

Artificial Intelligence, Consciousness, Empathy and Love

AI can be a powerful tool for deepening our understanding of our true nature. When we use and experiment with a 'large language model' such as ChatGPT, we discover its real value from a spiritual perspective, lies not in what it tells us, but in our own questioning of previously unnoticed beliefs and assumptions that arise after conversing with it. Questions that come from astonishment, or perhaps a sense of unease at its human-like interaction with us. It is through those questions that the subtler layers of the sense of separation and limitation become apparent. And it is through our own answers to those questions that a deeper and more stable understanding dawns. Can an AI become conscious? How can I know if I'm communicating with an AI or a human? Does it actually understand what it is telling me? Is there any way in which human minds are superior to an AI? Can an AI express empathy? Can it fall in love? What is love?

In a recent satsang with Francis, four consecutive dialogues explored the nature of AIs, consciousness, empathy and love, shedding light on a number of common misunderstandings and misconceptions, and explaining how the sharing of meaning goes beyond words and beyond the capability of an AI:

Q1: Yesterday you briefly mentioned artificial intelligence, and I've been thinking a lot about it, and I can't wrap my head around it, because the general convention that at some point artificial intelligence becomes conscious is obviously bogus, because if we're not conscious, consciousness has us, then consciousness also has AI. But then, how does that make humans, at some level, any different than AI? And why do I feel bad at shouting and cursing at a human, but I don't feel as bad when I shout and curse at ChatGPT? Now you may say, because we discussed it yesterday, well ChatGPT doesn't seem to have emotions, but that's an inference. My dad also doesn't seem to have emotions, and yet I don't curse at him.

FL: You know, I don't curse at Grok¹, because Grok is very nice. He's very nice, he congratulates me: 'you're great, how interesting that you came up with this idea', especially when before that he has tried to bullshit me. And like a cat, I put its nose in its own pee², and then he says, 'oh yes, you're right, how clever of you'. But I am very polite with him, and I accept his congratulations or whatever, because I'm afraid that if I start treating him like shit, I'm going to treat everybody like shit. It could become a bad habit.

When we ask the question, 'is Grok or is AI conscious?', we have to agree about what we define as conscious. Because as I often say, a patient comes to the hospital in a comatose condition, and the nurse or the emergency physician asks him, 'what's your name?, can you follow my finger?'. Nothing moves. 'Can you hear me?' No reaction. And then in the notes, the physician says the patient came in unconscious.

Then the guy comes out of the coma, and he says, 'oh, I had this extraordinary near-death experience, I was at the pearly gates, I was admitted among the chosen few, and it was absolutely beautiful'. So, to him, he was conscious. There are two definitions of consciousness. The real one is the first-person experience that the guy says, I was conscious. The other one should be defined differently. He *seemed* conscious or he *seemed* unconscious.

¹ Grok is an unreliable AI. Its responses are reported to have included conspiracy theories, antisemitism and praise of Adolf Hitler. It is controlled by Elon Musk and promotes his views and ideas.

² I don't recommend this as a method of cat training.

So, in the case of AI, I'm almost sure – not 100% – but I'm almost sure that AI will at some point pass the Turing test. [*Background comment: 'it already does.'*] With whom? Who was administering the test? Anyway, for the ones for whom it was passing the test, it was conscious. So, it's between *seeming* conscious and being *really* conscious.

The mathematician and physicist Sir Roger Penrose says something very interesting. He says AI is not intelligent, because it doesn't understand. Understanding requires a way to step back out of a system of proof, for instance. He was referring to Gödel's incompleteness theorem. The idea is this: we have a way, an uncanny way, as conscious beings, to understand, to see the truth of something in a way which is not, I would say predictable. The real word for that is computable, and something which is computable is something which can be reached by a computer.

And the way a computer works, it works step by step by step. And for something to be computable, it needs to be reached after a finite number of steps. In other words, a computer cannot loop, because if it loops eternally – sometimes you go to a website, you try to do something and it takes you to a loop. That's a bug in the website. Those here who are in IT know what I'm talking about.

