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

I see a number of questions saved up here.

Okay, here's a good meta question from Hunt.

What will doing science look like in 20 or even 50 years compared to today?

Okay, so here's something I get for being an old fogey, which is, I was already doing science 50 years ago.

And the one question to start off asking is, how's it changed over the last 50 years?

Well, the…

For me, 50 years ago, I had started to use computers to do science. That's gotten a lot easier. For me personally, it's… I have built a huge technology stack for doing science with computers. It was… I was already on that path even 50 years ago now.

What about some other things?

So, in terms of, for example, communication about science, finding out what's going on in science.

Well, there were already weekly journals that I used to look at 50 years ago, and the, gosh, I'm realizing it really is, this year, it really is 50 years since I was, first employed to do science as a 16-year-old high school kid, but… but that's,

at a government lab in England doing theoretical physics. But, you know, so at that time, you know, I was using computers to do science. I was looking at the weekly, kind of, what's going on in physics kinds of things. There were physics preprints that were getting sent around.

as some things that would arrive, and they'd get displayed in a library, or you'd get them yourself personally. There wasn't, at that time.

1976, there was just starting to be some email. I used to use a computer on the ARPANET that had at least local email, and actually also the computers at this government lab, place called the Rutherford Lab that I worked at, they also had some kind of email-like system.

the,

When it came to doing things like searching for stuff, searching for scientific results and so on, turns out that was already possible 50 years ago. There were databases that collected the abstracts of papers, databases like INSPEC and so on, that collected the abstracts of papers.

And you could use online systems, there was one called Dialog, that you would dial up to, and although, actually, at the government lab that I worked at, that… those computers were connected to the ARPANET.

So it wasn't even a question of dialing your computer up to a phone line to connect to the ARPANET, it was… the computers were all connected to it.

So you could go and do searches, at least in the abstracts of scientific papers. Also, if you're a little bit more, kind of, energetic, you go to the library and go to the Science Citation Index.

Which is a… was a printed index of all the papers that referenced a given paper. And it also had a thing, a so-called permuterm Index, where it actually would just take titles of things and kind of index by all the successive words in the title.

Lots of, kind of printed on thin paper, thick books, but you could go through and you could find kind of by hand searching, you could find papers that were about particular topics. So those things, you know, it's gotten a lot more streamlined in recent times, but it's still the same basic story.

Now, in terms of… of,

other kinds of things about, sort of, doing science and presenting science. Well, in those days, you'd write papers, you'd send out preprints, you'd send the papers into journals, et cetera, et cetera, et cetera.

One thing that was a little bit different.

is the papers were typed on a typewriter, and most people didn't type their own papers. They would handwrite them, I did that too, handwrite them, and give them to somebody to type them, and then they would show up typed, and then the copies would get made, get sent into journals, things like this. That was a transition that happened,

In the 1980s, mostly, is people starting to,

use to type their papers themselves. And back in the 1970s, most people who were doing science didn't know how to type. And that was a big competitive advantage for people like me, who actually could type fast. That, decayed rapidly over the years.

Well, of course, the other thing that's changed in 50 years… well, there are things that have changed about the social dynamics of science. There are things that have changed about, the… about the tools that one has available, particularly computational tools, and I guess I've spent a significant part of my life building those tools, but, to talk about, kind of, the more of the ambience of doing science. One thing to realize

50 years ago, being, you know, a physics professor or something was a tippity-top profession.

It was kind of a, kind of that's… that's a high-aspiration type profession.

I would say that the professoring business has definitely degraded in its tippity-topness over the years, possibly because it's sort of expanded a lot, and it's not quite as elitist as it used to be, for better or worse.

I think also, back in those days.

kind of being an entrepreneur or something was a grimy business, and not something that was thought to be kind of a… somehow, at least in the intellectual world, that wasn't thought to be kind of a path that one should follow. I kind of…

Broke from that tradition for myself.

And now it's certainly a much more kind of the being a tech entrepreneur at times has been seen to be, you know, a tippity-top kind of thing to do. It kind of goes up and down as the world changes, but that's a thing that, you know, the idea that there could be a serious entanglement Of, kind of, the entrepreneurial world with the academic intellectual world has definitely changed.

Another thing that has changed, I suppose, is, the…

the kind of, sort of, extra communication that one can have about science, whether it's through social media posts or live streams, other kinds of engagement. I mean, the idea of you go, you give a talk about something.

that's been the same forever and ever. I don't think that's changed, you know, occasionally I'll give talks at universities, and I realized just recently I happened to be giving a talk in Oxford, as it turned out, and I realized I hadn't given a talk in the physics department in Oxford since 1977. And I realized, by golly, it's…

just the same kind of schtick. The only difference is, instead of using overhead projector transparencies, I'm connecting my computer to a projector. But other than that, and I'd like to think the content, at least of things I've said, are rather different, but other than that, it's the same kind of dynamic of, you know, you go, you give the talk, people show up, you go out to dinner afterwards type thing.

in terms of the kind of social dynamic of science, I think, we were talking about this a couple of livestreams ago, something that's changed is a lot more collaboration in science.

between a lot more… when you write a paper, it's like there are many authors on that paper. It always used to be, or most of the time, in the past.

these… much of the time, at least in a field like physics, for example, there were many solo author papers, for example, in theoretical physics. That's much less true today. I don't really know what the cause of that is, whether people are actually working differently, or whether they are,

They're doing things that are, whether it's just a convention for, sort of, whose name goes on the paper.

I will say, you know, you can say, well, in today's world, you can do remote collaboration. That's a new thing. Well, it's not really a new thing. I did quite a bit of remote collaboration back in the 1970s, even though in those days, it was mostly you'd write letters and send them back and forth.

And there weren't even very many phone calls, actually, in those, in that, in those collaborations.

I would say… so we're… we're talking about, sort of, where's… what's science going to be like in 50 years, and I'm sort of talking about what hasn't changed in the last 50 years. And I would say that the, perhaps…

The level of cynicism about the processes of the scientific establishment has greatly increased, and probably deservedly so, over that period of time.

I mean, in those days, there was sort of a belief that, you know, peer review was kind of a worthwhile, sensible kind of dynamic, whereas today, there is deservedly deep, deep cynicism about what it really means, and whether it's achieving sort of progress in science, or whether it's just kind of a form of anonymous corruption, so to speak, in the dynamics of science.

I would say that the, so…

You know, the one thing that I would say, for many people has

changed, well, actually a couple of things to say. The big thing that's changed is ubiquitous computing. Everything is on computers, and

For me, that hasn't changed that much in 50 years, because I was doing things on computers 50 years ago. But for most people who do science, that has changed. For example, in theoretical science, in the kind of more mathematically-oriented sciences, I would say there was kind of a step change in 1988,

when we released Mathematica.

Before that time, people who were doing theoretical science, it's like you're working out this algebraic computation, it's like you got out your pencil and your pad of paper, and you were writing stuff by hand. After that time, there was a pretty rapid transition to, well, you use a computer to do it.

And I think the idea that, that that notion

that you can kind of use a computer to do things is sort of everything is now on computers. I mean, it used to be the case people would record information on index cards. Now it would be always in a computer, ever since the 1980s, probably, in a computer.

