Hello, everyone. Welcome to another episode of Q&A about science and technology for kids and others.

So, let's see… Gosh, there's a question here I'm noticing first, from Tank.

Why do we have a brain when all of our cells have a brain, aka a nucleus?

Well, nucleus of a cell works in a very different way than a brain of a critter. The things that have brains tend to be animals.

Animals have the feature, unlike plants, that they move around.

And I'm pretty sure that the number one function of the brain of animals is to make, kind of. Decisions about what the animal as a whole should do.

And so, the animal is noticing all kinds of things are coming and attacking it, or it's got food in some direction or something. It's got to decide, does it… does it go for the food on the right, the food on the left? What's it going to do? It's got to have some way of deciding that.

And how's it going to decide that? Well, it's got to have some way of fairly quickly, taking all of the sensory input that it has from its eyes or whatever else.

And then making a decision, you know, move the muscles on the left, move the muscles on the right, or whatever.

In… so, how do decisions get made in biological systems?

There are different levels of things that happen in biological systems. There are things like muscles.

Where the… it's an electrical signal that makes the muscle contract. The muscle is made of protein molecules, and those molecules are such that if you put essentially an electric current into that muscle, it will contract.

And that happens very quickly.

So that's a thing which… there's no change in the material, it's just when there is electricity, electrons, and so on, it will make the muscle, get… get shorter. And that's what happens. Every time we move around, that's what's happening, is our… there are electrical signals being sent to our muscles, and our muscles are moving.

But there are other things that happen. Like, for example.

We are generating different molecules, and inside cells, there are new protein molecules being made all the time.

So, in, there are probably, oh, probably every second, there are probably trillions of… is that… yeah, that's probably right, trillions of new molecules that are getting made somewhere in our bodies.

How does that happen? Well…

there's a whole sort of chain that starts from the nucleus of the cell containing the DNA. The DNA is a program that specifies what the sequence of elements that will be added to proteins. Proteins are long chain molecules, and the DNA

program is specifying, do you add this unit, this amino acid unit, to this protein, or this one, or this one? And it's specifying, for sequences of thousands or more of proteins, how do you assemble these protein molecules?

And the synthesis of proteins can start happening very quickly in the cells in our bodies.

And that's sort of a slightly longer-term thing. So things like muscles, it's just electrical signal, the protein moves, but then there is make a new protein, which takes a bit longer to happen.

Or there's things like make a new cell.

For something like a bacterium that has fairly simple cells, you could make a new cell in half an hour. For something like us humans, I think it's… it takes a day or something to make a new cell.

So if it's a question of, you know, are we… if you've had some wound or something, and you're going to make new cells, that can take a little while.
Same with if you're, but…
So, so there are these different timescales of things that… that tend to happen in biology. And I think the idea of kind of a centralized brain
like we have and other animals have, I'm pretty sure that the reason that that sort of evolved in the history of life on Earth was the need to make this kind of general decision. Do you move left or do you move right?
Plants don't have kind of a centralized brain in the same way that animals do.
And, although, you can, you can end up with, I mean, plants will respond to their environment, you know, a plant, I don't know, a sunflower will open, will, will point towards the sun and things like this.
That's a little bit like what animals do, but the mechanism is quite different, and it doesn't involve the same kind of, sort of, central, sort of central decision as with, with brains.
So I think the,
The thing that, the individual cells
make their decisions in some sense separately in one's body. The cell has… is… is aware of, kind of, what is locally around the cell, the concentration of chemicals, maybe whether there are cells that are right up next to the cell that you're dealing with, but
The kind of the… most of the time.
the kind of decisions the cells make are local like that. There are only… there are some things in biology where there's sort of global
things that get communicated to cells. So, like, hormones, for example, are chemicals that sort of go all over the body, and sort of are general instructions to cells about things to do.
And, so that's… but that's, again, a, you know, that's a… that's a pretty slow process relative to, kind of the… the time scales of, moving muscles and things. I mean, brains
for folks like us, brains respond in about one-third of a second or something, from the time you kind of see something to when you start to move your finger or whatever else, it's about a third of a second. If, for other kinds of animals, there's sort of less
kind of going around inside the brain to do, and so the reactions can be faster, maybe a third of that time, maybe a tenth of a second, maybe even faster than that. I suspect for very small animals, the reaction times are even faster than that. But basically, for us humans, it's like the brain can do all kinds of things with the input that we're getting.
And it sort of takes a long time to sort of rattle around in our brains, and there's a lot of neurons in our brains, whereas for other kinds of animals, that's not the case, and it's a much more direct see-it-do-it type thing that, takes… takes less time. I think the,
But, but, the thing that, that's kind of… the idea of, kind of, the central control is… is coming from the brain, and in fact, many,
the brain has glands like the pituitary gland that secrete chemicals that do have, sort of, global control, but as I say, the timescales for that happening are significantly longer than the timescales for the activity of neurons and electrical activity that leads to, sort of, move muscles and things like this.
Well, a few thoughts on kind of an interesting question there.
Let's see…
Oh gosh, there's a question from Tiny here. If I made a very tiny robot the size of a dust speck, what would be the hardest thing for it to understand about the world?

Well, I don't know. If it has eyes, then it can see photons of light coming from wherever. It could see photons coming from a distant star that have traveled for, you know.

hundreds of years to get to it, and so on. So I think by the time it has eyes, it can… it can kind of sense a lot of things about the world. I would say that there are aspects of kind of,

the experience of the world that might be a bit different if you're sort of very small. So, for example, for us.

The,

For example, things at the scale of molecules, where… so, in… in our normal experience, at the normal temperatures that we live in, so to speak, there are… the… that… that temperature molecules are sort of bouncing around all the time.

In all kinds of materials, but we don't notice that, because there's so many molecules that, for us, kind of, the average is something very stable.

If you're very small, then you might actually be sensitive to sort of individual molecules bashing you from one side or the other, and it might seem that there wasn't the same kind of, oh, there's just a solid object that is just sitting there, and nothing's moving about it, so to speak, because at a microscopic level, still the molecules in it are kind of bouncing around, at least at sort of room temperature, that would be the case.

Nice.

And I think… I think that's a, that's sort of a difference in… we're used to certain things being very stable and other things not being. I mean, we're perfectly used to, you know, the wind blowing this way and that way, or ocean currents or waves going this way or that way, but we're used to the idea that there are solid objects where they kind of just sit there and nothing is

Really moving in them, but if you go down to a small enough scale, the molecules are still moving, even in what we think of as solid objects.

Let's see…

Well, BioBob is asking, why did evolution settle on cells as the basic unit of life? Could there have been a completely different architecture?

It's a very reasonable question. I don't think we really know. I think that, life as it exists on Earth today is

is all derived from a single common ancestor that lived maybe 3 billion years ago or something, that already has a cell wall, and it already has DNA and things like this.

