

Hello everyone, welcome to another episode of Q&A about Future of Science and Technology. And I see a bunch of questions here.

Boy, you asked me very philosophical questions.

There's a question here. Could there be entire branches of science that humans can't discover until we invent new kinds of senses?

I think the answer to that is, sort of.

Almost certainly, yes. I mean, if we look at history, the invention of the telescope led to all kinds of science, not just astronomy, but also mechanics, understanding that there were moons of Jupiter that kind of worked kind of like things... that there were sort of laws of mechanics that would work there as on Earth, so to speak.

microscope led to, sort of, the discovery of biological cells and many, many other kinds of things in biology. I think that the technological extension of our senses has been simply critical to the development of science.

I think electronic amplifiers are probably necessary for the discovery of quantum mechanics, things like that.

So, I think the answer is definitely yes. Now, what kinds of new senses might we have? I mean, so one type of thing is things that...

sense of structure on a molecular scale. That's a... that's a type of thing that we could... we might expect will be possible. One of the things that's been a feature of, sort of, data taking in modern times as opposed to ancient times, is back in the day, data taking was somebody looking at a dial and, you know, writing something in a, you know, in a lab notebook.

And then there started to be automated data taking, and there started to be high bandwidth automated data taking, you know, where there's, you know, one image and it's a gigapixel or something, or there's one set of detectors from a particle collision, and it's a terabyte of data, or something like this. Being able to ingest giant amounts of data, sort of in one gulp.

It's something that's new in modern times.

And I think that, in many cases, what one's getting from that data is still sort of a fairly... the thing one already knew one could get, like one's getting an image from that gigapixel. What is one getting from the particle collision? Well, it's a bit more complicated, but it's fundamentally mapped into an image where there are a bunch of particle tracks going here and there.

I think one can imagine that there are features of, kind of, the... what one can see in, kind of, molecular structure of things, or for that matter, things that happen at the level of particles and particle physics and so on.

where there's sort of a large gulp of data that comes in, and then the things one's doing to it might be a little different than the things we have traditionally done. I mean, if we look at our human senses.

We have, you know, we have vision, which is very two-dimensional in character. It's very much sort of mapping our, sort of, light levels in this two-dimensional grid.

We have sound, where we're dealing with things where the frequency spectrum is important, and we're sort of taking the spectrogram as kind of the thing we're analyzing. Then we have things like taste and smell, where it's still a bit unclear what it is that we're exactly picking up.

I mean, for example, it's, in the case of, particularly of taste, you know, it matters not just what molecules are in the thing you're eating, but how those molecules are embedded in the thing

you're eating, how the thing gets sort of dissolved into the taste bud, or whatever. In smell, well, it's still not completely clear what

we are sensing in smell. You know, in vision, we have maybe... we have three kinds of color receptors. In smell, we have a few hundred kinds, it seems, of olfactory receptors. But exactly what aspect of molecules those are really sensing is not clear. Maybe they're just 400...

sort of proteins that have random shapes in them that molecules randomly stick into, maybe a little bit like the way that the immune system works with antibodies, or maybe there's something we can say that's more, sort of more structured about how smell works. But, for example, if there was a lot to be learned

from the detailed pattern of molecules that could be detected in a chemical way by something like smell, I don't think we would know it yet. And, you know, that's a thing. If we look at areas, well, for example.

One area where there's clearly things to be learned is in biology, where we know, yes, if we grind up this biological cell or something like that, it will contain this protein and that protein and so on. But, for example, knowing how those things are arranged in the cell is not something that we can immediately get

good information about them, and we can kind of freeze the cell in various ways and get some kind of frozen version of what's happening, but even there, we can't sort of gulp in a big chunk of chemical information and immediately know... and know at the level of sort of how things are arranged. We can do that at the level of what's on the genome, or what

proteins are active, but we can't do it when it's a question of how the thing is laid out in the cell, and that's probably another... that's probably an example of something where if we had that information, if we had that extra sense, so to speak, to have that information, then we would learn things from it. I mean, certainly the history of lots of science has...

has involved, well, when we can measure something, then we can learn things from that measurement. Whether it's X-rays, whether it's eventually various kinds of astronomy with different frequencies of electromagnetic waves other than light, whether it's ultraviolet astronomy, X-ray astronomy, gamma ray astronomy, these kinds of things, or radio astronomy for that matter. These are all places where we learn things we wouldn't have otherwise learnt. I mean, we wouldn't know about pulsars if it wasn't for radio astronomy. We wouldn't, and then...

when it comes to other kinds of media, like gravitational waves and neutrinos and so on, we get to see yet other features of the universe. So I think it is very much the case that, sort of, new senses, if you can call it that.

will lead to, sort of, new areas of science, and some of those senses might be things where you're detecting some, you know, gravitational waves or something, which you just couldn't physically detect before, and other things, I suspect, will be things where you could, sort of, one at a time, detect those particular molecules, but when you can kind of gulp in the data.

from, like you can sort of gulp in the data of an image, then I suspect there will be new things that can be discovered.

Let's see...

Pelenium is asking, could machines already be sensing things in data that humans simply can't perceive yet?

In some sense, yes. I mean, and we know that

with, kind of, machine learning methods, you can pull out features of data that humans hadn't noticed. I mean, we can even do that at the level of statistics from 100 years ago, or whatever else. But...

The question of, of, of, of sort of what

One question would be, are there general, sort of, features of data that...

We could be giving names to.

But we just never thought to do it. So, for example, in statistics, we might say, you know, the median of the distribution is this. You know, halfway lining up all of the values, you know, we go halfway through that, we get the median of the distribution. That's a thing where we sort of have an intuitive understanding of what a median is in a distribution.

Well, it could be the case that you've got a whole big pile of data, and there are certain kinds of features of that data which aren't just sort of black box things that machine learning can pull out to find an anomaly in that data or whatever else, but they're things where if we could identify that feature of the data, it would be something about which we could get intuition.

So, for example, something like a mean or median is

It's something where we know, well, that's the kind of the... that's the... that's often called the central tendency of the distribution of data. That's the thing that kind of gives you roughly where is the data, you know, where's the center of the data? Well, that's a good thing if you have fairly low dimensional data. If you have incredibly high dimensional data, that's probably less useful, but there may be features of very high dimensional data.

that are for typical data that one gets in the world, there may be things where, oh, it's typical that there's this or that pattern of arrangement of that high-dimensional data, but it's not something that we've ever sort of noticed on a large scale, but it might be something that occurs repeatedly. I mean, this is kind of the issue.