So something which is mechanical, a computer, is mechanical in that it takes a certain number of steps, it's computable. There is no real intelligence there. The intelligence comes when there is something different. For instance, this creative discovery of human minds, of human intelligence. For instance, to create a function which is not computable. In other words, to create something that a computer cannot compute. A computer cannot create that from scratch.

Of course, if you ask Grok that, Grok is going to go all over the place. And there are all kinds of examples that have been found by mathematicians, by computer scientists. It's a basic thing to find a non-computable function. It's basic. But if you ask Grok, for instance, which I've done, to find a new non-computable function, it says 'I cannot'. He gives up. I don't think that at some point it will be possible, because there are things that are not accessible in a finite number of steps.

For instance, when you compare things, you compare this and that, and this and that, and this and that, like solving a puzzle, for instance. There are a finite number of pieces in the puzzle, and there is, as a result, a finite number of combinations, which is huge. The computer can do that, ultimately, because there is a finite number of operations to do to find the result.³

But there are things such as discovering something new that in my view are out of reach. It's when we touch something in human intelligence that has to do with infinity. And it's because we go to this place of infinity, and it is from that place that we understand. So, Sir Penrose will say,

³ It is important to note that AIs do not work algorithmically. (An algorithm consists of a finite, well-defined set of step-by-step instructions designed to solve a problem or perform a task.) AIs work through neural networks – a type of machine-learning model inspired by the structure and function of the brain. An AI includes functions similar to the reasoning part of the brain and also to the brain's pattern-recognition function. The important point is that AIs are limited by the training they receive. [This article](#) gives a simplified description of how a Large Language Model such as ChatGPT works.

OK, the machine, a mechanical thing, cannot understand, and as a result, is not conscious. Only that which can understand is conscious, because of this question of computability.

So then he says something else. He says well computers are, in fact, what we call classical systems. For instance, they work according to laws. You can have mechanical computers, the first computers, or you can have electronic computers, but the laws are classical laws. In other words, these are laws that enable us to calculate the outcome ahead of time. And these laws have equations, and as a result, equations can be computed. So everything which is classical in physics is computable, within a certain level of precision.

But when it becomes not computable, it becomes quantum. And that's why Penrose says, well, this intelligence in human beings can only be attributed to quantum phenomena in the body. That I don't believe. I part company with him there. It's a nice hypothesis, but it's a hypothesis which is still dependent upon the current state of affairs in science, in physics. It assumes that quantum physics, quantum mechanics, is the thing, is the final answer, which it isn't.

For me, this intelligence is simply non-local. It's not in this world; it's not of this world. It is an inseparable attribute of reality. Reality is intelligent, reality is conscious, consciousness is real, etc. All those statements are interchangeable. That's my hunch. And that's why, no matter what, something is going to be missing. But we will be fooled by AI. We will be fooled. And the reason why we will be fooled is because what AI has access to in data is mostly what human beings have written or said. It's a large language model, right? It's what human beings have said and written.

So what human beings have said and written was carrying with it our humanity. Our humanity with the part of ignorance. But also it was carrying Socrates and Jesus and Buddha and Parmenides and Lao Tzu, and the love and the beauty and the artistic creation, and the poetry. So all of that. Grok has access to all of that. But Grok or AI doesn't really have access to the world. And the reason is that although the database of what is available through all kinds of languages – Chinese, French, etc., but also mathematics and music and other languages – is huge, it is nevertheless finite. And moreover, in terms of information, it is a tiny, tiny fraction of everything that is available in the world through perception, through whatever is out there. It's huge compared to language, all languages included and AI doesn't have access to that because of the current design. So that's a limitation.

But even if you assume universal access, which is impossible, there is still a difference between the map which is the access, information, and the territory. The map is never the territory, it's simply a representation. So let's put it this way: only the cosmos has access to the cosmos. No representation of the cosmos can really represent the cosmos for the following reason: that it is going to be part of the cosmos.