When it comes to experimental science, again, you know, gone are the lab notebooks where people would handwrite observations. In what's come in, instead, is everything is on computers.

Now, the fact is, lab automation has not progressed quite as rapidly as one might hope. There's still a lot of instruments where you're still pushing buttons on the front panel.

To make it do things, even though it might be able to dump its data into a computer.

I think, it's,

It's a question, but nevertheless, the types of data, the volumes of data which one can accumulate because computers are in the loop is much higher
Already back in the 70s, in fields like particle physics, there was a substantial amount of data coming in, and it was on computers. But, you know, there has been a progressive change that it's gone from megabytes to gigabytes to terabytes to petabytes, and beyond of the kinds of data that get collected in serious scientific experiments.
So, kind of looking towards the future.
I think there are things that I've thought should happen for 40 years, and they haven't happened, and maybe they finally will.
So, I'll give you some examples. So, one is kind of using computation as part of the presentation of science.
When you read a paper today, it's like the person says, yakety-yakety-yak, and if you want to reproduce what the person did, you're pretty much stuck kind of trying to understand, what do they do, let me write my own program to do that, etc.
There have been a bunch of initiatives to get scientists to sort of include raw data, include things that let their science be reproduced.
those efforts have been quite unsuccessful. People, will… people are sometimes reluctant to share that kind of inside information. If they share it, they share it rather badly in systems that are very hard to actually get the data from. It's… it's a kind of a sorry picture.
It's, for me, for example, for many years now, I have tried to make everything I do sort of perfectly reproducible and perfectly kind of computationally transparent.
And so, you know, if you read things that I write, they'll have all kinds of pictures in them, you can click the picture, you'll get a piece of orphan language code, you can run it, and you better get that picture again. And that's…
An important piece
in kind of both the reproducibility and the ability to follow on from the science that I've done, is this idea that you can… the code is right there. You can just take it and go and build on what's been done in the science. Now, the truth is, that's an idea that I first started talking about, oh, 45 years ago.
And the fact that it hasn't happened to date in a serious way is kind of a disappointment in both the social dynamics of science. It's no longer a technological issue at all.
I mean, it's a… it is… what we've built with Wolfman Language is kind of the best instance of being able to do that.
Both in terms of what we can present, what we can… how we can be useful in the science.
Also, the fact that we have this kind of unified language so that the kinds of things you want to put in your science paper, they'll be just a small piece of Wolfram language code, and it's all defined. It's not like, oh, and it goes to use this library that might have been discontinued, and might have been changed, and so on over here. It's all kind of a unified thing. The other thing we've achieved, and we're talking about long timescales.
is we have essentially perfect compatibility of orphan language from when it was first released in 1988 to today. So I will routinely take pieces of orphan language code that I wrote in 1990, for example, and fire them up today, and they'll just work.
And same with the notebooks that we have to… that represent, kind of, the workflow of using that.
So, I think one of the things that I'd like to see happen, and maybe eventually it will, is kind of this reproducibility and transparency of science, where sort of every picture has code behind it.

And when the picture involves data, that that data is in a… is in a kind of usable form. So we built our Wolfram Data Repository quite a few years ago now, which is a place where you can put data where it has been put in a form that is standardized enough

that you can just pick it up and use it in another piece of code. So quite a lot of things that I've written that involve data, the raw data is in the data repository, the code that you get in the thing I've written references a resource data object from the data repository so that that code… so that that data can just be pulled in.

Now, you know, making data fully reusable is not trivial. I kind of defined years ago kind of 10 levels of making data computable.

From just, it's in digital form in some way, to various kinds of standardization of entities that appear in it as standardized, kind of the way that missing data is represented is standardized.

to, you know, dates and times are standardized, all those kinds of things. In… we've built a bunch of tools, actually, we're building some more tools now, to help people kind of take, sort of, the unwashed data and put it into a somewhat reusable, readily computable form.

LLMs and so on help a bit in doing this. In our data curation efforts, it's tended to be kind of a factor of two speed-up in what we can do. It's not like you just feed it to the LLM and it does a perfect job, it doesn't, but it can kind of get around certain kinds of things that weren't there.

So I think that another piece to this whole story is having things where the… everything you're doing is kind of reproducible, computable, and so on.

Another thing is kind of the way that you communicate science in terms of how you actually describe what's going on.

And I would say that there has been, there hasn't been much evolution in the style of scientific papers in the last 50 years, disappointingly so. Because back when scientific papers were first a thing in the 1600s.

they read a lot like the blog posts of today. They were pretty chatty, they would have things like, you know, and I was studying this goat, and it ran away type thing.

Whereas in today's world, that would be turned into, you know, I was studying, I don't know the Latin name, for a goat. You know, we were… the specimen of such and such thing became unavailable in, you know, whatever it was, which means the goat somehow escaped and ran away, or whatever. Back in the 1600s, they would have basically just said the goat ran away. Today, they would have to fancify to make it sound more kind of professional and so on, which tends to lose a lot of information.

And I would say that different fields of science have evolved differently. I would say there's a trend in the last, probably, 5 to 10 years of, particularly, I would say, in the life sciences, more push

towards saying, what's the point of this paper, having a section about the significance of the paper. That's a good thing to have happened, because otherwise what can happen is that the paper reads kind of like a legal filing of a patent or something. It's like, this, this, this, this, this, and this, and it's all very formalized.

And it's like you're just sort of scratching your head about what's the point of this?

And sort of the idea of kind of this… this stilted kind of academic writing, that was a thing 50 years ago, it's a thing today. There have been some things that sort of break out of that. I would say on the web, kind of started, I think, with things like our Math World website.

which was then copied by Wikipedia and so on, as kind of a more direct way to sort of express things. In the things I've tried to write, I've tried to sort of express things directly in pretty much the things I've written, starting from about the, well, starting from sometime in the 1980s.

I mean, before that time, I too was writing kind of the stilted, academic paper type style. I kind of changed that, and by the time I was writing, well, for example, documentation for Waltham Language, that was in a much more direct style.

So, you know, it's like the actual language that we speak has changed and evolved over 50 years, so, to some extent, has scientific writing and scientific communication.

But then there's the question of what is the formalism? How are you actually describing what you're doing in science? You know, if you look back.

to, let's say, oh, I don't know, even the 1500s, for example, or 1400s, you'll find people were not really formalizing. Everything was in words.

And then along came things like algebra and so on. More formalization of mathematics. By the 1600s, there were lots of the notations of mathematics that we have today were already there, and so people were able to start kind of writing things in this more formal, structured way that mathematics allows you to do, and mathematical notation allows you to do.

So, we have that kind of mathematical structure now. A lot of things are still quite narrative. We also have had, for a long time, kind of, computer pseudocode.

What we now have, and what I've spent lots of effort on, is a kind of a computational language that allows you to represent, in a kind of standardized way, lots of kinds of things in the world.

So if you're trying to describe, kind of, the geolocations of different species of penguins or something, there is a canonical representation for those penguin species, there's a canonical representation for the geolocations.