Is that, you know, a lot of what we do in science, or in language, human language.

is we come up with concepts that are repeatedly useful, and then we give them words, and then we sort of build knowledge on top of those concepts. The question is, when is there going to be a reusable concept?

If it's just, oh, this particular data set, we could pick out an anomaly in this way, in this kind of black box machine learning system, notice this or that thing.

That's not something about which we're probably, unless it's a very significant data set of some kind, it's not a thing about which we're going to invent a word and be able to build a whole sort of tower of intuition on that basis. But maybe there are things which are general features that have that property. Let me give you an example.

So...

For forever, people have been... have noticed sort of periodic aspects of data. The data repeats every so often. The data can be, represented by some Fourier, set of Fourier coefficients that say what the kind of amount of data that has a certain frequency is.

So those are things that were kind of known for, well, in some form, for many hundreds of years, but in a pretty clear form for maybe 150 years.

But another regularity of data is nested data, practical data, data in which there's kind of a tree of pieces in the data, so that there's... so there's a big piece that's similar to the little piece.

That's kind of the feature of sort of nested data. That wasn't something that people kind of noticed in data.

Until probably the 1980s, with kind of the idea of fractals and so on, and the promotion of fractals and things like this. People started noticing, oh, that's kind of a... a fractal-like piece of data.

sort of a precursor to that had been kind of hierarchical data and thinking about data in terms of trees, but certainly at the level of sort of visual appearance, that was a thing of the 1980s, to think about things as kind of nested structure. So that's a thing where we,

We didn't have... we sort of didn't have that

That... those words, that kind of set of things that was kind of a... a way of describing a general class of characteristics of data.

I mean, I was very struck by the fact that I had found, at one point, a very early example of a fractal pattern that was done in maybe 1210 AD.

And, was done in a mosaic, and those mosaics spread into many places. And if you looked at art history books that talked about those mosaics, the nested patterns that were in those mosaics were just not commented on in the art history books. It was like people were blind to that phenomenon until it had a name. Fractals, nesting, whatever. And then people are like, oh, this thing has

It has a nested pattern in it, or something.

So, you know, and I think there are things that, yes, you could ask the question, would a machine learning system, would an AI system, have discovered nesting even if we don't know about it? I think the answer is quite possibly yes, and then the question is to know that it was something that occurred often enough that it was worth sort of giving a name to it and building intuition on the basis of it.

There have been a number of places where people have sort of thought, well, maybe this will be a thing that is very widely... occurs very widely in the natural world, or the world at large.

But it didn't pan out that way. So, for example, solitons are an example of that. They occur in some places, but they're not so ubiquitously common. Or wavelets, another example.

Which can be useful, and so on, but it isn't like, sort of, they show up all over the place as a clear, obvious description of things in, for example, the natural world.

So I think the, you know, are there more of these? There might be. I mean, in my way of thinking about these things, these are the kind of slices of computational reducibility, places where you don't have to do every step of the computation to know what's going on, where you can kind of jump ahead and, and see the results of the computation. Finding...

The pockets of reducibility in the world is sort of a key piece of what one is trying to do in science and in analysis.

Let's see...

Let's see...

Random is asking, could there be any patterns in the universe that our current senses filter out without us realizing it?

It's a good question. I think the answer is quite possibly yes. I mean...

You don't know what to look for until you've seen some, you know... Okay, an awful lot of discoveries in science have been the result of, sort of.

Closely looking at something that one hadn't bothered to look at before.

certainly some of my own discoveries about the fact that very simple programs can produce very complicated behavior. People had seen such phenomena before, but they just kind of ignored them. They didn't have a sort of conceptual framework for thinking about those kinds of things, and so they were just like, oh, that's kind of a boring case where the thing

produces noisy results. We don't really care about that.

That turns out to be sort of a pillar on which a huge amount of science can be built. But until you have that kind of conceptual framework, until you realize that, just seems like, oh, a sort of an irrelevant effect.

And that's something that has showed up over and over again in the history of science, and that is a thing where you just have to have, for whatever reason, you just have to get that little thread where you can start pulling, realizing that something sort of isn't as you thought it was, and then you start to potentially be able to realize, oh, we should look at this in a different way, we should look at it at all.

And, and, and find these things.

Let's see...

Well, let's see, there's a question here from LC. Do you expect the near future of human employment, if any, will be bring-your-own-personal agentic swarm, akin to the bring-your-own-device transition some years ago?

Well, let's see. There's several things to say here. I mean, I think... that...

Automation is something that has progressively increased in the history of civilization, and, you know, it's been one of the real differences over the course of history. I mean, human nature has not changed much in the course of human history, so far as we know, and but technology has changed.

And, you know, a lot of that's about, sort of, more automation of lots of kinds of things. We have another round of that happening right now with AI systems, and lots of things that, in fact, one hadn't imagined would be automatable turn out to be automatable.

Some things are not. And, I think the, the thing that,

Is, the thing that ends up happening is, once things become automatable, you can then build further on top of the things which are now kind of squashed down because they've been automated.

I mean, the classic example of this, if you look at sort of the history of jobs and so on, is, oh, things like agriculture and such like. You know, there was a time when that was very non-automated, and people kind of had to, go and,

You know, plow the fields themselves, and then that got automated, and the result of that automation

Was people had more time to do other kinds of things, and lots of other industries got started that weren't, that were enabled by the fact that people... that that piece of what people were spending, you know, a large fraction of time doing got automated away.

So, I fully expect that as we see various kinds of automation come in, that will enable other types of industries to... to arise. I mean, in...

There are all kinds of examples of things that one does which are the result of, kind of, many layers of automation, and without those, one just wouldn't be able to do those kinds of jobs and activities and such like.

I think that, this question of, sort of what, what is the typical

Wherewithal that people have you know, it's pretty much assumed these days that people have smartphones.

You know, a lot of the world is made up now of, oh, the, you know, you download the app to do this or that thing.

that was something that came in, somewhat, you know, it came in fairly quickly. I think this, oh, you know, let me,

you know, let me have my agent deal with that. Let me have this kind of, some amount of automation. We've already got tons of automation. I mean, to give an example, I mean, you know, when you, sort of buy something, and you do it by having your watch communicate with NFC, you know, with near-field radio, with some little box on a counter, and so on. That is an... it's an amazingly automated process relative to, let me reach into my pocket and count out some number of banknotes or something to give to somebody.

