thing. You might be able to actually see, you know, that there are these, you know, of the few hundred olfactory receptors that we have in our noses and so on, you might be able to see, possibly, this person has lots of that kind, or this person doesn't even have receptors for this kind of scent, so don't bother giving them that.
Alright, I'm led to maybe one more question from Henry.
Could future cities control their own microclimates?
Well, that's an interesting question. I think, you know, that relates to things like cloud seeding and so on, and, you know, cloud seeding is still a bit of a magic story of whether, you know, it's quite possibly the case that every raindrop and every snowflake
has a little grain of dust at the center, or something like that, a thing that nucleated it that allowed it to grow. And so, you know, as you spray more of that stuff
into, you know, into the air in places where there's lots of humidity, let's say, potentially you can have it rain. And for… for 100 years, people have been doing experiments on that. There are… there's sort of increasing evidence that one can make it work by even methods that were already known
you know, close to 100 years ago, of, you know, you spray, what is it, silver iodide, I think, and some other things that… that are, you know, seed clouds. Make them… make them actually rain. So, you know, those are things that one could do.
for, A,
sort of to control… to control something of a microclimate. I'm not sure what else one gets to do. I mean, you can have sort of aerosols, but I don't think you can put an aerosol in the atmosphere and expect it to stay where you are. It's not going to hang out right above your city, it's going to get dispersed if it's in the upper atmosphere, dispersed very widely, so I'm not quite sure how that works. Okay, I have to tell one last story here.
So back, 40 years ago, I was, a professor, and, I was,
Professor in physics, Math, Computer Science.
And, was like, I was getting assigned a course to teach, and I picked Teach the Physics for Non-Scientists course.
So, I started off doing this.

And,
I, I had very, kind of, idealistic, goals of, you know, I knew there would be a bunch of, of non-physics majors taking this class, and my, you know, non-science majors even, and my, sort of, goal was convince some of them that, you know, that science is really fun, and so on.
Well,
I would say that, I… it was a slightly funny situation, because there I was as a, you know, fine, full professor and so on, but I'd actually never taken an undergraduate class, never even been in an undergraduate classroom in the US.
And so, you know, at least in those days, there wasn't, you know, it wasn't like there was sort of training for professors about what do you do, because people assume by the time you're a professor, you've kind of learned what you do. Well, I hadn't.
And so it was a little bit of a funny situation, because my expectations about, kind of, what your average college student, would be able to do, and would know, and would think they should be able to do, and so on, and what their homework should be like, and so on, were distinctly out of whack.
Let's say. I will say that I think, the, the, the, I think the… the lectures I gave, I think, were pretty decent, but the, and, and some sort of locals started showing up and just hanging out at the… at these lectures, although I think the students
felt that, you know, they were much more expecting to be spoon-fed. This is what will be kind of in the homework type thing. And, and there were things that happened. This was 1980…
Early 1988.
And this was at University of Illinois.
And, the, One of the… one of the things that was just… I just didn't…
didn't think properly about it, was, I… one talk was about how airplanes work.
And I sort of assumed that most kind of college-age kids
Would have been on a plane at some point.
And eventually, I realized from questions and things that people really had no idea. And so then I eventually asked, you know, how many people here have been on a plane? The answer was about half the class only.
So that was something I just was completely blind to, didn't see coming. I'm sure that will be different today, but, that was, you know, one of many things that went wrong. Anyway, why am I telling this story? I'm telling this story because one of the things that was, I had these kind of, sort of, everyday physics-type
type questions that I came up with that were on homeworks, maybe. And, I think my… the very first question, the very first homework was, estimate the number of ears of corn in a typical cornfield.
Which I thought was thematically relevant, given where University of Illinois is, and so on. And that was, like, a total disaster, because people are like, how are we supposed to calculate that? You know, how are we supposed to work that out? Well, use some common sense. You know, roughly how big is a cornfield? You know how to multiply, you know, how far apart are corn plants? You can go look if you don't know.
And… but the idea that one should do that kind of thinking was… was, I… I discovered, deeply alien to the American college student. It was much more, you know, you're told this… this is the formula for what you should do. In any case, my mistake.

But in any case, one of the questions that I remember in something that I… I forget what it was, homework or… there wasn't… there weren't tests, but something final thing, or something like that, was, describe some of the issues associated with building a dome over a city.
Which obviously would be relevant to controlling its climate, but that was kind of the, and, that, like, like many other questions I asked, was a total disaster.
And I… the only thing I can say… I can say a couple of things, that I did office hours for that class, and almost nobody would ever show up. But occasionally, students would show up, and one of the things I learned is… is how little information college students, at least at that time, were sort of getting about the world.
And, you know, students would show up, and I would sort of say, so, you know, what's your story? What are you studying? What do you want to do when you're grown up, more or less?
And I'd end up with these interesting conversations, and I thought, you know, okay, it's, you know, I've tried to… I know a little bit about the world, I know more today, but I knew a little bit back then, and, you know, I can be slightly helpful in suggesting this makes sense, this doesn't make sense, etc.
It's been very…
very gratifying that decades later, a whole bunch of people, you know, I don't know how many, maybe, you know, a whole sequence independently of different people who happened to show up to, sort of, office hours, for that class.
and talking about things having nothing to do with physics for non-scientists, have contacted me and said, hey, thanks, you know, you gave me a good piece of advice that I took, and it worked out great.
So that was very gratifying. I'm not sure that wasn't sort of the supposed point of the class, but that was a good thing. I have to say, okay, my last story is at the end of that… when that class was sort of winding down, I decided to have kind of a party for the students from that class.
at this research center that I had.
that, have been the reason I'd gone to the University of Illinois to start this research center there.
The,
And it was right at… just before version 1 of Mathematica came out, and they were already starting to be media articles about it and so on, and there was a kind of a funny article with a title, like, Physics Whiz Goes Into Biz, which had just come out, literally the day that I was doing this… this party.
And so the party was sort of a disaster, with a… nobody was saying anything, and I was just kind of going around saying, you know, what are you gonna, you know, what are you gonna do this summer type thing, and people were answering with one sentence, or whatever.
And, the, somebody who was working with me at the time was, like, noticing that this was kind of a disaster, and, thought, you know, let me just copy it, make a bunch of copies of this article that just came out, and sort of hand it out to these students. So, you know, hand it out to the students.
They sit, it was a kind of a long article, so they sit there for a long time reading this article.
And eventually, there's this kind of,
kind of quintessential fraternity guy with, you know, the backward baseball cap and so on, who, kind of looks up from reading the article and says, hey, with all this stuff going on, why the heck are you teaching college?
So, I had the opportunity to say, well, you guys are probably the very last cohort of people that will get the benefit, such as it might be, of my efforts to teach college.

And, of course, so… so it has been. But it was kind of a… a nice moment for me, and I thought it was, everybody thought it was, thought it was kind of, kind of amusing, and, you know, perhaps, not a good match. I… I… I will say that I…
I… I… I might do a better… I don't know, I think… I think… I think if you looked at the, the actual lectures I gave, they probably were pretty decent, although my expectations about the level of thinking expected were… were… were not right.
Anyway, I see somebody saying, thank you for the office hours. That's appreciated.
That's, I always enjoy
chatting with you all, and I enjoy the questions. They get me to think about lots of kinds of things. So thank you very much for that, and look forward to chatting with you another time.
Bye for now.