There were presumably earlier forms of life that might have been very different, that may have kind of set up the environment necessary for the life of the kind that we have today to exist.

I think one of the things that tends to be the case is if you've just got molecules bouncing around, there's a question of how often will those molecules actually react with each other to… and have chemical reactions, and produce other molecules, and so on.

In most situations, molecules just don't interact very much. I mean, unless you're having some molecular process where there's an explosion going on and molecules are just, you know, interacting very, very rapidly and making something explode, most of the time, molecules just don't interact very much. They just… they bounce off each other, or they don't,

And they don't… and there isn't a way in which they chemically react and change the form of the molecule. So, there are various sort of tricks for making things react more. One of them is just to concentrate the molecules more, put them in a small container so that there's just more of those molecules closer together. And that's kind of one of the things that cells presumably do.

I mean, another thing is more elaborate, which is to have one molecule that kind of gets another molecule positioned in a certain orientation. So one molecule gets stuck in a groove in another molecule, and that positions that first molecule to be in just the right position so that another molecule, when it comes along, will interact with the right part of that molecule.

That's in the case of industrial chemistry, that's catalysts.

Often things like metals, which will have a certain surface, and molecules will get stuck in that surface and be oriented so that other molecules will come and actually sort of hook onto them in the right orientation so that they can have a chemical reaction.

In the case of biology, that's enzymes, which are big molecules, big protein molecules, that have a certain shape, and they have holes and grooves and things in them, where other molecules kind of get positioned in those holes and grooves and get oriented so that

So that molecules that are going to interact with them can actually do so. If the things were free-floating, they would have a much lower chance of being sort of oriented in the right way to actually interact.

So, those are some of the things that,

that sort of make biology possible, and I think the idea of having sort of enclosed regions of where things are interacting, which is what cells achieve, is something that seems necessary to get enough sort of chemical reactions going on to be able to produce new chemicals, and so on.

I think, the… the question of whether You could imagine…

oh, I don't know, for example,

I don't think one can imagine cells that are completely flat, two-dimensional. I don't think you'd end up being able to get things to sort of move around enough in that situation.

I think you could, you know, cells in biology come in all shapes and sizes. I mean, there's tremendous variation in the size of cells in us humans, from muscle cells, nerve cells, they're big, red blood cells, they're small, you know, there's vast variations in the size of cells. There's vast variation of the shapes of cells. I mean, there are…

cells in plants that are the shapes of polyhedra and things like this. There's,

there's a lot of diversity in that, but the idea that there is a definite sort of container in which things happen seems to be a fairly universal thing. Now, you know, the idea that that universal container can break apart and make two of itself when cells replicate, again, it seems hard to see how else that would work.

Well, it's always bad to say that, because I think it's sort of just a limitation of one's imagination, that one can't see how that could work. I mean, I suppose I could imagine something where you have, for example, a solid surface

or something like a clay, where in clays, there are sort of gaps in between solid… pieces of solid, and in a sense, you have… I suppose one could imagine something where you have a large number of sort of very tiny cell-like openings. Instead of having a whole big cell with its big cell membrane and so on, you could imagine something where there's sort of a thing where chemical reactions

happened, but they're in lots of tiny pockets in a material, and maybe that's what happened.

There's a sort of one of the theories for origin of life

Implicates clays a lot in being able to bring together chemicals to be able to do those kinds of chemical reactions.

Well, Guadeloupe asks, why did evolution stop at one brain per organism?

Well, I think… that… The… the main thing is, if you're an animal, and you move.

you have to collectively decide to move the whole of yourself. So you better come to a collective decision. I want to go left or go right. You better not be having a debate about whether you go left or right, because you've got to decide one or the other thing to do.

Now, it has to be said that… so, you know, something has to make that overall decision.

That's not to say that there aren't plenty of brain-like structures that make, sort of, local decisions. So, for example, in our spinal cord, there are ganglia, which are clumps of nerve cells, that are sort of deciding some kinds of things. For example, I think… I don't know to what extent this is true in humans, it's certainly true in some other animals. The basic ability to walk and go, you know, one step in front of another, that is…

That already can be achieved in… in just, you know, nerve cells in the spinal cord without having to have, sort of, central brain control, so to speak. The same is true in,

There can be reflex arcs where, you know, you touch something hot and you pull your finger away, that didn't have to go all the way to your brain to get you to start moving that muscle. I think it's, and in… in…

Even in very tiny animals, like, like,

Animals with a very small number of cells, a thousand cells or something like this, you will, you'll… you'll see, sort of, these kinds of particular circuits, where if you poke the creature in one way, something else will happen.

And that's kind of independent of its full, kind of, centralized brain, so to speak.

So it's pretty common to have these kind of extra pieces, and, you know, people say, like, if you're a very big dinosaur, and you're moving the back of the dinosaur, and so on, you can have said the brain of the dinosaur can send a signal saying, move the back legs, but it could take so long for that signal to reach the back legs that, you know, those back legs might have had to move

by that time, otherwise the dinosaur would have tripped over, and so that means that the control of those back legs has to be in some ganglion that is separated from the brain, and where the ganglion was basically just told, keep walking, and it starts generating kind of the signals to get the muscles of the back legs to move, or whatever. And so that's a, I think, a fairly common thing. It's interesting that when one makes computers.

Computers have, also often have, for…

input-output devices, they often also have peripheral microprocessors and so on, where the main microprocessor says, you know, start reading this data, and the peripheral one will do it, and will maybe have some loop where it's pulling in more bits of data and so on, and sort of… and then later, the sort of central

processor will tell it, stop doing that, or something, but it will have done that sort of somewhat independently, much like the ganglia controlling, you know, the back legs of the very long dinosaur, and so on.

Now, I think the,

it seems to be something pretty fundamental to the way that, sort of, animals exist, that they have one main brain that is making, sort of, collective decisions about what the creature should do.

When we're dealing with computers, for example.

In… it is, that one can often have many, many processors, many cores in a processor, many different, sort of.

Threads of execution of programs and so on. And it is a very non-trivial thing to deal with the coordination between those different, those different kind of, those different places where computation is going on.

with… With biology, what…

one has is, you know, we're all separate organisms, but somehow we get coordinated as a society, or if we're an organism that lives in herds, the herd of whatever, sheep or something, is coordinated in what it does. The school of fish is coordinated, the flock of birds are coordinated in what they do. So those are cases where

There is a, you know, each bird has a separate brain, but there's somehow some kind of collective activity of all those birds together.

And, exactly, you know, how does that work? Well, there's a sort of active effort of the birds to coordinate what they're doing. So, you know, birds are looking out to the… you know, they have eyes on the sides of their head, and they're looking out, and they're looking at the neighboring birds, and it's probably a pretty good model of how the flock of birds moves around that, essentially, a given bird

It's keeping neighboring birds a certain distance away.