These are things that become kind of standardized and computable. That's what we've tried to achieve over the last four decades with Waltham Language. And kind of expressing yourself in computational language, not only

not only, kind of, data, but also computations, algorithms, what you are actually doing expressed in computational language. So I think a big thing is the, sort of, the rise, and it's happened to some extent, not as fast as one would like.

The rise of using computational language as part of the thing that carries the message of scientific communication.

And I think we've had, sort of, plain English, we've had diagrams, we've had math formulas, most…

when you're using traditional programming languages, putting big blobs of low-level code is not a thing intended for humans to read. With Wolfram language, with our computational language. It is, if you do a good job writing the language, it can be very succinct, and it is something that humans can read and learn from, just like math formalism is something that humans can read and learn from. So…

That's something which I would hope

that there will be increasing, kind of, use of computational language as part of scientific communication. So you're not… so the things you can say become formal and precise and reusable in the same way that… in the way that one can do that computationally, so to speak.

So, another thing I might say about science communication is things like livestreams. How much can you see about the actual ongoing doing of science?

So far, you can see almost nothing about the ongoing doing of science. You know, those… all those meetings and discussions between scientists, those are kind of behind closed doors. I have done a fair amount in the last 5-6 years.

of kind of livestreamed, science, a little bit less in the last few years, and I really need to start it again some more, of actually showing, you know, of actually sort of letting the world look in on the actual dynamics of science being done.

I think there's a very interesting thing, it's very kind of educational to see how that works. It's something I tend to think I'm prepared to do it, because

I feel confident, and I'm not too embarrassed when I say stupid things, but my academic friends have mostly… well, I believe there's basically zero of that that's gone on in the world to date. That's something that might change.

The sort of increasing transparency of how science gets created.

I think, Another kind of… other questions about, sort of, how science gets done.

I mean, a lot of the systems of doing science are a little bit under attack as a result of LLMs and AI. You know, everybody can generate a thing that kind of looks like a science paper. You just tell an LLM some prompt, it'll make a science paper-like thing.

anybody can use LLMs to write, reviews of papers if they want to. These are things that are sort of somewhat poisoning

the system that had existed. I mean, it used to be the case that, you know, preprint servers and so on could just accept anything anybody sent them, because nobody would bother to send them anything unless it was an actual science paper that made some sense. But now you can kind of churn out

AI-generated stuff with incredible ease. And that's kind of making some aspects of kind of the machinery of science creak

quite badly. I think that the, in terms of, sort of, the use of AI and the doing of science.

I think one of the things is kind of on this continuum of how do you find information about what's been done before, from kind of… you'd look through the Science Citation Index, you'd use an abstracting service online, to you can do a web search.

Well, LLMs and AI provide another level, which is you can kind of search thematically for things. And that's an interesting possibility that I guess I've found myself using increasingly.

And something where you can kind of drill in, you know, this specific question.

And it's not a question of keyword search, read, try and find it. It's more the LLM will do the drilling for you. And that's a thing that I think is, you know, will make certain things more efficient. It's, in a sense, one of the things you hope for in science is that the whole sort of corpus of all those millions of papers that have been published, that it sort of is additive in some sense, and that it's not the case that most of the time you're kind of going back

And I don't really know whether anybody did this before, I have to do it again for myself, type thing. So as these kinds of thematic searching capabilities get better, so one can expect sort of more of, yes, you can build on what's already been done to a greater extent.

I would say that, another thing

to mention, actually, about the… about the doing of science is… is, well, actually, a couple of things. Let's talk, first of all, about the content of the science, and then we'll talk about, sort of, who gets to do the science.

So, in terms of the content of the science.

every… you know, there are these bursts of kind of activity that come about when new methodology exists, new methodologies exist for doing particular areas of science. And we can expect that over the next 50 years.

There'll be a bunch of those new methodologies, whether it's experimental methods, whether it's theoretical ideas, and so on.

I think that there are certain places, I mean, I've mentioned this,
The kind of excitement of my own efforts in the last 5 years or so that really do seem to be kind of a new level of kind of capability, methodology for theoretical science, kind of building on sort of the last half century of kind of development of sort of the computational paradigm for doing science.
I think that's something where we're seeing sort of a burst of activity in a bunch of areas. I would say that one of the things that is sort of a methodological change, I started doing this at the beginning of the 80s, but it's something which is sort of increasingly popular, is kind of computer experiments.
And experiments in, sort of, minimal, abstract experiments in the computational universe, so to speak.
and the things you can discover from that. It's sort of a different approach to science.
Where instead of going from the phenomena we observe and measure and drilling down and trying to sort of reverse engineer how the world works, we're sort of going up
More from the underlying phenomena of what is computationally possible, and exploring that, the field of ruliology that explores, kind of, the abstract, basic science of what happens when you follow rules and so on.
How that works, building that up, rather than reduction… reductionistically going down.
And it's somewhat more like the way that pure mathematics has operated, but I think that rumleology is a richer and there's more pieces to explore there, and it's also something where there's kind of the… there's kind of a big… a big collection of things to be discovered that can then be made use of, that can then be sort of pieces of theoretical science that are on the shelf
Ready to be applied when one needs them for technology, or for other areas of science, and so on.
I think that, the, sort of the computational paradigm as a way of thinking about the world
is something that I've been pushing now for 45 years or so, but it's something that we're seeing more and more, kind of, get into science. I mean, one of the transitions that's happened over that period of time
is the transition from, you're making a model in science, let's make it be, in terms of mathematical equations, to let's make it be in terms of computational rules. That's something that's sort of silently happened over the last three decades or so, a transition from math, it's all math.
To, it's all sort of computationally defined.
I think there are also questions about different fields of science that have had a certain paradigm for operating for a long time, and that paradigm might change.
You know, in something like physics, there's been a strong theoretical tradition, as well as the experimental tradition. In something like biology, it's been all experiment, no sort of systematic theory, or very little systematic theory. That might change.
In fields like, in fields of social science and so on, same kind of story.
how that will change, whether that will change, I think on the timescale of order 50 years, we can expect that that will change quite a bit, that there will be sort of a substantial branch of biology that is really the theoretical side of things that makes more general statements than finding all the experimental details.
And I think that's, and in mathematics, kind of the opposite way around, I think experimental mathematics, that's much closer to ruliology.

exploring mathematical kinds of ideas, but by doing computer experiments, I think that will become increasingly popular. Again, it's been something that I've done for 50 years now, but it seems it's only gradually kind of gaining in popularity, so to speak.

In terms of, of,

Well, so, as I'm… as I'm saying, the… kind of the dynamics of, sort of, the scientific establishment, the kinds of things that get done, the units of scientific activity, these have all been pretty constant over the last 50 years with sort of incremental change. I don't really expect that to change.

Now, there's a question of, sort of, who gets to do science, and how do they get to do it?

in the last… there's a question of, for example, doing basic science. Why should one do that? It's payback is 100 years in the future, where payback… if by payback you mean something of, kind of, technological, medical, whatever significance.

And so, it's been an interesting question. How does that get done? Where does it get done?

For a while, kind of, it's been, sort of, universities and governments are the main places where, sort of, you get to do, kind of, this, sort of, long-range intellectual work.

I think that it's been a long time question whether there are other models, and in my own life, I've obviously followed a somewhat different model of, sort of, doing science as a spin-off, in a sense, from doing technology, although the science also informs the technology and so on.