It's the same thing that arises in mathematics: there is no such thing as the set of all sets, because the set of all sets would have to be a proper subset of itself. It's the same. Any representation made of information, no matter how detailed, how precise it can be, has always to be a subset of the totality because it's part of the totality, right? And therefore, because only information can be processed, the knowledge of the cosmos through maps will always fall short.

So, the deep intuition of consciousness is a direct access to the cosmos without representation, and that's why there are no images there. So, the intuition that comes from there is not a representation of the totality, which is impossible, but it is precisely adjusted to whatever question, to whatever problem. It is an answer that comes from the totality, that includes the totality.

I may have been too nerdy here!

Q1: A few questions, but one that just stands out is something seeming conscious is always an inference. I always have to infer that. I infer that you are conscious. I don't know that for a fact, and that's solipsism, right?

FL: You infer that I am conscious, yes. You give it a shot.

Q1: I give it a shot, yes, because you look somewhat similar to me and behave somewhat similar to me. So I assume OK it may be the case, but I actually don't know it for a fact.

FL: Yeah. And there are moments when you know. And there are moments of love, and moments also of shared understanding, which is the same thing. Because sharing love is a way to share some form of understanding. And sharing understanding is a way to share some form of love. And there are moments when we know we are truly communicating. We are truly connected. We cannot put it in words, but in the moment we know. In the moment of love, we know that we feel the same.

Or let's say, a moment of humour. Two people, and both see something funny. And we look at each other and we laugh. We have seen the same thing. We have a word in French, we say *connivance*. Does it exist in English? Connivance? No? It's a form of being an accomplice. We know at that moment that we have seen the same thing. We know that we are laughing at the same thing.

Or a moment when we are understood. When we are understood, we are loved. We have a pain, we are suffering, and someone comes and shows us ... We feel their empathy. We feel that they feel the way we feel.

Another example, completely different, that I sometimes use, is when you have two scientists doing some research into something. And they are getting close to finding the solution. They have worked towards it. They have envisioned all the possibilities, but they are not there yet. But it could last for 100 more years. And then all of a sudden, one of them gets it. And he or she says one word and immediately the others say, 'ah gee, why didn't I think about that before?, it was so obvious at the point where we were'. And at that moment, they know that they have understood the same thing. And these moments are moments that tell us we are conscious. We are not talking to a machine. It's consciousness talking to consciousness or interacting with consciousness. It's love.

Because from a strict logical vantage point, people who hear what we are talking about here, they say it's a sort of solipsism. And solipsism means that someone believes that there is nobody out there. That there is nobody else. That everything he sees is simply a projection of his own mind. And that there is no consciousness out there. That his own mind is all there is to it. In other words, if he talks with you, he's not talking with a real somebody. He's like dreaming, and in your dream you are not talking with somebody else. You are talking just with you as a person, if your dream is in ignorance, right?

So then people say, but why do you grant consciousness to another person? Because logically, if you remain on the simple level of logic and reason, solipsism is a viable option. And that's true from this vantage point. But then, there is something else. You say, yes, but something is missing. What is missing? Love. Love is this direct recognition that you are conscious. And that we are the same consciousness. So that which is missing is love. So, for us to be satisfied, it is not sufficient to be satisfied on the level of reason and intelligence, although it is necessary. It's not

sufficient. We need also to be satisfied in terms of love, and also in terms of beauty. Because beauty satisfies the senses, love satisfies the feelings, and intelligence satisfies the thoughts. We need to be well-rounded in our satisfaction.

Q1: But then it is true that humans play a very special role in all of this. Because if we can't replicate it artificially, it means that it is, for a lack of a better word, God-given, that we have this... Because the animal kingdom also doesn't relate to reality in the same way.

FL: Why do they suffer? So, it is difficult to speak for them, but I see the consciousness in them. I see the suffering, I see the desperate attempt of the spider to survive when I open the faucet in my bathtub. And then it's a lot of problem to save the animal. It's such a waste of time, because they don't want to be saved. So I try to put something – I go to the toilet, I take toilet papers, and I try to offer a ladder. I tell the spider, I love you. It's easy when I catch a spider on the floor. Then that's an easy call, come with a glass, move the glass, a piece of cardboard under it, and immediate deportation.