Interesting question is when you see a flock of birds and they change direction, how does that happen?

Was there a bird in front that sort of made the decision? The leader bird or something made the decision? Apparently, that's not how it works. Apparently, it's always the case that there's sort of some wiggling around of this flock of birds, and somehow the decision sort of… at some moment, there's enough of a decision, it can start somewhere inside this flock of birds, that that sort of… that that, oh, I'm going to move to the right

That sort of starts an instability in the flock of birds that eventually causes the whole flock to turn around.

So I suppose that's a sense in which there isn't, in some kinds of animals, in some sense, there isn't just one brain, because if you think of the whole, kind of,

the whole kind of flock as being a single kind of thing. There are many brains involved in that, but they have this kind of coordination mechanism. I have to say, when it comes to computers. the… this whole question of how the coordination of things works is a challenging one. Sometimes you can do like a flock of birds does, and a whole bunch of computers will all be doing the same thing. That's what happens in, well, for example, in GPUs, there can be,

sort of multiple pieces, multiple different computations that are going on, but they're all kind of going on, in the same way, so that it's like all the… like, all the birds are flying in the same direction, all the computations are happening, in the… sort of the same computation happening to different

Pieces of data, for example, but The general problem of

Seeing how you get lots of different, lots of different

processors in a computer or something to do a problem in a coordinated way is quite difficult. And it's something we have a hard time… we as humans have a hard time kind of wrapping our brains around, probably because we're used to this sort of single thread of experience, single set of decisions that… that get made, and it's a little bit hard to, sort of do parallel programming. It's very common that there are bugs that get generated because it's like you didn't quite… weren't quite able to foresee. If this happens and that's happening at the same time.

time, and one thing ends first, and so on. It's hard to tell what's going to go on.

Let's see…

Lazy is asking, if humans had evolved during the time of dinosaurs, would we have had giant humans?

You know, the idea of megafauna, very big creatures.

there are times in history when, sort of, very big creatures have evolved. Usually, what happens is.

there are, you know, if you have a species where, I don't know, elephants or something, where there's a very big creature, or whales, or something like that, you… you know, it's sort of a trade-off. You can have a small number of very big creatures, or a much larger number of smaller creatures.

And usually, I think the total, if you just added up the weight of all the very big creatures versus the even the many more very small creatures, the very small creatures sort of always win, in practice, in what happens in the ecology of the Earth.

And I think, you know, the problem with a very big creature is how's it going to eat enough to survive? It's got to be sort of scooping up lots of small critters to eat, or lots of small things to eat to survive is a big thing. I think…

I'm… I'm trying to remember. There is… there are theories about why megafauna have evolved at certain times in the history of life on Earth.

I'm not sure, I don't think, I mean, it would be a reasonable guess

It had something to do with things like the oxygen content of the atmosphere, but I don't think that's what people have claimed about that.

So, this question of, I mean, it's worth emphasizing that even in the time of the dinosaurs, there were plenty of small creatures as well. Just like today, there can be, you know, giant whales and elephants and things like that, along with plenty of small creatures, you know, and and you know, little tiny algae and plankton and things like that in the ocean.

and yet very big whales. And it's… it's, so it isn't the case that, you know, in a given time in biological… in the history of life on Earth, that it's like everything was big or everything was small. No, sometimes there's a tale of very big things, even though lots of stuff, most things are small.

And I think what,

I know in… in Ocean, well.

It's… it's always complicated, because it's like, in… in the ocean, it's like, big fish eat little fish. And there are often many, many, many levels of that, like 10, 12, I don't know, levels of so-called trophic levels of… of one thing eats, another thing eats another thing. I think on land, there tend to be many fewer levels, like 3, 4, 5 levels, and I'm not sure the, and that may have something to do with the extent to which you can have, you know, you can… you can have, sort of, support these bigger creatures, although, you know, whales do eat, you know, krill and things like that, which are very

tiny.

So, I'm… I'm not sure… I'm not sure that that's a complete story. But anyway, so I… I… I don't know whether the… I mean, us humans have,

there are all sorts of different features of biology that make us sort of suitable for the size we are. I mean, you… but every time you change the size, you've got to have other mechanisms, like, you know, if you're a giraffe with a very long neck, you've got to have an elaborate setup with your heart to be able to pump blood up to your head, up that long neck, and so on. But, you know, biological evolution

tends to find a way to solve problems like that. And, similarly, you know, if you're a dinosaur that's really big, you need… you need bones that are strong enough, and you need to be light enough that you can kind of support yourself.

And it can't be the case that the, that you have, you know, bones have a certain strength.

And, so if there's a limit to how big you can be, and not have your legs, break every time you walk, and so on.

Well, let's see…

There's a couple of questions. There's one from Gregory. Is evolution still happening to humans in a meaningful way, or have we stepped outside it?

Well, I mean, the question of what new humans are made has to do with who has children and, how do they choose their mates, so to speak. I mean, that's the… that's, that's a… you know, that is ultimately the thing that determines the future of the species, so to speak.

You know, right now, we have about 8 billion humans on the planet. In the history of our species, there have probably been about 100 billion humans that have lived at some time in the last few hundred thousand years that our species has been around, and what those

Those… who those people have been, what those humans have been, depends on, you know, what children… what children were born, so to speak, which depends on what parents got together to have children, and so on.

And so, it's a… it's an… there are… there are different dynamics that determine that. I mean, one thing is, if you die before you could have children, well, then you're not going to have children. And if you have certain traits that

That make you, sort of, not be able to survive, not be able to eat, move around, do whatever, you know, fight well, whatever it is.

If those traits kill you before you can have children, well, then the traits that you had, that those traits won't tend to be propagated, won't tend to survive in the species.

But, for…

Us humans, you know, in lots of parts of the world, lots of things that would have killed people very early in life

are now medical issues that can be dealt with. And, you know, just take a pill every day, and you won't die from this or that thing. Whereas in the past, you know, if you had that this or that thing, you'd be dead within a year. You wouldn't be able to pass on that trait.

And so, you know, we've done a certain amount.

To, sort of, prevent evolution by the fact that traits which might make you unsuccessful in a world where, kind of, where there isn't medicine, will,

We'll make,

you know, then those traits will be… will be kind of wiped out of the population by evolution, because those… the folks with those traits just won't survive to have children and pass the traits on. Well, now the question is, okay, if you have other… there are many other… so…

You know, medicine is making people survive various kinds of physiological traits.

Then the question is, well, who decides to have children? And, you know, who… who has children with whom, so to speak?