But it's… it's a… there's sort of… I've wondered for a long time whether there are sort of fundamentally different dynamics of how basic science should be supported, who gets to do it, the extent to which most people who do it are professionals.

versus not. I mean, back in… if you go back a few hundred years, the majority… there were a few professional scientists, but the majority of people who were doing science were not doing it for a living.

And, they were doing it, they made their living some other way.

And I think it's been surprising to me the extent to which there isn't much of that happening today, that most people who are doing science are doing it professionally. I'm not quite sure why that is, because I know for myself.

you know, doing it, quotes as a hobby, I think it's fair to say that I managed to produce a lot of fairly substantial science

even though I'm, quote, doing it as a hobby, CEOing a pretty active company and so on, that you might think would mean that I couldn't possibly have, you know, science as a hobby, but at least for me, it's worked well.

And maybe more of that will be seen in the future. I think the efforts of, sort of, involve everybody in science, kind of the citizen science approach.

It's had some successes. I don't think it's been a resounding kind of, you know, transforming the face of science kind of idea. I think there are fields of science where it is more realistic

for a broader range of people to make contributions. You know, that's been… that was true at times in natural history. It's true today in natural history. You know, if you're looking at bird observations or something, the vast mass of those are made by… by amateurs, so to speak.

And are being aggregated in a way that people who are trying to do systematic things can use.

I think that ruliology is a place, so the study of simple programs and so on, is a place where one could imagine a very substantial kind of amateur set of contributions, where things get discovered that are sort of point things that are discovered, that's kind of on the level of, you know, I saw this lesser-crested bird in this particular place. It's a computational kind of thing.

Like that. And in the end, if those observations are aggregated in the right way, the result can be something that is useful.

I think that, in terms of, what else to say about the doing of science, I mean, I think in terms of experimental science.

There will be more and more automation.

of how data is taken. That's been very slow to happen. Surprisingly slow to happen. And that's a consequence of the details of the kind of social dynamics of science, and the way people are trained.

And the instrument makers, you know, really wanting to have, kind of, instruments be islands and so on, but it's coming.

And I think that it will be, you know, the idea of going into a biolab and seeing lots of people with pipettes, you know, measuring out little pieces of fluid, that really is going to be kind of something that looks very ancient.

In the future, just as seeing that person, you know, writing out that integral by hand looks very ancient to us post-1988.

Well, anyway, so that's a… that's a few thoughts on that.

Let's see… Well, there's a question here… Several questions.

Justine asks, will future scientific knowledge be stored as papers, simulations, executable models, or something totally new?

I think a lot of it should be stored as computational essays.

things that have narrative tests that humans understand, together with computational language that humans understand, and computational language that's also understood by computers. So, it really is something that is where the computational language is carrying meaning

Both for humans and for automatic execution on computers. And I think there's a sort of a question of what's the unit of scientific communication?

And that's changed a bit over time, and it's still kind of a mess. I mean, a paper is a certain unit of scientific communication. A book is a certain unit of scientific communication.

a kind of… it has often been difficult to have smaller units of scientific communication than papers. I mean, sometimes there'll be tweets these days, sometimes there'll be… somebody makes an observation and puts it in a big database.

I would say that having that sort of sub-paper communication is something that, for example, in ruralology, one can imagine that being pretty useful, and I don't think we fully invented the methodology for being able to store things that aren't just

Yet another bird observation to be put in the database, and isn't at the level of, oh, you have to explain things in words.

Now, there's also the question of, sort of, the book versus paper, etc, and there's a question of, sort of, how incremental is science? You know, in my own efforts, a lot of things I've done

I can't seem to write things that are less than 100 pages long, because it turns out there's quite a bit to say, and there's sort of… you can… you could break that up into many pieces, but then each piece would have its own separate introduction, and you wouldn't really be able to see the big picture.

I would say that that way of doing science is the exception rather than the rule. The vast majority of science that gets done is quite incremental, and it really can say, you know, the joke of science is always, you know, that all papers begin, you know, Smith & Jones recently did this, and now we're going to do that.

And that's sort of explaining that the thing you're writing fits right into the flow of what's been done.
And I would say that that… the, you know, one would like it to be the case that
Perhaps one would like it to be the case that there's sort of
pieces of science where there's really big, you know, intellectual structures that get built. I think there's a limited rate at which those can get built, and those are the things that get communicated in books and so on.
I would say that the idea of,
For example, it is useful in communicating science that you can have a thing with… of a delimited size, a paper, a book, whatever else. The website with tentacles everywhere
I think it tends to be a less satisfactory form of communication for sort of a new science thing, because it's kind of like, well, I know this piece, but I don't know that piece, and I don't really get the whole picture, and so on. So I think this very traditional kind of view of this unit of scientific communication is not so far off.
Although, having something for smaller-scale communication would be a good thing. Now, it's also the case that when it comes to the development of tools.
You know, computational tools and so on. There has been sort of an up-and-down story of whether people consider that to be sort of appropriate academic activity.
I don't know the current state of that, but for a long time, if you wrote some computational tool, you couldn't sort of put it on your resume and have it be something that counted towards your academic tenure. It had to be sort of papers and peer-reviewed journals type thing. I don't know if that's changed. It certainly should, and it would be… because that's a very valid contribution to kind of the tower of knowledge that gets built.
And if it isn't recognized as such, it's a mistake in the recognition, not in the mistake in the doing of that activity.
Let's see… Gary asks, how do we decide what science is worth pursuing when almost anything becomes possible?
Well, I mean, that could be said about many things. You know, you could write any piece of music by putting the notes together. Which piece do you want to write?
That's a matter of human choice. Science… what drives science?
Science is sometimes driven by technology, sometimes driven by things like medicine, sometimes.
And I have to say, this is unfortunately not the most common. It's driven by purely, kind of, aspirational, kind of,
questions. I mean, it's… it's driven by sort of a desire to understand the world, something sort of very pure and almost aesthetic, about kind of, you know, we are… what… what…
what kind of a thing can we build? I mean, just as it's satisfying for us humans to have built cathedrals or something, so I think it's satisfying for us humans to have built, kind of, the giant edifice that is modern mathematics, for example.
But that's something, you know, realistically, that is probably a minority, kind of feeling, that those things are, are sort of, are things worth aspiring to, so to speak.
But I think there's a certain sense of that, a certain aesthetic sense, a certain sense of, sort of intellectual closure that can be reached through those kinds of things.
That's something that, for some set of people, is really a driving force.
And I think, you know, one can ask the question, who wants to understand fundamentally how things work?

I know I do. I don't know how many other people want to. I think that if you say to people, would you like to figure out the fundamental theory of the universe? Some people would say, hey, that's really cool, we really want to do that. Other people would kind of shrug and say, yeah, I don't know, you know.

Doesn't really, you know, let me, eat better or whatever else.

it's… people will say, what's the point? What do we get from doing that? And the answer is, we get something that kind of, in a… in a kind of almost theological way, kind of takes us sort of closer to kind of the… the ultimate nature of reality and so on, and the ultimate nature of sort of our place in the universe, but in terms of, you know, does it

put, does it put better food on my plate? The answer is no.