It's, it's something where there are many layers of authentication and this and that.

going on, you know, at your beck and call, so to speak, to make that transaction happen. And I think there will probably be some additional things of that kind. I mean, another area of that is things like dealing with things like spam.

You know, the fact that, you know, it could be the case that everybody looks at every raw message that comes in to them, but in fact, there's a certain level of automation, a certain level of agenting, so to speak, that's going on there, of saying, only let through the things that are plausibly what this person actually wants to see.

And one fully expects that there'll be a bit more of that kind of thing. I mean, I think, you know, there's also then the question of, you know, how does, sort of, society adapt to some of those types of things? How does society adapt to more and more things being automated in certain ways? I mean, in other words, for example, right now.

you know, I get a lot of mail from all sorts of people about all sorts of things, and when does it become how socially acceptable to have the first round of response to that mail be obviously an automated response?

And when does it... when is it kind of like, oh, I've got an LLM that's running against previous responses I've made, or against things I've written, or whatever else? When does it sort of socially acceptable to,

To do that. I mean, that... that was a thing that had to come in slowly for society. Like, there was a time when, you kind of, you know, if you wanted to have certain kind of conversation with somebody, you'd always have to do it in person.

And then that became a thing that you could do by phone.

Or you won't have another kind of conversation with somebody, you know, again, you get them on the phone, and then it was like, oh, I'll just send them an email, or I'll text them, or something like this. There are things where it becomes sort of socially acceptable to do that, and it wasn't before, and there will become a question of, when is it socially acceptable to have kind of my agent talk to your agent type thing, as opposed to people interacting for themselves, so to speak. I mean, I think an interesting question is in the area of things like sales.

when, when it's kind of agents talking to agents, and I mean automated agents talking to automated agents, to what extent, you know, how will that actually work? I mean, because a lot of...

You know, when you're buying high-priced objects, for example, there's a lot of, kind of, human interaction that ends up being typically part of that process. And if that's all going, sort of, through the agents, I don't know what that will look like, and probably new things will develop that are new and different ways to, for example, sell things, as have developed with, you know, people have their cool things on Instagram or something, rather than coming to my store and see it, kind of thing.

But I think,

you know, the question of what will happen to human employment, the thing that tends to happen, my observation, is that, sort of, automation allows more possibilities. More possibilities allows more choice.

choice is something that ends up getting made by humans. It's like, you know, you can have, oh, I don't know, something where,

You... well, both... there's,

It's like there are many things you could do, and well, like, for example, let's take podcasts as an example.

One could say, oh,

Automation is coming in that area, okay, the AIs will just generate all the podcasts. I don't think anybody believes that that's a plausible thing. It's, you know, podcasts, somebody identifies, there's this thing that I, as a human, think people will care about, and I'm going to make a podcast about it.

It's not something where it's like, let's have the AI just sort of decide what all the podcasts are going to be, because that's not the kind of thing that an AI usefully does.

Just, like, it's... it's... there's a... there's a certain, sort of, in what direction do the humans want to go? And that's something that there's no right answer to. It's just a question of the direction the humans actually happen to choose to go in.

And I tend to think that the kind of automation makes possible many more choices, and the choices end up being sort of, in some sense, arbitrary human choices, which end up getting made by humans. And my guess is that that will represent a larger and larger fraction of the economy.

I mean, that's one piece. The other piece is, when are there roles where humans feel that they need to be interacting with other humans?

you know, will teaching, for example, be a thing that becomes a thing that can be automated, or is it something where it really matters that the person who's learning stuff, the kid or whatever, actually has a human that they're trying to convince of something? Do they care, or do they not care?

I mean, there was a moment when, you know, various kinds of trams and so on started to not have human drivers.

And people were sort of okay with that. I think...

That the question of a, you know, it's... and now there are starting to be self-driving taxis and so on, where...

You know, people seem to be okay with that, too. I don't know whether people are going to be okay with the self-flying airplanes and so on, where, you know, you're there as a passenger, but there's no pilot on the line as well.

I don't know. You know, these things, it's... there's questions about where will... where will humans feel the need to have other humans involved versus have something where it's... it's just a machine doing it, so to speak.

I think...

I mean, I have to say, I think that the current round of, kind of, endangered occupations is a lot of, kind of, I would say, kind of, light judgment.

Heavy on, on, sort of, not quite bureaucracy, but on... on kind of heavy on, on sort of structure, on form, you know, in... in the kind of,

white-collar sector. I think a lot of, kind of, the kind of... the sort of mid-range of a lot of professions, there's a lot where it's kind of like, we know this is a thing you can do by turning a crank. It's been quite difficult to turn that crank, but now we can automate doing that.

I would say the high ends of professions, much less automatable, because that's a place where much more is being done, sort of, for the first time. That's a place where, sort of, the high ends of professions are places where being able to use that automation lets people reach further, but it's not something where it overtakes what's possible there. Now, of course, the order of things can change, because there will be people who just don't use the automation, and the people who do use the automation have a big advantage. I've seen that in things I've been involved in building, like the whole building of computational tools for doing science and technology and so on. It's, you know, there was a time when there were people who, you know, oh no, I'm never going to use, you know, mathematical or Wolfram language or whatever, I'm just going to do my math by hand or something.

Or I'm not going to use, you know, I'm not going to use computation for things, I'm just going to do it some other way. And yes, you know, using the latest and best tools and sort of always being at the leading edge of automation is a way that you kind of stay ahead if you're kind of leading a field. If you're doing something where fundamentally the thing you're doing is something that can be automated, then that's something where that part of the, kind of, the mid-range activity in the field might just go away. I mean, we're seeing that now in a fairly dramatic way with programming.

I think that, you know, for somebody like me, I've spent the last close to... close to half a century kind of working to automate the process of programming. I mean, in a sense, you can view Waltham language as being something that is building a huge amount of automation of what would normally be, sort of, specific line-by-line programming. It's kind of identifying the lumps of computational work

That are repeatedly used, building them in a nice way, and putting together many thousands of them, and having that be the thing that you're working with, rather than the kind of low-level, traditional programming language picture.

Well...