And, you know, if all the nerds get together, so to speak, maybe if there's… maybe the children are nerdier from, you know, pairs of nerds getting together, and if the nerds always get together with the non-nerds, then you won't have a sort of higher degree of nerdiness being produced in children and so on. Not that we know, really, what the genetic

you know, what the traits associated with what we might describe as nerdiness, really are. But so, there's,

there is an effect, as a result of, the way society works and so on, there is evolution going on in, in human populations. I think it's, it's, it's not, sorry.

the, it's not, I don't think we really understand very clearly

what that evolution is like. I mean, some aspects of that evolution one might think of as being kind of societally divisive, and maybe one doesn't even want to know how that's working, but I think it's something where
Where maybe not as much as understood as could be. We do know things like, haplotypes. Like, for example, you can identify
certain sort of genomic patterns that are associated with particular, I don't know, ethnicities, places of origin, and so on. Those are things which one can trace, and there are lots of maps showing kind of how, sort of, our species started in Africa and then moved to lots of other places, and how, you know, groups that have been sort of
more or less hanging out in one part of the world for the last, you know, 10,000 years, sort of concentrate certain kinds of genetic characteristics, in that part of the world. I mean, this is pretty obvious from, you know, if you go visit different countries, you know, the
the typical who's wandering around on the street looks a bit different, and that's, you know, that's a consequence of that phenomenon. But how, you know, for example, one thing that is probably the case
is that one piece of human evolution is that it was the case that people stayed very localized. People weren't, you know, they were not moving very far from where they were born. Now people are moving much further. That means there's much more mixing in the genetics of, you know, it can routinely be the case that somebody who was, you know, whose ancestors had lived in, you know, Japan for however many generations.
has children with somebody whose ancestors lived in, you know, France or something for who knows how many generations, and the result is, you know, children that are kind of new genetic, you know, have a new genetic content that hasn't been so, you know, hasn't really been seen before. So that's a way in which the human population evolves,
Perhaps in a rather interesting way.
So a few thoughts about that.
let's see
Mary is asking, could there be fundamental computational limits to how complex biological organisms can become through evolution?
I don't think so. I think that…
one of the things, and I've looked at this a fair amount, actually, in recent years, about, sort of, the computational analysis of biological evolution.
for biology, it's, you know, DNA gives us sort of a program for building biological organisms. So, you know, we each start from a single cell that has a certain
DNA molecule inside it.
It has, and that molecule has these 6 billion base pairs, which are essentially a program
for building the rest of us, so to speak. And as more cells are produced and so on, different kinds of cells get produced, and then within each kind of cell, different proteins get produced and so on, and that makes everything that is us. So we are the result of that DNA program
that we got started with, and then all of our tens of trillions of cells use that DNA program to determine what they're like.
Well, you can do the same kind of thing just with raw programs in a computer.
very simple programs, for example, where you just say, I've got this very simple program, what… how does… what does it do when that program runs? And you get some whole pattern that you can display on… on the computer or whatever, and you can say, well, with this program, it does this. Okay, let's make a tiny change to that program. What happens then?

And then we can ask the question.

if we make tiny changes to the program, how do they affect some feature of the… of the output from that program that we want to, kind of, that we want to get? Like, we want the output, the program, to kind of run as long as possible before stopping, for example. And how do we tweak the program to make that happen?

Well, one of the things that is very surprising is if you have even fairly small programs, you start tweaking them sort of one bit at a time, they'll end up being able to do very surprising, very complicated kinds of things.

There's sort of a question of… of, if you… if you say, well.

You know, I've got one program to do some complicated thing, can I always go further and get another program to do an even more complicated thing? The answer is, sort of, for fundamental reasons, the answer is basically always yes.

It's kind of like there's a certain… there's this kind of idea of universal computation. The idea that once you have a certain kind of… once you have essentially a computer, a general-purpose computer, you can program that computer to do anything you want.

If your computer… the only thing your computer, in quotes, could do is add numbers. It could do nothing else.

then there'd be a limit to what you could get your computer to do, basically, just to add numbers.

But as soon as you have a bunch of different things that a computer can do, add numbers, subtract numbers, decide to

go to the next instruction in a program, decide to jump back to another instruction, all these kinds of things. Once you have a… it's a rather small collection of different kinds of things that your fundamental piece of hardware can do, then it becomes the case that you can essentially write software of any kind to run on that hardware.

and get that computer to do, sort of, anything you want. And so, that's,

As soon as you've gotten to that point, where in principle, you can program your system to do anything you want, then the question is, if you're… if you're not… when we think about doing programming, most of the time, we're thinking about a human writing lines of code and deciding what line of code to write next.

Or maybe in modern times, an AI doing that.

But the other thing that can happen, and what happens in biological evolution, you just have this program, and just through the sort of random things happening in the world, when there'll be, sort of, when a new organism is being created, there'll be some bits that change in its program, and it will have some differences from the parent organisms, and so on.

And the question is, sort of, how far can you get by just making, sort of, one bit at a time changes to the program for a biological organism?

And the answer seems to be that you can get as far as you can get. In other words, as far as… that you can eventually reach the point where you will be able to do what any program can do.

It's not obvious that would be the case, because it could be true that if you try and make changes, sort of one bit at a time, you always end up having to go through some sort of valley of death, where the organism just couldn't make it out. It couldn't, it couldn't… you made those changes, you really want to get it to the other side.

where the creature has, I don't know, you know, wheels or something.

But as you go from where it is now with legs, there isn't a way to sort of make a small change to the legs that then turn them into wheels.

And that you'd always have to go through a situation where the organism couldn't survive. I actually think that will essentially never happen.
I think we may not be able to foresee how do you get from legs to wheels or whatever else, but I think that it will essentially always be the case that there is a path to getting from one thing that's computationally possible to another thing that's computationally possible, a path that goes through only these very small kind of mutations, kind of one bit at a time.
That's something that is sort of a… that conclusion is something that I would say is a very recent thing coming out of a bunch of things that I've done in trying to understand the foundations of biological evolution, but my guess is that that's the way it will work.
And that means that, in some sense, in principle, there's sort of… once you have a universal system, there is no limit to what it can, to sort of how far it can evolve. Now, you know, you can ask, is the way that biology is made an adequately universal system?
We're all made up of certain kinds of molecules, proteins made of amino acids, and so on. And, you know, we're all based on DNA and things like this.
there are certainly things that it is implausible that you could make. I mean, I think it's… I say nothing, you know, you can't say it's impossible. I might say, you know, a creature with a big golden horn is probably impossible. I say.
And then I think, is that really impossible? Because, you know, you can't… a big horn made of gold
doesn't… there aren't cells in it, it's just a big lump of metal, it can't grow on its own because it's just a big lump of metal, it doesn't have biological tissue inside it. But does that mean it's impossible to have a creature with a big golden horn? I don't think it's impossible, because
What happens in other situations is, for example, we accumulate, well, for example, there is a… what is the situation?
There is, A case where some organisms concentrate metal
in… in a collection of cells. And so while you couldn't get the sort of horn to grow by itself, you probably could imagine a situation where there are,
Biological cells that concentrate a certain kind of metal, and, you know, you've ended up concentrating this metal, and eventually the sort of biological piece falls away, and then you've just got that piece of metal.
One of the things that's done, for example, in hemoglobin, the protein that's used to carry oxygen around in our red blood cells and so on, the way that works is that it's a protein made of amino acids and so on, and carbon and hydrogen and oxygen and phosphorus and a few other things, but it…
Has a kind of hole inside the molecule that's sort of just the right size to fit an iron atom in.
And so, when you just get iron from eating all that, whatever spinach that you eat, or anything, or lots of other things that you can eat.
that have, have iron in them. Then, you know, the atoms, the iron atoms are going all over the place, and one of them sort of gets,
gets near the cage in that molecule, it gets stuck in that cage, and so you're kind of accumulating iron, in that case one atom at a time, in the biological organism. But there are other cases, and I'm forgetting the example where this happens, where you have
A sort of a large concentration of metal in… in… in… inside, sort of, cages in a protein.
And, you're kind of deciding which metal it is by how big the hole is, because these different atoms effectively are different sizes.
So…