And so that, that's a, it's kind of a, you know, to what extent do you pursue that? I mean, that's been a question for many fields. I mean, you could ask that question for art. Why do we pursue art?

you know, do we pursue art? I think we pursue art because it is… well, some people would say we pursue art because it is ennobling of us to kind of experience good art.

Other people in a more modernistic point of view will say, you know, it makes us think in unexpected ways, or whatever else, or sometimes it's not really clear what the point is supposed to be.

I think that the same can be said of science. Pure science, basic science. It's kind of an activity that's a little bit like those kinds of aspirational activities in art and so on, and to have

Made progress in basic science, and to have understood more about how the universe works and so on, feels like an ennobling activity.

Let's see…

David asks, from your perspective, is this a good way to learn programming so the knowledge is stored in long-term memory and remains usable years later, instead of being forgotten after, sort of, short-term practice?

How did you personally figure out how you learn best when mastering complex subjects? I think I've talked about that second thing quite a few times, but I'll say a few things here.

You know, I always learn best by trying to actually figure things out.

by asking questions, often quite foundational questions, about different fields, and saying, what do I need to figure out to be able to make progress on this question?

And for me, you know, I think

using our computational language. I'm… when I'm thinking, I'm typing notebooks with computational language and some English text, and that's my kind of medium for thinking.

And it has the great feature that I can go back to notebooks that I created in, you know, 1989. I can, you know, fire them up again, and they'll still run.

And that's a very useful thing. I mean, it used to be the case, and people don't do this really anymore, but, you know, when people go through school, they would take notes, and the notes that they took would become the things that they would refer to

for the rest of their lives. I mean, that was true, for example, particularly things like ciphering books used in American education in the… a century ago, or something, where people would write out, kind of, math problems

And then when they were, you know, being farmers, trying to figure out, you know, this percentage of that plus that, plus whatever, they would go back and look at their ciphering book that they wrote when they were 11 years old, or whatever, to see what to do. Now, of course, textbooks came in, the web came in.

And it is less necessary for sort of traditional knowledge to have to go back to your own notes. Of course, if you were figuring out things for yourself.

then your own notes are what you have. I mean, for me, one of the things that I tend to do is that because I kind of make a lot of stuff, publish a lot of stuff.

it's, it's convenient for me that things that I knew at some time, I probably put out there in the world, and I can find them because they're out there in published resources. That's particularly true of my big book, A New Kind of Science, which came out in 2002. Big, thick, 1200-page book.

that contained a very dense whey, just an awful lot of stuff that I had figured out.

over the 20-something years before that time. And so, for me today, probably on most days, I'll go look something up in New Kind of Science, usually abbreviated NKS.

And that's very convenient for me. It happens to be something that's sort of widely available in the world, but it could as well be something that I'd stored carefully in my file system and was just looking at for myself.

certainly the source material, which we're finally putting on the web after all these years, are finally sort of cleaning it up to the point where it's generally usable for all of NKS and so on.

Let's see…

You know, I have to say, one of the things that's sort of a new game in town that I haven't really tried is learning, interacting with an LLM.

I mean, I think…

It's, for example, if I was, you know, in school trying to memorize things today, I would get an LLM to ask me questions about the things I'm trying to memorize.

And, you know, I would… I would prompt the LLM with my… with the thing I'm supposed to remember, and say, ask me questions about this. I suspect that would work well. I'm not… haven't been in that business for

more than 50 years now. I left high school, yeah, 50 years ago now.

So,

And haven't been in the, doing tests and memorizing things world ever since then, for better or worse.

The, I think,

in, but I think that's a… that's kind of a new game in town. I mean, in fact, we've been working quite a lot on, sort of, AI tools for education, and it's the interesting question, whether… if I wanted to learn a subject, we… we're not… we've… we've built

a course, an AI-based course, for algebra. I'm not in the business of learning algebra right now. If we'd built such a course.

for things I'm trying to learn, like various areas of economics and philosophy and so on, then I might try that as a way to actually learn things. I haven't had the experience of doing it yet. I don't know how well that will work. I know that for myself.

it's once I reach a certain level of knowledge in a particular field, it's super useful for me to talk to people who know that field really well.

And sometimes that's a double-edged sword, because you find people who are real experts in a field, and they'll have their own particular schtick about that field, and sometimes it's hard to get the big picture of the field.

And for me, I have to already have some of the big picture before I go talk to this particular expert.

And then, because I have some view of the big picture, I can know what the things, how the things that they're telling me fit in to other things out there in the field. What I tend to find is that some fraction of the way through learning a field, I'll go from… every time I talk to somebody, they tell me things I didn't know.
to when I talk to people, the things they tell me are things that fit into things I already knew. And that's a… it's always a good feeling. I've done that many, many times with many different fields, and it kind of… you feel like you're drowning when, you know, you keep on learning new things, and, you know, this new thing is coming at you, and it's… it's kind of like I'm… I'm not going to make it for this field, but then at some point.
you… it all comes to closure, and the things you hear start repeating and repeating and repeating, and then it's like, yup, I've got my arms around this particular field. And that's the time when I feel like some confidence in making progress in that field, and so on.
Let's see…
Let's see, Serbin is asking, the context of LLMs, how important is it to provide access to curated databases or labels of answers on particular topics? Is this something that our company aims for? Well, we have put out a great deal of, kind of, information about the world, and algorithms, and so on and so on and so on, and, I think we've been, lLMs have, have eaten that all up, and we've helped them do that in many cases.
I think, that… The level of structuring that one needs to do is something that's a question.
I mean, at a time, you know, in practical terms, you know, it's like you've got a document. Do you convert it into Markdown? Yeah, we're doing that kind of thing. Is that really going to be necessary? Probably not. Probably, if you just feed the LLMs, kind of, any old, sort of, documents in any format, they're going to be able to untangle that.
That's a… that's a matter of, sort of, the harnesses around LLMs, that's something which I think is sort of readily… is readily coming, and it's not really the… the effort to let me prepare it for the LLM is not going to be so important. Let me have it be there in the first place, that's important.
actually having the material exposed in some way, creating the material to begin with, that's important. I mean, I know, for example, these livestreams that I do generate a lot of transcripts that contain, you know, answers to things.
which LLMs can then go and use. But if I hadn't exposed those answers, and often when you guys ask me all kinds of questions where I've never thought about that question before, and I have to try and sort of come up with the answer here.
those answers are sort of created for the first time here, and then they can be scooped up by LLMs to use for what they do.
Justine asks, how do we prevent students from outsourcing all thinking to AI without banning it outright? Should schools explicitly teach students how to work with LLMs, not just restrict them? Look, there's a really simple principle. If you're going to have a tool that you're going to be able to use for the rest of your life, you might as well learn how to use it.
And if… if everybody can stop working and just turn it all over to AIs.
then, well, I don't think that's going to happen, for many reasons. But the kinds of work that get done change a bit.
The kinds of things, you know, it's no longer useful to learn that sort of… well, to learn that kind of

particular low-level thing, because, yeah, that can be done automatically. It's no longer useful, you know, in my life, the transition happened between when, you know, any real programmer would learn a single language.

to then any real programmer will learn C.

to, you know, we've got now sort of high-level computational language, well, from language, for example, where you can just learn that. Lots of kinds of things you want to do, you just have to learn that. You don't have to go to the lower levels and learn all those kinds of things. And you might have said, let's outlaw, you know, we really need people to learn a similar language, let's outlaw any higher-level language, because it makes it too easy.