The... it is interesting to me, and it's been an interesting thing to watch, that if you look across lots of different fields.

People who are sort of leading people in those fields are people who end up using our tech, using alternate language, and so on, to... to make progress in what they're doing.

their kind of typical workflow is, I have an idea, I want the best tools I can for taking that idea and making something real out of it. And... and we've been the very dominant thing that's been used for that particular purpose.

Well, then, it's like... and then once they've actualized one idea, they go on to the next idea.

But then, there are millions of programmers. What do programmers typically do? Programmers are given some ideas, some spec, and their job is turn that into code.

It's a much more kind of, it's not a thing where there's a high density of ideas. It's a... it's a... it's a... it's a trade where one is turning ideas into code.

I don't think it should have been necessary to do anything like the amount of that that's been done over the last, I don't know, several decades. You know, we'd already automated a lot of that stuff. The fact that there are millions of programmers grinding out, you know, millions of lines of code, sort of line by line, I just think is crazy.

But the fact is that the reason that's happened is because there's a whole profession that has to do with, I'm given a spec, now I grind out code for that spec. It's not about I make up the spec, it's about I grind out code for the spec.

And so, often, you know, if you feed programmers, sort of, kind of mid-range programmers, kind of wolfram language, they might say, well, you know, they can do what they're trying to do in, you know, what was supposed to take them a month, they can do it in 2 hours or something, but then they're like, oh, but it's really difficult. What I have to do next is really hard, because the next thing I have to do is have an idea about what to do next, and that's not my job. My job is to go from an idea that came from somewhere else.

And grind out code.

What's happened now is it's clear that that grinding out of code, I have thought was unnecessary for decades. I mean, to be clear, I've thought that was a thing which we had successfully automated a large fraction of a long, long time ago.

Now, there's sort of a new game in town, which is have AIs grind out that code. and go from sort of vague description to the actual ground outlines of code. And the activity of sort of step-by-step finding out the code is going to go away.

Now, I have to say, I think that the thing that's sort of interesting from our point of view is that Wolfram language is this way of describing what you want to have happen. You can say, okay, I'm going to describe in English roughly what I want, maybe the system will ask you some questions about, do you want this or that?

And for some purposes, that will be completely adequate. If you're trying to build, you know, some, I don't know, some basic user interface where there's no... you're not going to... you're just... the user interface is the final result. You're not going to build some tower of functionality on top of it, just like, I want to have a slider roughly here, and this roughly there, and so on.

You can probably just describe that, or you can make a picture of it, you can draw a picture of it on a napkin, show it to the AI, and it'll go and build that code for you.

I think that, but if you have something where it is a component in what could be a very big tower of functionality, you need something which has a better definition. And you kind of need a way of having the AI say, this is what I think I'm building.

You know, one definition of what's being built is the code.

But if the code is kind of low-level programming language code, it's like, well, I just built this thing, it's very impressive, I built, says the AI, a thousand lines of code for you. Is that what you want?

well, people can't read all of that code. Maybe they can try running it and see if it works in one case or a few cases, but they can't know all the corners of those thousand lines of code.

The thing that is really an important and unique characteristic of orphan language is that we've tried to build this kind of notation for computation that allows one to get sort of a succinct summary of what the computation is supposed to be.

And if the AI surfaces that, you have a fighting chance of being able to quickly read that, glance at it, and understand at a precise level what the computation that's going to be done is, and you say, okay, that's right, now put that as a sort of solid brick you can build your whole tower with.

And I think that, I mean, it's kind of like...

Imagine if we didn't have mathematical notation.

And we were just using AIs, and we were trying to have them do mathy kinds of things, and it was, and they were describing everything in words. One would get lost very quickly, one would have no idea what, you know, what was being talked about, and so on. But we have, in math.

from 400 or 500 years ago, a fairly compact notation for lots of kinds of things that allows us to sort of tell at a glance roughly what's being talked about mathematically. What Wolfram Language is doing is to provide that same kind of capability for computation.

But it also has the feature that once you have that precise specification, that precise kind of symbolic specification of what you're doing, you can immediately also execute it and have it be an active piece of what you're doing.

So, I mean, I think that, these, these areas, like, you know, the sort of mid-range areas of a bunch of kinds of, I don't know, trades, professions, I don't know what you exactly call them, yes, they will go away.

And, you know, it's an interesting question, what people have the skills for those particular trades, what

How those skills can be refactored into something that

is valuable, even when the AIs have automated that particular thing away. I mean, my general theory of things tends to be that us humans are pretty flexible in our capabilities, and there's usually sort of a core thing that is what somebody is good at. Like, it might be problem solving, or it might be sort of keeping a team together, or something like this.

These are things that are independent of the particular details of writing this line of code and that line of code, and so on, and those are things that are sort of refactorable into other places, even if the particular trade of writing lines of code isn't something that sticks around.

I mean, I think that the thing that we will see is, you know, eventually there will be, sort of, successful humanoid robots, and eventually a lot of, kind of, more, sort of, blue-collar-type tasks will end up being automatable. I mean, the thing that's sort of a surprise to everyone with... that came with a sort of chat GPT surprise was that suddenly it looked like a bunch of sort of white-collar-ish jobs, were automatable. What one wasn't expecting would be the things that were automated.

But I think that there will be a lot of, kind of, more physical jobs that get automated once humanoid robotics really works, and it will eventually work. How quickly and how suddenly it will come in, I don't know.

And then the question is, well, okay, what do the humans do then? And I think the thing that is sort of remarkable is if you look over the course of time, and you say, well, do humans end up with nothing to do?

Well, no, humans find all sorts of different kinds of things, you know, they become podcasters, they become influencers, whatever that might mean. They become, you know, professional esports players. They become people, you know, all sorts of different kinds of roles.

That didn't exist 20 years ago.

And, and have sort of come into existence because of the automation that's been possible, and so on. And I fully expect that to continue. I think that... I think it's very hard to predict what are the kinds of things that,

we humans will decide are worth doing in the future when our world has changed a bit. I mean, there are plenty of things which

In the past, you know, in a couple of hundred years ago, one might have imagined, is somebody going to be... I don't know what,

doing, oh, what's an example?

well, all sorts of social media kinds of things. Is somebody going to be doing that? Well, nobody would imagine it. It wouldn't be a thing you could... you could conceive of. It wouldn't be a thing you would expect. Now, looking back.

you realize that the things that happen on social media and so on are reflections of human nature, and they're things that, in a sense, it's not surprising, work that way, because you can kind of see an analog of them in the way people were in villages back in whenever.