you know, in the question of whether you… what can you get to with biological evolution, I think the answer is you can get a surprisingly long way, and we don't know how far you can get, and even these things where you say it couldn't possibly make a this, that, and the other, it might be able to.

it sort of has an interesting resonance with the world of AI, because in what's done right now with training AIs, you're just doing this process of saying, okay, I've got this AI, it has this particular form, I'm going to try and tweak the details, the weights inside the neural net.

so that the AI will be able to, sort of, reproduce my training data. It'll be able to take kind of text that it read on the web or something, and be able to… if you mask out a word in that text, it'll be able to guess with reasonable likelihood what that word should have been, and so on.

But, you can also ask the question, if you define some sort of overall objective for the AI, could the AI do something more like biological evolution, where it keeps kind of tweaking the program for the AI, and then kind of, eventually, and then when that program has been tweaked, the AI will do something different.

And then, can you sort of navigate it to do the thing you want?

And that, that's a thing that's kind of… it's related to reinforcement learning, which is a method that, gets, gets used for sort of the final training of your average modern AI.

But there's also kind of different directions there, where you're sort of, kind of,

repeatedly kind of looping through, I've got this form of AI, now it does this, now I want another form, what is it going to do, and so on. And I don't think we know yet, exactly how that… how that ends up.

it's important to define, kind of, what the… you know, given an objective, you can say, can I evolve my AI to get to that objective? Often, as in biology, the way it will achieve the objective is incredibly hard to understand.

just like in biology, you know, can fill thousands of pages of textbooks with all the details of how biology happens to achieve certain things that biology does. And it's not like there's a… it had to be that way. It was, you know, a clean engineering design.

It's like it happened to be able to be put together that way and worked, and so biology kept on using it.

So…

Let's see… Tom is asking.

Why has evolution repeatedly converged on similar structures, like eyes and wings and social behavior? Does that suggest deeper laws constraining biological design?

Reasonable question. I think the thing to understand is when there are these cases, like eyes, for example, that have evolved separately, I think eyes have evolved, like, 5 times in the history of life on Earth, it's like, what do we exactly mean by an eye?

Do we mean something, a molecule that's sensitive to light? Well, there are tons of those. Do we mean something which takes the sort of… the things that are, kind of.

take light falling on some part of an organism, and that leads to an electrical effect in a brain-like thing, that's a slightly different kind of thing.

my feeling is the things we describe as eyes are often really quite different from each other. I mean, in, you know, there are light-sensitive pieces of, I don't know, worms and things that really are just directly go from, there's light there and a chemical is produced.

For us, we go, you know, there's light there, an electrical signal is produced from our photoreceptors on our retina, the electrical signal goes through a nerve, the nerves to our brain, and then eventually, often a chemical is produced as a result of that.

or even more directly, we will have our circadian rhythm, day-night rhythm. There are… there are chemicals that we are producing in our brains, actually, that are sort of sensitive to the day-night cycle.

In other organisms, possibly also in us, there's sort of a direct effect of light on the production of those chemicals, not going through the brain. Just a purely sort of physical chemical process leading to different, different levels of chemicals and so on.

So I think it's always a little tricky to say it's an eye, because we have that level, you know, we have that description of eyes, but I'm not sure that it's an eye in quite the same sense in all these different cases.

I… similarly with flight, for example, the way an insect flies is pretty different from the way a bird flies, and the way some,

You know, I think, oh, I don't know, a pterodactyl from the dinosaurs is, again, different from the way that a bird flies.

Now…

you know, is it… it goes through the air versus it, it walks on the ground? I suppose we could say, you know, something discovered that it could go through the air as well as walking along the ground. There are plenty of things when… there are things which

seem like they're not physically possible, like… like a creature that boars its way through a rock. Having said that that's not possible, I realize there are creatures that will at least get into cracks in rocks and so on. And maybe there are other ones that will actually etch away a rock, I'm not sure. So, I… I'm… I guess I'm… I'm not… I think…

that there's sort of different ways to move around, different ways to survive in the world. Once you've defined that you want to move around in the world, it becomes sort of unsurprising that you end up with things that fly a bunch of different times, I think.

But it's an interesting question.

Okay, pale blue comments, about my comment about making gold horns.

the, perhaps pearls could accumulate gold instead of calcium. I, I think, I mean, one feature of, Biology doesn't make use of most heavy elements. Biology actually just has a very bad time with heavy elements like lead and so on.

And by heavy element, I mean, you know, in the periodic table, it goes hydrogen, helium, methane, beryllium, boron, and so on, and those successive elements have successively larger nuclei, and usually biology doesn't get… I don't think it ever gets past iron.

Is that true in things that it actually makes direct use of.

if I remember my periodic table correctly, I think that's right.

And usually, heavier metals

tend to be quite poisonous. They tend to accumulate, and so biology doesn't know what to do with that lead atom or something, so it just accumulates, and it doesn't get moved out of the body, and eventually it starts to cause trouble.

Let's see… maybe a very different kind of question.

Okay, Alan is commenting on my comment about centralized brains, that it might be an anatomical coincidence, and that octopuses are often thought to have somewhat independent brains in each arm. I think that's similar to what I was saying about dinosaurs and ganglia and different… and dinosaurs. I would suspect that for an octopus.

It, And my impression about octopuses

is that people sometimes say, octopuses are incredibly smart. Look at all these cool things they do.

the repertoire of cool things they do is a bit limited. I think one of the famous ones is to be able to unscrew the top of a jar and get out.

The, but I think…

That there are some very impressive ones.

that are a bit alien to us, because octopuses have rather different, sort of, sort of structures and nervous systems than we do. They look very impressive, but if you kind of try to get an octopus to do something else that might be quite easy for us, the octopus will absolutely not be able to do it.