I think that that's kind of a mistake.

That's… it's kind of a bad way to think about it. Now, the question is, what is it that people really need to do? They need to be doing things that aren't just, oh, you were asked a question, there's a standard answer, can you trot out the answer? It's more, you know, how do you think about something? What questions are worth asking? All these kinds of things. These are things which, the more that one can teach people to think.

as part of education, I think the more valuable that will be, and the kinds of thinking that one does, there's a lot of thinking that is just

for what we choose to do as humans, which, other than the AIs reflecting what we've already chosen.

it's something that we get to do for ourselves, and that's sort of what one should focus on in education. I mean, there's a lot, you know, learning the mechanics of things that have been automated by mathematical and modern language for 38 years.

is not necessarily useful. It's, learning, you know, how to do some algebraic computation is… is not necessarily the point. I mean, that might be an interest in its own right, but the more important thing is, now that you can do that.

what can you build? What's the tower you can build that goes taller? And yes, there's a taller tower you can build with AI and so on, and, you know, now we have to figure out what do we teach that makes use of that taller tower?

And how do we teach, kind of, you know, AI criticism, so to speak? How do you tell… how do you read AI with the appropriate level of skepticism? How do you leverage what's possible there? That's my view of that.

And if you've got an exercise, I mean, this happened when, particularly when Wolf Alpha came out, there were exercises that people were doing in areas like math and chemistry and things like this, where it's just, look, I'm sorry, you can just do it. It just… it just… Wolf Alpha just does it. Now, did that mean nobody did those exercises anymore? No, it didn't. You know, there were people who were interested in doing those exercises, there were people who were never going to be interested in doing those exercises, and would find some way to get the answer without doing the exercise, and Wolfram Alpha was a much richer ways, a much richer way to do that, with a lot more, kind of, actual… you learn something because you're seeing things other than just

oh yeah, the guy last year was doing the same homework and left it for me, and now I can just write down the answer. You get to see the whole thing in a much more contextualized way, more of an alpha.

And so on.

Max is commenting, with LLM prompting, you have to be critical of every answer and play your own devil's advocate. Yes, indeed, and it's quite educational.

You know, it makes a very glib statement. This is true. You look at it, you think about it, you realize, no, it's actually total nonsense. I mean, it's… in a sense, it's a little disappointing in terms of, sort of, the social dynamic of it.

Because…

people confidently saying things that are complete nonsense is something that doesn't happen that much in the world. It's like people, one likes to think, and I think it's true, that people don't just sort of lie about things most of the time. Most of the time you have, you know.

kind of more or less the straight story. And with LLMs, it's kind of become more taken for granted that people will just talk nonsense, and… or LLMs, at least, will just talk nonsense. It kind of socializes the idea of talking nonsense, which is not a good thing.

It certainly doesn't make the world an easier place to operate effectively in.

Let's see…

its comments. I've seen so many examples of companies trying to use LLMs as search engines, and never seen it turn out well.

Is it possible for LLMs to reliably retrieve factual information, or is it a fool's errand? Well, I think that it's absolutely useful to use an LLM to kind of do thematic searching. Where are you looking?

What is there at the end?

It's something where you want to pull out the thing that was actually there at the end.

the piece of database, the computation, whatever else. LLMs are a great kind of linguistic interface to things. They're a great way to have kind of this fuzzy linguistic interface. I mean, it's kind of like graphical user interfaces.

you might have said, I want precise control over what's going into my computer, and so on. I don't want to just move a window around on the screen, and it's kind of, you know, I'm moving this pixel from here to here, and I… it's not precisely right. How can I live with that?

But in fact, graphical user interfaces are very useful, even though, in some sense, aspects of them are not as precise as if you were typing every character to say exactly what you want to have happen. And I think the same is true with a linguistic interface that's provided by LLMs, and it's really just a question of sort of getting used to using that in an effective way.

I think that ultimately, it's a, it's a way of accessing tools

and the tools can be precise. I mean, we've… we built, kind of, the sort of biggest example of a computation and knowledge tool in Wolfram language that LLMs can access, and LLMs can sort of be the… a human tells the LLM roughly what they want.

the LLM kind of goes agentic and figures out, you know, kind of, it needs to get this, it needs to get that, and needs to get the other thing. That's actually a fairly shallow thing to try to do.

And then, when it really gets into depth of, oh, I want to compute this thing, I want to search all these rows of this database or whatever else, that's done by… by sort of hard tools that are not part of the kind of interface provided by the LLM.

Let's see… it's also asking, do you think there's an opportunity for development of an educational LLM, i.e. a dedicated system specifically designed to aid discovery and understanding, as opposed to simply providing an answer?

Yeah, it's an interesting question. I mean, is there a,

is there a kind of Socratic LLM that you can build that never answers your question, it just asks you other questions? It'll be an interesting experiment, it might work quite well.

I mean, I think that, you know, it's a little bit… people go a little bit wild on that kind of idea.

Like, people say, let's… you know, there's certain kinds of things in, sort of, philosophy and the way you think about yourself and the world and so on, that you probably can discover, you can educate in the Latin sense of to draw out, so to speak.
That you can imagine it's all in there, but we have to draw it out.
In some way, and that you can do sort of Socratically by just asking questions, and an LLM could probably be quite good at that. But there are other things where it's just, like.
There's no amount of drawing out that's gonna tell you what the distance to the nearest star is. unless you heard it somewhere, or know it somehow, that's not a piece of knowledge that is ex… that is… that an individual person can intrinsically know. I mean, back… it's sort of an interesting thing to realize. Back in antiquity, back in, you know, ancient Greek times. when, you know, when Socrates was plying his trade, most
Things were things that an unaided human could know.
And or could figure out.
Whereas today, there's enough depth of knowledge in both the, kind of, the amount of knowledge, and the level… the extent to which knowledge has been accumulated, not by direct human experience that anybody could have.
But by sending spacecraft out there, or by, you know, running computers for a long time, or whatever else.
And so it becomes something where the notion of education, where you're pulling out from what's already there, isn't quite as realistic, and you have to kind of explain things that are part of the corpus of human knowledge and so on, before you can pull things out. You have to push things in before you can pull things out.
But, so, you know, I think that's a… that's an interesting question. If you can give, you know, some pieces of information, and then you sort of try and pull out the links between them, I don't know. I think it's a reasonable thing to try.
Let's see…
There's a question here from Technical. What's a significant scientific discovery you'd like to see being made during this year?
Oh my.
It's always so difficult to know timescales for things.
I mean, I think that,
I mean, I can say for myself, personally,
there are a bunch of things that are happening around our physics project and so on, where I'd really like to know what some experimental implications of our model of physics
is, I'd really like to know, you know, is dark matter, in fact, space-time heat? And how could we know that? These kinds of things.
those are discoveries that might just get made one day. You know, we're primed, we're ready, there's just work to be done, and it might just happen. Hard to know. I think…
In the world at large.
there are plenty of areas where, sort of, there are sort of breakthrough discoveries that could be made. I mean, in biomedicine.
You know, there are sort of breakthrough discoveries, like, how do you reprogram cells to go back to being stem cells?
How do you clone mammals? You know, one that's certainly a coming attraction there is cryonics. How do you freeze humans and bring them back from being frozen type thing? My guess is that will be a point discovery. When will it happen? Who knows?