And, but it would be very hard to predict, going forward, how... what that would look like. And I feel like the same thing is going to happen now, that there are, you know, it's... well, it already happened to some extent with AIs, and, you know, there were the, you know, prompt engineering was a thing, and then that got sort of somewhat automated.

And it became easier to write prompts. I think, sort of, the AI psychologist, I don't know if that's going to be a big thing or not, of understand why did the AI do what it did.

there was a period of time when it looked like, sort of, AI ethics was going to be a big thing.

Turns out that's a very difficult field, and it's not so easy to really have useful things to say about it. It will be important, though, and that's a thing that, you know, there'll be more of, I'm sure.

But

So there are lots of kinds of things like that. Now, you know, the question, I think this original question asked, in the future of human employment, if any.

Well,

it's a question whether, sort of, sort of, the robots are going to make everything for us, the AIs are going to do all the figuring out we need to do, et cetera, et cetera, et cetera. Will there come a time when we can just sort of sit back, relax, enjoy the ride, whatever that means? I don't think that's very consistent with human nature. It's just like kind of these schemes for society and for economics and so on, where people say, oh, we'll make this kind of utopian setup.

And everything will be... will be wonderful.

But it just turns out that's not very compatible with human nature. It isn't the way that... that us humans

for whatever reason have, you know, evolved our biology and so on, we're stuck with the way we are, and I think that it's something where, you know, you can imagine, kind of, worlds where, sort of, we just sit back, relax, and enjoy the ride. I just don't think that's something that people will end up doing.

I think that...

things will happen that, people will find things to do. They may be things to do that, to us, even today, look like, wow, that's a thing people are doing? I never would have guessed.

I mean, I remember back in the day, people saying, oh, you know, cable television is coming in, you know, there'll be so many channels, there'll even be a channel that's showing the weather all the time.

Of course, that became a very popular channel. Probably there are many such channels, I don't know, I don't watch television. But, you know, it's like it was inconceivable. It was like that people would just, like, sit watching a weather channel, so to speak.

And I think that will happen again with things that seem inconceivable to us now, and which are maybe things that are very small effects now, where it's like, yeah, that's the thing one does, but it isn't a big thing, but then when it becomes easy enough, it will become a big thing.

So, for example, I mean, things that, I remember, you know, when search engines were coming in for the web, you could say, well, how many search engines would anybody need? Well, how many reference librarians are there in the world? Well, not very many.

But the fact is, it was a big effort to go ask a reference librarian a question. By the time you have web search, it becomes very easy to ask that question.

I mean, the same has happened in science, a lot of science I've done, other people have done, with the tools we've built.

that science has become possible. It's kind of come over the horizon because those tools exist, and the science becomes possible. I mean, this whole every riology of studying simple rules and what they do, that is entirely enabled by tools we've built.

And without them, that would just never, never be a possible thing.

I think the, The, the... this question of sort of what, you know, is there...

Well, another thing to say is.

People will complain about the modern world. They've complained about the world as it is at their time, forever, I suspect.

But the fact is, in many respects, you know, most of us have it pretty easy, relative to what the world was like 500 years ago, a thousand years ago, whatever else. You know, it's, there's a lot that you kind of don't have to worry about.

in most of the time. There's a lot where you could have to spend all your time, you know, doing this or that thing, growing your food, whatever else it is. But we don't. You know, we can just buy food at a store type thing. And I think

It's, so you might say that relative to 500 years ago, oh my gosh, you know, we have a cushy time, we're just sitting back, relaxing, and enjoying the ride.

you know, we have more automation coming that will allow us to do more of that, but I don't think it's going to be a situation where it's kind of like people have nothing to do, so to speak.

You know, it's always a funny thing when you try and interpret anthropological research.

And people say, oh, there was some society somewhere where people... and the story was something like, oh, it was a place where there was abundant food, and, you know, for thousands of years, people could just sort of hang out and enjoy the ride.

And then what did they do? Well, the anthropologist's description is they engaged in a lot of ritual behavior.

Well, what does that mean? Well, they were doing a bunch of stuff.

And they were, you know, who knows what they were doing. They were, you know, arranging things in this way, and chanting this or that thing, and so on and so on and so on.

And, you know, from the outside, not knowing that culture, you might just say, oh, they were just, you know, it was meaningless. They were just doing ritual things, because they had no, you know, because they... because life was easy and they were just sort of enjoying the ride.

But I suspect if you're inside that situation, it's like, no, no, no, you've got a lot of effort you're going to, you've got to figure out, you know, how do you get ready for the harvest festival, but this way, or that way, or whatever it is. And it's... it's kind of, you know, just as much hard work, so to speak, as one's doing in other situations.

I suspect that the perception of what needs to be done may be somewhat constant over time. I mean, look, I think that, you know, as a... it does tend to be the case that in those

In those situations where people are like, you know, you make... life has been made very easy for you, and you haven't kind of figured out what you want to do, somehow... somehow the fact that it's been made easy is...

you know, hasn't... hasn't opened up new possibilities. I think that tends to end rather poorly.

And, you know, one sees that in a lot of,

A lot of unfortunate societal kinds of things.

I think that's not... you know, the fact that things were made easier for us does not necessarily lead to that. That's kind of a cultural, societal choice.

And perhaps the thing that is a consequence of... I mean, I think there are... there are always things to do. People may not choose to do them. So I... I don't know, I think, it's,

It's an interesting question as I... to think about it, you know, transitions that we've seen happen, like at the beginning of the Industrial Revolution, you know, the invention of the Jacquard loom and things like this. What happened to the Weebers?

What happened to the individual weavers? I think that might be an interesting story, I should research the history more carefully, but here's my guess. There were people who were doing generic weaving.

And those people were sort of put out of business by the automation.

Then there were people who were kind of designing the fabrics and figuring out, you know, the... where do the birds and flowers go on this fabric, and so on. My guess is those people were just fine. In fact, those people probably had more to do and could produce more designs, more fabric, and so on. I think it's the kind of... the things that... and, you know, the analogy to the programmers, so to speak.

is the people who are grinding out the lines of code, they're like the weavers, the people who are figuring out the designs for the textiles, they're kind of like the CTOs and the architects and so on, who are doing sort of the computational thinking necessary to figure out, sort of what to build, rather than doing the step-by-step weave the next piece of the fabric, so to speak.