So, it's kind of,

the… I think the intelligences of… intelligence of octopuses might be a little bit overstated. Not that I think there's a kind of linear scale of intelligence. I mean, if you put us in a situation of many kinds of critters, we will definitively fail. And yet, we might say, have I seen that critter be able to solve math problems? Absolutely not.

But nevertheless, that critter is going to be an absolute winner relative to us in climbing a tree, or something like this.

But I think in, in Octopuses, the,

My guess is that there are ganglia in the arms that are sort of, yes, they're somewhat independently controlling them, but there probably are features of the octopus that are sort of, you know, have some central control. Like, I think, you know, if the octopus is going to be walking in one direction, or moving in one direction, or swimming in one direction, there has to be some coordination in the motion of those tentacles, otherwise it's,

Otherwise it's just not going to, you know, if the octopus is just every arm for itself, so to speak, the octopus isn't going to go places. It's got to have some collectivity to that motion.

To actually be able to go somewhere.

Now, things like octopuses and things like cuttlefish and so on have an interesting feature that we don't have, which is that they have cells in their skin that change color, and they can change color under nervous system control. So it's kind of like we could think

I want to be green today, and suddenly our, you know, hands would turn green or something. And, you know, that's kind of what the octopus gets to do, or cuttlefish and so on get to do. And in fact, you know, sometimes, well, it's like chameleons have the same feature that they have, color-producing cells in their… in their skin, and for a chameleon, it's like, it looks around, and it's like, what kind of thing am I on? Am I on something green? Am I on something brown? I'm going to make my skin turn

the same color as the thing I'm on. Of course, if you put a chameleon on a mirror, it does terrible things. I don't think the chameleon is very happy in that situation.

But so this idea of, you know, you can sort of communicate, well, you can, you can put colors on your skin, so to speak. I don't think chameleons use that to communicate. I think cuttlefish, do. And you can kind of see videos of cuttlefish that are, doing things together and so on.

And they'll… you'll see these, sort of waves of color going across their skin. And it's sort of an interesting thing, because we humans, you know, we get communicate… to communicate by talking to each other and things like that, and maybe waving our arms around and so on. But cuttlefish get to actually kind of have instant

sort of visual communication in a way that we don't. I mean, if we wanted to communicate busily, we're going to have to draw that diagram to show what we're doing.

sort of an interesting case where AIs can sort of rather easily generate images, and that's a case where you could imagine sort of an AI communication situation where it's synthesizing one AI synthesizing an image, the other AI is looking at that image, deciding what to do, communicating back by synthesizing an image.

It's a place where, kind of, the standard kind of path of human language and so on

is not being used. The standard path of, sort of, doing things with language is not, is not, is not what's,

what's being done. Instead, it's this kind of, this kind of visual communication.

Let's see…

the,

There's a question from MX here. Do you think little brains creating for computing in labs now have consciousness?

Well, so, one of the things that's become possible is to grow…

Bunches of nerve cells. A million, 10 million, 100 million, maybe. Nerve cells.

When you grow a brain.

In the development of an organism, like us.

The brain… there's lots of other things going on in the development of the organism, and the brain is growing in the way that human brains grow. Assuming nothing goes wrong, you're gonna have a brain with all these different structures and lobes, and all this kind of thing.

when… People kind of grow…

So that happens by us going from the original, sort of, fertilized egg cell through many different levels of many different kinds of cells, maybe a dozen kinds of cells, eventually getting to nerve cells. But there's all kinds of scaffolding of how the nerve cells get arranged, and other kinds of support cells, and blood vessels, and all kinds of other things.

That are sort of providing an environment in which those nerve cells can be produced.

And so that's how an ordinary brain gets made.

Well, now one of the things that got discovered about 20 years ago is how to make stem cells that are kind of like fertilized egg cells, that are like cells which can turn into any other kind of cell.

So-called pluripotent stem cells.

And it got discovered how to go from basically any cell and reprogram it to be like a stem cell. Occasionally, things go slightly wrong, there's some sort of instability in the cell, and it kind of does remember that it used to be a skin cell, but in a first approximation, you're erasing the eye as a skin cell and putting it back to, I can be any kind of cell.

And so then, there are all sorts of protocols of kind of feeding it certain kinds of chemicals and doing this and that to it, which take, you know, for things like nerve cells, take two or three months of, sort of, daily, feed it this, feed it that, and so on. And eventually.

you get a cell, an individual cell, so to speak, or a whole clump of cells that are all nerve cells. It's not the same as when you grow a brain in an actual organism, where you're growing all those other kinds of cells, and you've provided the environment to really be a brain.

Instead, it's just a blob of cells.

And, that blob of cells, when… whenever you put nerve cells together, they will start to show activity. They'll start… because nerve cells are kind of like little batteries, they kind of have electrochemical activity. You put them together, and they'll start producing little sort of pulses of electricity that go from cell to cell, and so on.

And so it's kind of as if the thing is kind of lighting up and doing brain-like stuff, but it doesn't have the same kind of coherent structure that a fully grown brain has.

What does that sort of jumble of nerve cells, what does it think?

Well, my guess is…
It doesn't really have the same kind of experience that a human like us can have, because the way humans like us work
We are taking in, you know.
input from our eyes and so on, we're seeing all these different things that are happening, then we are… our brains are kind of concentrating down all that information into a, what do we do next? And in the structure of the brain, I guess it's the basal ganglia at the base of the brain that are sort of the final sort of… that are associated with
kind of making that consensus decision. You know, there's been 100 billion nerve cells in our brains, and maybe a whole bunch of them are saying different things, but in the end, something has to come to a consensus decision about do this versus do that. And there's definite structures in a fully grown, sort of fully structured brain that sort of enforce that single thread of experience.
If you just have a blob of nerve cells and they're firing randomly, you don't have that same kind of single thread of experience, which seems to be a crucial feature of consciousness as we experience it, so to speak.
So my feeling is that with most kind of definitions of what it means to have consciousness and a memory of the past, and kind of a thread of behavior and so on for the future, I would think most of those kinds of definitions and experiences of consciousness would not be shared by kind of the random jumble of nerve cells in the jar, so to speak.
Doesn't mean to say that those nerve cells in a jar can't do some kinds of computations that are useful, and because there's plenty of computation, even that we do as humans, that doesn't go through the whole sort of conscious loop of remembering the past and having a thread of experience and so on.
I mean, you know, when we move our eyes in response to some kind of, you know, signal off in the… off in the corner here or something, that happens, faster than it goes through the sort of whole loop of conscious thinking.
And I wouldn't be surprised if… we don't do that, but I wouldn't be surprised if there were critters that didn't, like, move their eyes separately in different directions. Maybe a chameleon does that, I don't know, because it has separately, you know, separately movable eyes.
And, where the chameleon didn't probably have a single conscious thread of experience when its eye just happened to move around to that place, happened to move around to the other place.
And I think, so, you know, this idea insofar as sort of the conscious experience is this thing that is concentrated down into a single thread of experience, it's, it's something where there can be, sort of, other lower-level structures that are having things happen at the level of nerves doing things, but not, kind of, bringing it down
To sort of a single thread of experience.
Let's see…
Simon's asking, do you ever think there are types of evolution that might be more efficient in computing? In nature, an organism has a genome, but digital strategies don't have to follow the same rules. Just curious what you've thought about in that area.
this whole question of how you might evolve things. If you take a program, you make changes in it, and for example, one thing that is always a question of can you evolve the evolver?
So, one thing you might do is say, I've got this program, I'm just going to make random changes. I'm just going to zap this one bit here, one bit there, and so on. Or you might say, I'm going to learn how you should make changes, what kind of changes are worth making?