But that's the thing where, you know, you find the right cryoprotectant, you find the right, you know, high-pressure magnetic field, something or other thing, and it's just going to work.
Or maybe it's physically impossible, but I don't think it is. And it's certainly encouraged that, you know, there are fish that get frozen all winter, and so on. It's, but that's an example of something where there could be a sort of point discovery in that area.
I think in biomedicine, there are also more kind of, you know, that's a cool hack.
that will allow you to cure a bunch of diseases and so on. I can imagine those things being discovered. I think that there are things in… I mean, this is sort of the spectrum of different kinds of things that might happen.
You know, I think one of the sort of coming attractions is going to be sort of humanoid robotics that really works. There'll be a chat GPT-like moment for that, probably. I don't know when it will happen, but that's sort of another coming attraction.
there'll be things to do with nanotechnology, where it's kind of, can you build things at a molecular scale, or can you even build a universal constructor that operates at a molecular scale?
These are things that I think are sort of things that might be kind of point discoveries and so on.
I think, well, those are a few kinds of things, but it's, you know, would they happen this next year? No idea. It's really hard to tell those kinds of things. Like, if you'd asked, when will ChatGPT happen?
Nobody knew it could happen.
Nobody knew it would be a point thing that it happened, and certainly nobody knew it was going to happen in late 2022. That's just not something you can sort of pinpoint in thinking through how things develop in science.
there's a question here from Ricard. What do you think about the concept of technological singularity?
Oh, my.
I mean, I think, sort of, what is the concept? The concept is things happen faster and faster and faster, you know, things are growing, you know, going exponentially faster, and eventually it'll get to the point where things sort of get infinitely fast, and infinite things are discovered infinitely quickly.
Not gonna happen that way.
I know that from a very foundational theoretical point of view, this idea of computational irreducibility that I talk about a lot.
Even when you know the rules for a system, you can't work out, you can't kind of jump ahead and immediately work out what it will do. You have to just sort of follow each step and see what happens. One of the consequences of computational irreducibility is there's an infinite chain of discoveries that can be made.
You'll never be able to make the kind of the master discovery that discovers all discoveries.
You're still… and you're still stuck.
Being… doing things in our physical world, which itself is kind of executing this irreducible computation which corresponds to the passage of time.
So, I think the idea that, you know, the moment will come when it's just… it's all figured out, and we can just jump and say the answer is 42 and we're done.
That's just… for theoretical reasons, that's just never going to happen. That's not going to be the way it works. There will always be… there is an infinite frontier of invention and discovery to be made.

The question of whether we, for example, care about those discoveries, that's a different question. There are discoveries to be made about all the possible theorems of mathematics. Most of them we don't care about.
There are discoveries to be made about all kinds of weird materials that you could create. Some of them will be useful, some of them won't be.
It's, but I think that the idea that, sort of, things speed up to the point where it's infinite speed
So, you know, that's partly a consequence of saying, you know, the AI will figure out how to figure itself out, and it will spin itself up to infinite speed. It's just not going to happen.
That's what computational irreducibility, undecidability, the halting problem, Godel's theorem, they're all kind of saying the same thing, which amounts to, that's not going to happen. You can never have something which sort of infinitely jumps ahead infinitely quickly. It is not the nature of… of not the sort of… it's not a question of what's physically buildable, it's a question of what abstractly
As possible would abstractly make sense.
So, you know, it's, the,
So I think, kind of, the self-improving AI that's suddenly going to jump to infinitely improving itself, I just don't believe in that scenario.
I think, I find it kind of amusing that there are people I know who are, on the one hand, say, the world's getting exponentially faster.
And other people who say, I can't believe everything is happening so slowly, and that we have such, sort of, lame advances in technology. How can both of these things be true?
I think part of the point is that in well-developed fields, there's just more and more of the same, and there's more and more and more of the same. It's sort of exponentially larger numbers of papers, exponentially larger numbers of patents. I don't know if there are exponentially larger numbers of companies, I wouldn't be totally surprised.
But within a well-developed field, you can just pile on more incremental stuff, and it looks like more and more and more is happening. But the real progress gets made between existing fields, and those are things that are quite invisible if you're looking at measuring things by the particulars of some particular field.
And so, if you're measuring things by the particulars of that field, you're like… you're like, well, nothing's really happening, it's just incremental more of the same. But if you're looking just within that field, you say, it's more and more and more of that incremental stuff, it's exponentially increasing amounts of exponential stuff.
But then, as you look, sort of, the real progress is getting made, not in what you're looking at, because it's between these fields. And so that's why both of those points of view about things are speeding up exponentially and things are incredibly slow, they're both sort of right and they're both sort of wrong for this reason.
Let's see…
JH is asking, if an LLM were trained not to guess the next words, but to guess the previous words.
Could you give it an end state and go back from that? Is that harder? No, actually, there are many situations in which you want to do
kind of, you want to say things like, I've got this context, fill in what's in the middle, I've got what's here, what's likely to have come before it, and so on.

One of the things that I've done experiments on, so have other people, is kind of the planning LLM. You say, I've got this, you know, an LLM is just putting down, essentially, one token at a time, and then looking at what had come before. I mean, in earlier times.

there used to be this amusing thing that you could do with an LLM that would say it would produce some result, and then you would say, that's wrong, you know.

You know, look at what you just did.

And it would say, oh, I'm terribly sorry, I can see that what I just did was wrong. The reason it didn't know what it was doing was wrong was it was always just… always going out in front, it was always just putting down another word, and it was only after you say, take a look at what you wrote, look back.

that it had that whole thing in its context, and it could then say, oh, wait a minute, that doesn't make sense. The next thing I'm going to say is, that doesn't make sense.

So, there are…

there's sort of a question of could you, in an actual LLM right now, it's following a single path of, I'm going to say this, I'm going to say this, I'm going to say this. It's building a single future for itself.

You could imagine building a whole sort of quantum, kind of set of parallel paths of the future, and somehow pruning those parallel paths to say, I'm going to look at what… which of those paths lands in a place where I want it to land.

There's… there's various kinds of work that's been done on that that's closer to diffusion models and so on. It's not clear how that will work out.

Let's see… There's a question here…

Will we ever, from Davey, will we ever have cybernetically enhanced memory?

would this be a great leveler for people with otherwise similar intelligence? I mean, we already have cybernetically enhanced memory. You know, somebody like me has stored, you know, every message I've sent and received for

what is it now? 35… more than 35 years, and, you know, we… I can…

immediately recall things, oh yeah, I had that conversation with this person 20 years ago, and, you know, I wrote notes about it at the time, and now I can remember that. And it's very cybernetically enhanced, and it's very, very useful.