It's an interesting question what the people who do that should be called, because it's not really programming. It's not...

you know, when you're doing the functional design for some system, you're thinking through, kind of, how does this really work computationally? That's not really programming. You don't have to be able to write a single line of code to do that well.

But, so what... what is a person who does that? It's kind of like, you know, I was thinking it's kind of like a computationalist.

At least that's my best current idea for what such a person should be called. But there's going to be a lot of jobs for computationalists. As the actual, you know, cost of programming, so to speak, goes down, it should have gone down years ago. If people had been using our tech more widely, it would have gone down more widely years ago. It's gone down for plenty of people who have been using our tech.

But it's now firmly going down, and that means there's sort of more opportunity for the computationalists. By the way, if anybody has a better idea than the word computationalist to describe a person who does that role, I'd be really interested, because, as I say, I think there are going to be lots more of them.

Well, okay, a few, a few comments there.

Let's see...

Mayor is asking.

Do you think future scientists will spend more time choosing which questions to ask than actually solving them?

In some respects, I mean.

the original seed question is often not the only strategic thing you have to deal with. By the time you are, sort of, by the time you've finished the strategy, and now it's just mechanism, now it's just, like, solve the problem, then, you know, that's increasingly automatable.

But my observation from doing lots of science, for example, is that there are many strategies

strategy decisions that have to be made along the way. Typically, what happens is, you have a seed idea, and you say, okay, I'm going to go see what I can discover along these lines. And then you discover things, and they weren't the things you necessarily expected to discover. And then you have strategic decisions to make about where you go next.

That's been what I've seen happen. I've seen that happen both with science and with technology. When you're building something, you've got a piece of it built, you start using it, it's like, oh, actually, the thing that's important about this thing isn't the thing I thought was important about the thing.

So I think that it's not a kind of fire-and-forget type approach that we can expect to see. It's the strategy points that become more and more important. I have to say, people who've been using our tech for years, that really is the thing that's happening a lot of the time. It's like, the actual doing of the implementation of the thing is not terribly hard. It's a question of, okay, what do I do next? How do I conceptualize

what's going on? What do I build next?

Let's see, Isaac is asking, do you see the end of AI coming with news about NVIDIA pulling back investments and so on?

Look, this is such a typical thing. It's such a typical kind of hype process, so to speak. that some new thing comes out, it's surprisingly successful, people are like, oh my gosh, it's gonna go to the moon. And people start making big investments in it, they start saying, you know, every company has to have an AI strategy, and a chief AI officer, and so on. And it's, and that gets overdone at the beginning.

And then... and then sort of things... people say, oh, no, no, this isn't really working out, it's all... it's all nonsense, and the thing sort of crashes down for a while.

But actually, there was something real there. It just wasn't... it wasn't the full go-to-the-moon on day one thing that was real. But there was something real there, and over the course of years, it becomes the case that there's a very solid thing that will build over time that is a methodology that's useful. And for sure, AI and current AI technology is useful for many purposes. I mean, I think in...

In our own company, you know, we're seeing, as one has better and better harnesses on the AI, the core AI technology hasn't changed that much.

But better harnessing has happened over the last few years, and for us as a company, for example, more and more things we do, we are able to put AI in the loop, so to speak. And we're able to let the humans kind of, sort of rise to a higher level.

And, or we need... we need less humans to do the particular things we were doing, and the humans who were doing those things can go and do other things now. So I think... but I think what's, you know, what tends to happen is these things overrun and then get pulled back, and so on.

the detailed dynamics of what's happening here, and the detailed, sort of, personalities of companies and so on, is... is a whole interesting circus. I mean, I know many of the people involved, and I've known them for a long time, and

it's, you know, it's sometimes remarkable that they sort of landed in a place where there's suddenly this huge success around them, and, and I think

The, the sort of the question of

of, sort of how... how sensible are the things that are being claimed? How much is it all kind of a little bit of a scheme, so to speak, as opposed to being kind of all on the intellectually straight

and narrow? Those are interesting questions. I mean, when you see all these benchmarks, you know.

the amount of, kind of, weird cheating on benchmarks is amazing. And it's like, what does this actually mean? What was actually done? It's reminiscent of of, you know, when you... when you see, things where you're doing, you know, clinical trials in medicine or something. Those are... those are things where there's a whole giant effort to kind of figure out

how to optimize those things. You know, there's something real there, probably, but you have to kind of puff it up and optimize it to present it in a way that looks the best it can. And I think this is sort of the same kind of thing with a bunch of these benchmarks and things like this that are not what they seem to be. And that sort of all feeds into kind of a PR story

About, you know, amazing things are happening.

Well, good things are happening. Whether the, you know, peak amazement things are happening, probably less convincing.

And, you know, that... that leads to people saying, we're going to rush in here and invest a ton, and like, whoops, that was a cheat. We're not. That was a scam. We're going to pull out of that. So, you know, I think one's seeing a bit of that turbulence going on, and it'll be, you know, there's real things there. It's probably not, you know, the maximum level of hype, not really correct.

But that doesn't mean there's nothing there. And one shouldn't equate, you know, excessive hype to zero content. It usually doesn't work that way. Excessive hype, there's usually something underneath it, it just isn't as big as what one was led to believe.

Let's see...

HIP is saying, in the AI revolution, is it important to distinguish between the concepts of intelligence and consciousness and self-awareness? And also, can there be intelligent machines, including those more intelligent than humans, but unconscious, as LLM models currently appear to be? You know, this consciousness thing, which obviously I've worked on a certain amount, is a slippery mess.

I mean, you know, we have a sense of our own, we have this own internal experience that we're having, which is what leads us to say, hey, we're here, we're conscious.

We can't say that about anything except ourselves, for sure.

We see other humans, we see that other humans kind of, kind of have the same general appearance that we have, and seem to have the same sort of,

Reactions to things that we have or something, and we say, yeah, if we know internally that we are whatever this conscious thing is, then we think they probably are too.

Now, if you're, you know, if you're talking to a human, but it's a text chat.

and you don't see the other human, all you have is their text chat, and then somebody says, oh, by the way, this thing you thought was a human that you were talking to, that was actually an AI. In modern times, I think much of the time, you won't be able to tell the difference between the human and the AI.