So, yes, there is a way in modern times that you absolutely can think about things like evolving the evolver.
So, for example, if you have an LLM, you know, a modern AI model, you can say to it, I've got this program.
Suggest 5 interesting changes you could make to this program.
that's a way of using that whole sort of thing that's been developed as a way to say, well, here are ways you could sort of make changes to that program. Then you can watch what happens with those changes that have been made, kind of feed that back.
as kind of changing the way changes are made, and you can imagine many levels of changing the way changes are made to changes, so to speak.
And so, yes, there are lots of things to do there. I suspect there are a lot of tricks that allow you to sort of more efficiently evolve programs in a way that's useful than we already know. So, for example, in biological evolution.
There are a bunch of tricks that are used. One has to do with different parts of the genome have different rates of mutation, and that's important because there's some things that have to be very stable, and without them, the organism would just die. Other things where you can be very variable. I mean, the classic example of this is in the immune system, the hypervariable regions that are literally just above
A bunch of, a bunch of pieces that get sort of randomly reassembled.
That's a place where it seems to be important to have… where it is important to have things be very flexible and very rearrangeable and sort of evolve at the drop of a hat type thing, versus things where, basically, if you don't have
let's say a ribosome that controls production of proteins and cells. If you don't have ribosomes, you're not going to be alive in the sense of life on Earth.
So that has to be kind of a fixed thing.
So I think, you know, in the history of life on Earth.
various things were invented. Probably the most famous one is sex. The idea that you don't just have one organism that is replicating, replicating pieces, and pieces break off and become new organisms, and so on.
From a single genome, asexual reproduction, but rather you have two parent organisms whose genetic material gets mixed up
Before, to get to the next generation. So, you know, we have, pieces of every chromosome come from, sort of, one parent versus the other.
And what you see on human chromosomes, you see, you know, if you do genetic analysis and you do it, you know, if you have multiple children, and you do genetic analysis of their genomes, you'll see that within a single chromosome, there'll be maybe half a dozen crossovers, where there'll be a stretch of genome where it all came from one parent, and then a stretch of genome where it all
came from the other parent, and they kind of, they're, they're typically maybe half a dozen of those in, in a, in…
In, sort of, in humans. I mean, it is an interesting question, the fact that there are these quite long stretches that all come from one parent, then it jumps to the other parent,
that kind of means that there are necessarily correlations between certain kinds of pieces of the genome, genetic traits even, associated with one parent versus the other. And I've sort of always wondered whether, if you have this thing with eye color and this thing with finger length or

something, that those are encoded close enough in the genome that they usually come together, so to speak. I've not seen an analysis of that.
But it doesn't sort of help the analysis that in the genome, it's not like the human genome is carefully organized. Oh, this is a section for the immune system, this is a section for the musculoskeletal system, and so on. It's all jumbled up.
And all the genes for different kinds of things are all sort of, you know, it's kind of a weird thing, because you're seeing… it's like you went into a library and there was no organization in the library, and it's just a big pile of books, and you know, like you sometimes see in used bookstores and so on, a big pile of books, and they could be about completely different things. That's kind of what the human genome is like.
But, you know, sexual reproduction and crossover is a mechanism that's sometimes been thought to be very important in, kind of, rapid
evolution of things, I don't know how important it is. My own experiments on that over the years have not suggested that it's as important as people have claimed that it is. But one could certainly imagine other kinds of mechanisms. I mean, you know, one could imagine instead of having, you know, two sexes, one could have eight sexes or something.
And that would lead to different kinds of, sort of recombinations of pieces of genome and so on. I don't know whether that would have an effect.
But I think this whole question about, sort of, given a program, how do you mutate it? Or given a population of programs, how do you combine them to make new programs and so on, there are probably a bunch of really good tricks to be discovered there that will speed up the process of, kind of, artificial evolution. I just don't know what they are yet.
Let's see…
Patches is asking, why do humans not have the ability to turn half their brain off to sleep, like dolphins, and so on? That seems like a useful survival tactic.
You know, the thing that… Tends to happen in biology is there are always trade-offs.
like… You might say.
Well, let's have an immune system that's really, really powerful, that gets rid of everything that doesn't even look vaguely right.
Well, if you do that, you end up with autoimmune diseases, and the immune system starts attacking your own cells.
Or, you say.
let's have cells that are, that live as long as possible, or that can replicate, you can have chains of replication that go as long as possible. Well, that's well and good, except that that's… you end up with tumor cells, then.
So, you know, biology is a… is a sort of a… very much a, a story of compromises between this and that. You, you know, it can't be all of this. It can't be that, I don't know,
You, you get to run extremely fast.
But yet, you also get to kind of, hang out and not get yourself in trouble by wandering around at night when the, you know, when the, nocturnal creatures are going to eat you, or whatever else. The,
there are… there are always these trade-offs, and I think
I don't know what it would take
for human-like brains to be able to switch on and off one hemisphere, for example. But it might be that the additional apparatus that you need to make that work, like, like, for example.

For example, I don't know… I assume that dolphins are motionless when they have half their brain switched off.

But for us humans, the motor area on the left side of our brain controls the right side of our body, and vice versa. So if we switched off the left side of our brain, it's kind of like, you know, we're not doing anything with our right hands, so to speak, at that time.

And depending on what situation you're in, you know, if you're just hanging out in the ocean, just… just sort of hanging out there, and nothing, you know, and there's no… and you're not having to move anything, then, you know, maybe it's okay to switch off half your brain, and, because, you know, you don't have to operate the… the right flipper or something, for a while. And then your part of your brain that's awake can notice, wait a minute, there's a predator attacking, switch on the other part of the brain, wake up, other part of the brain, and then you get to operate both flippers, so to speak.

But… but my guess is that there are,

there's a certain amount of, kind of, apparatus and mechanisms that would have to come with that, and well, you know, there are several possibilities. Maybe biological evolution for us just never thought about it.