And, you know, I have a fairly decent memory as a pure human, but me with my computer as my sort of enhanced memory, I think we're pretty good at remembering things. So, I think it's already there. I mean, whether… if it's like, oh, I'm just gonna think.

and it's all happening inside my mind, and with neural links or whatever else, that's a different thing, but we're so close to that now, I don't think that's going to be much different. I mean, the fact that we can store all this information, I mean, you know, I really notice… it's very interesting to me, it's kind of like things before the web, things after the web.

Things where you could know a bunch of stuff about the world, and things where you would have a very hard time knowing that stuff about the world.

before digital… before phone cameras, and after phone cameras. It was… there was a time, you know, I was… I was… like, there are things where it's, you know, I wonder, is there a picture of that thing from the 1970s? And the answer is probably not. Nobody was running around taking pictures of things.

And, you know, even though many people saw the thing, nobody took a picture of it. Nobody came in with a camera and took that picture. Whereas today, because we always have, you know, we all have sort of cameras in our pocket, or even on our wrist, or whatever, or whatever it is.

It's, there's much more of a tendency to just take a picture when you feel like taking a picture. I think… I think I'm not modernized enough, actually, in that regard. I take fewer pictures than I probably should.

It's, but that's another sort of form of enhanced memory, and as it becomes easier to sort of decode what's in pictures, it becomes easier to sort of have that even more memory. I mean, I think it will be the case. It's an interesting question.

whether, if we had those little, you know, cameras on our glasses or something, and we were just recording everything, how much would that help us? I mean, in a sense, there's always the question of when do you see the forest for the trees? You can have all this detailed information. But if you can't kind of summarize it in sort of a narrative of history, it doesn't do you very much good.

One thing, I know I haven't been doing this, unfortunately, two iterations of this company went out of business, but there used to be a company that would make little cameras that you could, sort of clip on, that would take a picture every, I don't know what it was, 15 seconds or something, and I used to take those to things like trade shows.

And then I would be able to, after the trade show, I would be able to see what booth did I visit, or what person was I talking to, because the camera's resolution was enough that it could read the badge of the person I'd been talking to. And it's like, yeah, I talked to that person for 12 minutes, or whatever.

And… but I have to say that it was sort of interesting, because it was very…

granular, and it was very kind of hard to know what was important out of it. I mean, I suppose it's like looking at security camera footage. Most of it is… is quite uninteresting. I mean, sometimes you can get, you know, there was motion then, and it wasn't… it wasn't just a, you know, a bird or something.

But I think this question of how do you extract from sort of a bigger stream of data, how do you extract things that are useful, is an interesting one. I mean, I personally collect all kinds of personal analytics, whether it's, you know, every heartbeat, every step.

every keystroke I type, all those kinds of things. What do you make out of that? I can tell you that I've occasionally made very aggregated conclusions from those things, but it's not like the immediate recall of what was the sequence of keystrokes I typed at this particular time. That's really important. It is interesting, when I've discovered things in science, for example, it's pretty cool to be able to say, this was the second when I made that discovery.

It's, you know, this was the set of keystrokes I typed.

This was the thing I saw on the screen. It happened at, you know, 11, 23, and 42 seconds, but that was what I saw, and that was what I had figured out at that time.

It's interesting to go and do that. I don't know, is it fundamentally significant to be able to do that? I'm not sure, but it certainly gives you this sort of new level of computational history.

But the problem is, how do you see the forest for the trees? How is it the case that out of all these little tiny point things of every keystroke and so on, how do you extract from that the overall narrative of history? It is an interesting, sort of, general conceptual point.

that there is typically an overall narrative of history. It could be the case that, you know, in the world with 8 billion people and so on, everybody's doing their own thing, there'd be no coherent narrative about what's happened in the world. People would always argue about which narrative it should be, but fundamentally, we at least have the… we imagine that there's a coherent narrative about what happens in the world.

And there's sort of a significant history that one can tell about what happened, from which one can make new conclusions. And the same is true with us as humans. It's kind of… it will be interesting when there are tools, and this is something I've been interested in being able to do, and LLMs help a bit.

To be able to go from all those, for me, you know, whatever it is, 4 million emails.

and be able to tell the story. What happened? You know, each email by itself, I jump into some random time in 1992 and look at what emails there were. It's like one drowns in all the detail. But, you know, what was the overall arc of that narrative? What was happening at that time? That's something one can now expect with sort of a computational history to be able to do.

And I think that that,

it's, to the question that was asked about whether, sort of, cybernetically enhanced memory levels things out. Look, I think that the presence of knowledge on the web, the presence of, kind of, all of the tools I built, other people have built.

the presence of more and more, kind of, personal data that one can record. Yes, there are lots of things that, in the past.

one would have been blocked from a particular area. For example, you know, for myself.

I was never a big enthusiast of doing math by hand.

wasn't… I don't know, I don't think I was particularly good at it, I don't really know whether I was good at it, but I didn't think I was particularly good at it. And that could have meant I'm blocked from doing theoretical physics, because I don't want to be doing all these complicated algebra calculations by hand.

But for myself, I realized, wait a minute, I can just make a tool that does that.

And I've absolutely not been blocked. In fact, my tool has given me a big platform to be able to do much more in that area.

And so I think this question… it is certainly true that there are things where the automation of those things will go from somebody that is blocked from working in that area because… because of something that they're not so good at.

to something where they don't have to be blocked at all. I mean, if you… if you fundamentally can't spell, spell checkers will save you.

You know, if you, an interesting one for me right now is I don't do physical experiments. I have done, you know, only a handful of chemistry experiments in my life. They mostly didn't work, etc.

So, an interesting question. I've worked for many years with a company called Emerald Cloud Lab that's been building a big biology and chemistry lab in the cloud. I mean, it's a physical, sort of factory-like thing where you program it in Wolfram language.

And it goes off and does physical experiments, and does all that pipetting, and all those kinds of, you know, mass spectrometry, or whatever it is that you need to do. So an interesting question is whether, given that level of automation, a person like me can actually do physical experiments and have them work out.

I did make some attempts there, and I discovered that there were so many details in how you set the PCR machine or whatever that I don't know. I think that's now improving by using machine learning to look at, kind of, what the defaults should be based on what people have written in papers and so on.

But that's a, that's an interesting case where, for me personally, I'm not…

I don't know whether I would be good at doing physical experiments. I haven't done very many of them. I have the self-image that I don't think I'd be particularly good at them. And it's, so, you

know, will the automation of these things allow me to be at the point where I can do more science because I can now routinely do physical experiments, or not? I think it's very likely that it will be.
So, you know, these things where we get, sort of, cybernetic enhancement
they allow people who would otherwise be blocked from a particular area, as I might have been, from theoretical physics, had I not had, kind of, computers as tools, that allows one to kind of level that playing field. I think that's absolutely correct.
And I think that, you know, when it comes to, sort of, memory of things.
you know, in the doing of science, yes, it is super useful to be able to remember things from here and there and pull them in when you need to. It's not necessary for… for doing… there are many areas of science where you don't need that. Yeah, I think that's a… but it's certainly useful when you can get it.
All right, I should wrap up and go back to my day job here, but thanks for some interesting questions, and I look forward to chatting with you all another time.
Bye for now.