And this is the kind of classic Turing test thing, which seems to have been pretty firmly passed at this point.

So now the question is.

do you then conclude, oh, whoops, it was a, you know, it was... I saw behind the curtain, it was a machine, it wasn't a human, therefore it can't be conscious in the same sense. No. Operationally,

from what you could see, it was doing just the same as a human is doing, and you should assume the same level of consciousness

Whatever that means, as in a human.

you should assume that its inner experience... I mean, as I say, the only inner experience that we know for sure is our own inner experience as individuals, and it is by pure extrapolation that we conclude that other humans have similar inner experiences. We don't tend to do that extrapolation to machines, because they're so different from us, they seem so different from us. But yet, if we're exposed only to the text channel or whatever.

they will seem very similar. I don't think... I think it's a slippery mess.

to say that, you know, we can say, oh, no, no, they're not conscious, why aren't they conscious?

Well, you know, because we can see what their neurons are doing inside. Well, eventually we'll be able to see what human neurons are doing inside, but will that mean, if we can see that you know, the next word I say originates in this part of my brain and then spreads with activation spreading to that part of the brain. We can see the whole story of how that word was made. Does that mean that I am not acting consciously to make that word? I don't think so. I think that the... the fact... it so happens that brains are still very much black boxes. We don't know what's happening inside them.

But that doesn't mean that if we did know what was happening inside, we could trace all the different neuron firings, that we would say, oh, that... that's a nothing. That, you know, that couldn't have this magic of consciousness, because after all, we can trace all the neuron firings. I mean, I think this is partly a thing that comes out of thinking about things like my concept of computational irreducibility, that

even if you know the rules by which something operates, it may be the case that the only way to find out what it does is to step-by-step follow those rules and see what happens. And if one's in that situation, it's not the case that just sort of seeing how all the neurons work, we can just say, oh, we're done. There's nothing to that human brain, so to speak. There can still be this sort of irreducible computation that's being done by the human brain, and that can sort of add up to something.

And that's, I think, sort of the... I think, actually, that the essence of consciousness as we tend to experience it, and this is getting a bit more philosophical, is very much associated with this sort of single thread of experience that we all feel we have. You know, we get all this input from the outside world, and we kind of compress it down to this thread of experience where we can remember the past.

we can figure out what we're going to do next, and so on. And we're taking all this input data, and out of it, we're abstracting the small sequence of commands, so to speak, about what to do next.

I, you know, I tend to think that the most likely theory for how that notion of take all the input data, grind it down, and come up with a small number of instructions for what to do next, the reason that originated in the history of life on Earth is probably when animals started to be mobile, they had to make a collect... the whole animal had to make a decision. Am I going to go left or am I going to go right?

If the thing was just sort of this static.

you know, object that was, you know, plant-like thing spread over a big part... big, big region.

It's like, well, one part can do this, another part can do that. It didn't have this... have to have this kind of, sort of

collective consensus decision-making. And my guess is that's sort of the origin of our sort of single thread of consciousness, single thread of experience that we think of as so significant in our way of perceiving the world, in our inner experience.

I don't think that's really so significant, but, you know, it's a thing that is a sort of vaunted feature of our own consciousness. But I don't think it's, it's,

To say, sort of, to make a kind of... okay, the thing to say is, yeah, the machines aren't conscious in the way that we're conscious, but they... to say that we're going to have this word, conscious, and we're going to have... that's the big...

trigger for whether, for example, the machines can have rights, or whatever else it is, you know, I don't think that's the correct

thing to be thinking about, because I think that word is a slippery thing that is... is... is... we only have a clear definition when it's an extrapolation of our own individual inner experience.

Let's see... Elsie is commenting. So, will one get employed as a human individual or as one's personal agentic AI stack included.

That's an interesting question. Do you come with your tools, or do you come just with you?

Well, I mean, there are many, many professions where you come with your credentials, for example. You know, we need a certified whatever for this particular role.

or you come with, you know, you come with your Rolodex for, you know, if you're going to be networking for some role. So I think it's a very well-established thing that, you know, you're not just getting the raw human, so to speak, you're getting the human with their accoutrements of various kinds,

When you hire somebody from the role, and indeed, you know, it's like.

do I come? It's interesting, I mean, a lot of the time.

for various reasons of, sort of, the flow of intellectual property and so on, when somebody comes to do a new job, it's not like, well, bring all your tools from your old job.

That was a thing as sort of, you know, in the sort of labor market of the tech industry, the kind of the, let's make it open source allowed people to move from one job to another, sort of carrying their tools with them.

Which hadn't been really much of a thing before that time. So I kind of think, yes, people will sort of carry their tools with them, but they've always been doing that.

And, I don't think... I don't think it will be anything particularly dramatically new. It's really a question of, is the agentic swarm that you built around yourself, is that something where that was entangled with your employer's IP, or was that something where you sort of privately built your agentic swarm that you...

Carry to the next place, and so on.

Let's see... Gosh.

Oh, so many.

So many questions, interesting questions here.

Well, Aniel asks, how can we automate data visualization and user interfaces? Does this relate to observer theory quantifying us as observers, so we can measure the intelligibility and aesthetics of GUIs?

Well, I have to say, in building Wolfram Language, we've done a huge amount of automation of visualization and user interfaces. I mean, it's kind of like when you, say, I've got this, this, this network.

Just show me the network.

That's a huge amount of automation to actually lay out the network in a way that is sort of appropriate and pleasing for that particular thing, or put labels on this network, and so on. this question of whether... of how you can take, sort of, raw data and just, sort of, say, show me a visualization of this data, that's an interesting question.

We built, 15 years ago, we built in Wolfram Alpha, we can use it today, the pro version of Wolfram Alpha allows you to just upload a bunch of data, and we'll use a bunch of automation to just show you reasonable views of that data.

It's... it's useful. It's not as useful as I might like it to be, because it tends to be the case that when you have data, you usually have something that you're going for, something that you want to see in that data.

And sort of our goal in Wolfram Language is to make it easy to be able to say what that is, and then there's a lot of automation in what you actually get shown, and that's a, you know, a lot of why

Wolfram language is very widely used for visualization tasks is because we put the effort in to do that automation, whether it's the way that colors are chosen and setting it up so it works on light and dark screens, and whether it's dealing with how the dynamic interface works and so on. These are all sort of pieces of automation, and I think

I think that's a... this thing about automated aesthetics, we've done tons of work on that over the last 40 years, of kind of figuring out, how do we

How do we optimize the appearance of this thing, given what it's trying to communicate? And that's been a big story in the way that we built, kind of, visualization involved in language to this time.