Maybe it tried it in some way, and things went badly, and it backed out.

maybe, you know, it's hard to tell, and maybe there is something really bad about it. Maybe it's something you can't get there from here, sort of incrementally, as I mentioned, although I think that will be a rare situation. But, I think, you know, usually these things are not…

a uniform plus. I mean, you know, one of the things to notice about life on Earth is there's not just one winning species. There are lots of different species, that 10 million of them, roughly, that all coexist on the planet and occupy different niches. I mean, there's some which

live in this situation. There's some which only exists because of something that some other organism had done, those kinds of things.

Let's see, I should probably wrap up in a second here.

Let's see… Questions about very different kinds of things, which I don't want to launch into right now.

Oh boy, behind is asking, how does the brain turn chemical signals into specific tastes like sweet or bitter, and why do certain molecules taste the way they do?

That's a complicated can of worms.

More is understood about smell than about taste. One feature of smell is when you have a molecule in a gas that's just sort of floating around, and it sort of attaches itself to a molecule in your nose that can sense it, the molecule in the gas is not… there's nothing about the gas that's destroyed.

Taste, you're eating things, and you're going crunch, crunch, crunch, and you're destroying the food.

And so it, and, you know, the molecules end up on… or the pieces of the food end up in taste buds in your tongue, but the food sort of is no longer, it's in a different state than it was before.

So it's probably easier to talk about smell than it is to talk about taste.

There's a lot that isn't known about smell.

I think the general belief is that there are maybe a few hundred kinds of different molecules that sense different kinds of smells. Like, for color, for light, we have just three kinds. Humans have just three kinds of photoreceptors in our eyes, for roughly red, green, and blue colors. Other kinds of animals have different numbers. Cats and dogs have only two kinds.

I think the mantis shrimp is the winner, so far as there's known, with 15 different kinds of color receptors sensing different, essentially different bands of energy of photons corresponding for us to different perceived colors.

But so… For light, it's, you know, a few, a handful of kinds of,

receptors that sense what we interpret as different colors. For smell, it's probably hundreds, and maybe for dogs, it's a thousand different kinds of receptors that all sense

Different, and each one of those receptors is roughly sensitive to a molecule of a different shape. The details of how that works are not well understood.

But roughly, what you have is a molecule in your nose.

that, is, has a certain shape of, sort of, has a certain hole in it. And then molecules that you are smelling, and, you know, a dog goes sniff, sniff, sniff to kind of move molecules through its nose, so do we, the,

And, it… it is,

What's happening is the molecules that are coming through that are, you know, molecules of some scent or some smell or whatever, those molecules have a certain shape. And probably what's happening is that those molecules are getting stuck in the molecules that are the corresponding shape.

Actually, there's some similarity between the way that antibodies and things work in the immune system, where they're trying to recognize the shapes of molecules of antigens that are coming to attack you, so to speak, and the way that the molecules in noses seem to have a wide range of different shapes that can, kind of, that bind to, that attach molecules of certain shapes.

So, the basic picture is that, the,

That your, your, your nose has these different populations of cells that,

that sense different kinds of molecules. Now, a question is, how does that get mapped onto the brain? I mean, if you look at… in the visual cortex, at the back of our heads,

The nerves from our eyes go to the visual cortex, and roughly, if you're looking at a visual scene where you have something in the top left, something in the bottom right, and so on.

Those different positions correspond to different actual positions of nerves at the first level in the visual cortex where you're interpreting a visual image.

how does color work in that regard? So what happens, for example, for our left eye versus our right eye, we have different

kind of receptive fields. We have, on… in one part of your visual cortex, you'll have sort of fingers of information coming from your left eye, and fingers of information coming from your right eye, and there's a definite… there's sort of nerves that are all coming from the right eye, and there are nearby nerves that are all coming from the left eye, but they kind of form this zebra-stripe-like pattern of… of what, what

Nerves come from what part of the visual scene.

The same thing happens with color. There are… locally, in the brain, there are sort of little pockets of, okay, there's a bunch of red sensing, you know, neurons that respond to red, blue, green, and so on, and that exists at every place on… in the visual cortex. There's sort of a little almost color wheel of red, green, blue, of different things that are sensitive to different parts of different

Different colors.

How does that work with smell? One doesn't really know. There's some information that there's localization of the same kind, just like that in smell space. So in… if we are looking with our eyes, there is a different you know, different neurons that are firing when we see something in the top left and in the bottom right. And they're arranged in a sort of pattern, in a two-dimensional pattern, just like the scene we're seeing.

So the question is, that's in sort of visual space. We know things are on the left, things are on the right. In smell space, it's not clear what's on the left, what's on the right, what should be positioned where in smell space, so to speak.

I mean, maybe if we were all dogs who had a very keen sense of smell, we would have some sense of the geometry of smell space. That's something a bunch of people have worked on the last decade or so, those kinds of questions about what is the geometry of smell space, and if there is a geometry to it.

Is it related, is it reflected in the geometry of kind of the way that the brain is laid out to sense different kinds of smells? The answer is, I think, not clearly known, although there's some evidence that there's some kind of spatial localization like that.

Actually, somebody was just, telling me about… pitching me a company that can,

is working on bionic dogs, where the idea is that the dog is smelling things, and you're picking up signals in the brain of the dog that are sort of localized, and a pattern of signals that corresponds to what the dog would be sensing if

It was smelling some particular kind of explosive or some particular kind of… of sort of medical… medically relevant, gas or something.

And it's an interesting, sort of philosophically interesting situation, because it's something where you're implanting an array of things that can read out the nerve firings inside the brain of the dog, near its olfactory bulb, near its… near the place where it senses smell.

And you're kind of taking the information directly from inside the dog.

And so people often ask the question, sort of philosophically, if you take two different people, are they seeing the same red when they see a thing that they describe as red?

Well, usually we don't get to know, because we don't see inside the brains of those people, but with this kind of setup, with dogs, you'd be seeing, sort of, you'd be directly able to kind of know, how did the dog sense this particular smell?

Then it's a big question whether there's kind of dog-to-dog transportability, and you can just take the configuration of what, you know, this kind of explosive looks like.

in the smell space of dog number one, and give it… and sort of, you know, use it in the smell space of dog number two? Or is it the case that every dog has arranged its brain differently in terms of how it smells things? That's a question to which we might actually know the answer reasonably soon.

But to this question about whether, how, you know, what the geometry of smell space is, it's not really known yet.

And, you know, in the case of vision, we can plainly say that's on the left, that's on the right. In the case of… we've got lots of molecules of lots of different shapes.

it's much less obvious how we put those in a kind of space. I have to say a bunch of things I've done in areas like metamathematics have this problem of how do you put

sort of things that are described in different ways, how do you arrange them in a kind of space?

It's an interesting problem, and there are lots of different solutions to it, and maybe some of those will be the ones that brains used to have some coherent organization of sort of smell space.

Alright, I should probably, wrap up there with, with a…

A nod to bionic dogs.

That's probably a sign we should, I should go back to my day job.

Where I get to think about, other very different kinds of things. All right, well, thanks for joining me, and
Bye for now.