Let's see... Zwei, Zwei...

is commenting, Freeman Dyson thought the great difference in the future of humanity will be genetic modification of both plants and humans. How feasible and how dangerous will this be?

I knew Freeman Dyson for a number of years fairly well. I saw him

close to every day when I worked at the Institute for Advanced Study in Princeton in the early 80s, I think Prima didn't like me very much.

I grew to not like him very much either. I... Freeman always had these kind of, you know, these sort of controversial views of things. Even if there was an obvious answer, he would always pick the non-obvious, controversial answer, so to speak, contrarian kind of, kind of point of view. Not always with great support.

But I think one of the things Freeman said to me, at least, was, oh, all this stuff with digital electronics and so on, that's all kind of irrelevant. The future is, you know, genetic modification and so on, and being able to make things biologically.

I think...

There's a part of that that is probably correct. That is that we're operating with electronics, you know, we have

Basically, where, you know, we might have a few nanometers of scale size for our features in a piece of microelectronics, in, you know, if we could do things truly at a molecular scale, it's, it's, you know, a few orders, a couple of orders magnitude smaller than that.

The other thing that's true about electronics is it tends to be very two-dimensional. We're kind of making that image that we sort of etch the chip out with, but it's a two-dimensional kind of thing. Whereas in biological organisms, we are

pretty three-dimensional, you know, our cells are, you know, blobs of stuff with all kinds of... all kind of, as I'm now calling it, bulk orchestration of molecular processes going on inside. Will we be able to control that bulk orchestration? We don't know how to do it yet.

Probably we will.

biology has a deeply inefficient way of doing it. That's the result of a few billion years of biological evolution. There surely is an engineering way to do that that will be, presumably, much better than the biological way to do it that was kind of gotten by this sort of series of steps of evolution. I mean, it's worth realizing that what's happened in biology through biological evolution, and through, kind of, progressively modify the genome, modify the program for the organism, get a better, you know, see if the organism is better or not. Keep going, kind of iterate that to optimize for some... for some fitness objective.

the genome of the organism. That's a very similar process to the one we go through in training a machine learning system. And in both cases, we're ending up with these kind of these lumps of computational work, these lumps of irreducible computation that get sort of stuffed together to do the kinds of things that we're trying to do. It's not what we get in a sort of designed engineering system where everything kind of fits perfectly together, it's just we happen to be able to put these lumps of irreducible computation together to do the things that we want to do.

And I think it's similar between biological evolution and machine learning. I've written about that quite a bit, actually. But,

The, the thing that, is... is sort of a... the examples we have right now of molecular scale, stuff are all in biology. And what we need in biology, we have machinery, ribosomes, and things like that, for making things on a molecular scale.

We don't have engineering machinery for doing that. We only have evolved machinery for doing that.

One day, probably, we will have engineered machinery for doing that, and then we will turn this sort of molecular level of things into something that is more traditional, like digital electronics.

As opposed to something where the only way we get to it is to, you know, cross-breed different kinds of plants and so on to get some effect. That's a... it's sort of similar to, do we use, you know, we found horses, so we can ride horses to get from here to there.

But the horse is a complicated thing. We're sort of, we're quite literally riding on top of, sort of, this natural system. When we build trains.

we have something that is engineered and is much more, kind of, optimized for its particular purpose. And I think that same thing will happen, when it comes to molecular scale... molecular scale capabilities. So I would say Freeman was... I will give Freeman credit for being sort of right

that the molecular scale of things is important, I think he's quite wrong about the fact that we will get there by sort of riding the donkeys, so to speak, and making use of life as it is right now.

We have been doing that so far, but there will come a time where we'll sort of solve the problem of making sort of molecular-scale devices that can create other molecular-scale devices and so on, and be able to sort of engineer with molecules.

I think, We should probably wrap up here.

You know, it's an interesting question. You guys are asking about, scenarios for, sort of, AI helping the science and so on.

you know, it's a good challenge. Can you invent a universal constructor? Can you invent a molecular constructor where you can sort of give it instructions and it will just start making things with molecules? That would be a... if that happens.

with, as a thing that's sort of an AI-generated thing, where the AI figures out how to do it, I'll be impressed.

I don't think it's completely out of range, out of the range of possibility, because I think it's something where you're kind of ingesting the literature that exists, and then trying to figure out kind of, you know, synthesizing that together to... together with properties of molecules and things like this, to see whether you can achieve that objective. It's an objective we have to imagine

And then we can say, can you fit the pieces together to actually get to that objective? That will be an interesting example of something that is kind of a potential sort of AI-generated major change. I mean, if that gets figured out, and we start being able to do molecular-scale manufacturing, then lots of things in the world will change. I mean, the thing I've often commented on.

is that, you know, it's only been 100 years that we've had sort of artificial materials, like plastics, and we make things out of those artificial materials. There will come a time when sort of everything is made of computers, where everything is made of, kind of, molecules that are doing known computations, not doing computations that sort of happen in the wild in nature, but doing computations that are things where

We sort of, in some sense, you know, the computationalists defined what it should do, and now something implemented that, and it's now actually doing it at a molecular scale, so to speak.

All right, I should wrap up there. Thanks for lots of interesting questions. I can see lots that are, can be queued up for another time. I will say, next Wednesday, I am starting my series of philosophy discussions.

with, professional philosophers, very scary for me, because, I'm, not in that business, in a sense. But,

talking about the relations between things I've been doing scientifically, and the kind of philosophical implications of those things which I've tried to work out, and particularly kind of classical thinking and philosophy. That's going to be the focus of a series of conversations, not least because I want to understand that stuff. You know, I read

fragments of what, you know, Immanuel Kant might have written, or Gottfried Leibniz might have written, or Aristotle might have written, for example, and I say, boy, that sounds an awful lot like what I'm saying now. I want to understand whether they really, sort of whether...

whether these ideas were sort of intuitively clear in sort of the... in philosophical history, and what we can learn from... from the other things that people

People like that figured out.

So, coming, first one of those coming next Wednesday.

All right, well, should wrap up here. Thanks for joining me, and talk to you another time.

Bye for now.