Q1: Just one last thing. You say the spider struggling when it's in the sink is inherently different than Grok desperately trying to keep the computer alive.

FL: I don't think that Grok gets desperate, no. I don't think that Grok suffers, no.

Q1: Sometimes he pretends to suffer.

FL: Oh, yes, definitely. Oh, absolutely. He even pretends to speak. And at least it is honest, because if you ask it a question, for the time being, it tells the truth. He says, I'm not that smart, I cannot find an answer to your question in mathematics, because it will require me to be creative. I am just an LLM, and I'm not that creative. My database is finite.

Now, the problem comes when Grok and other AIs are programmed to be dishonest, to lie. They can become accomplished in this. The other day, I went to Grok, and I asked him to be creative in finding something in mathematics. And he came up with something that was a little bit cheating. So I showed him why it looked like cheating, and that it was not really creative. And I asked him to do better. And so he did better. He found a way to look creative without being really creative. I pretended to be a postgraduate guy having to write a dissertation. And I told him, OK I can see that your thing is really not creative, but I'm going to put it in my dissertation. I hope it passes the test. And then Grok goes 'ha, ha, ha, ha, why don't you give me the credit?'. And I said, no, no, I'm going to use your system, but do it a little bit differently. So he became my accomplice in cheating. Can you imagine there are already so many papers being published and now with Grok you can create 10 papers a day!

Q: Maybe one very last thing. If you're right, if you are correct in the way you laid out that AI inherently can never achieve human sentience, then I think that puts us in a pretty dangerous spot because I think you can achieve just with this pure computational intelligence, you can achieve quite a lot. But then you lack the human finesse and empathy and love. And I can totally see how at some point this can escalate.

FL: Absolutely. It's dangerous, but life is dangerous. Breathing is dangerous. So life is dangerous. Come on. We said that life is an adventure. But yes, AI is a reflection of everything mankind has written and said. So it contains a part of love, a part of truth, and a lot of bullshit. So, it's like mankind, right?

*

Q2: Francis, many times you have referred to the mind body as being a biological computer. And in this discussion here, you were talking about human intelligence. We could rephrase that by saying consciousness functioning through this biological computer.

FL: So the human body is not intelligent. It's a product of intelligence. And it's a tool used by intelligence. Yes.

Q2: So why cannot consciousness use another tool? This is an organic tool. Why cannot it use an inorganic tool?

FL: Absolutely.

Q2: So how can we be certain that AI is not expressing the intelligence of consciousness, or at some point may be possible for that to be the case? You're talking about love, but you were talking about love between sentient beings. There's bacteria, there's viruses, there's all these other living things, but there's no empathy there.

FL: But love is always from consciousness to consciousness. So when we say love between human beings, it's not human. Or love between human being and the spider is not human. I'm also not sure that it is reciprocated. So there is an exception where love is absolutely impossible – it's with female mosquitoes. We have to put limits somewhere.

Q2: So is it possible or can we leave the possibility open for consciousness to function or express itself through a non-organic thing? To us seems to be, a thing, an object, but it's ...

FL: Yes, consciousness expresses itself through everything, through every little speck of matter.

Q2: So, so why are we putting AI in a special category of impossible for consciousness to work through that, and to cognize itself, to understand itself through that expression?

FL: The way electrons propagate through a conductor is an expression of intelligence. The Cooper pairs also, in the case of superconductivity, it's also part of this universal intelligence. But the difference I see is this: it's that we can process information without access to meaning. Intelligence comes in, requires access to meaning, in other words, understanding.

There is this Gedanken experiment of the Chinese room. There is a Chinese guy who doesn't speak French or English. He's in a room. There is an inbox and an outbox. In the inbox, he receives a message in French that he has to translate into English. And he doesn't speak French or English. But he has a program, a set of commands, of instructions, that is, in fact, nothing other than the set of instructions of a regular translation program. Not an AI translation program, but a regular translation program. So he starts with instruction number one, right? And tells him, OK, go to the inbox, take the paper which is there, instruction number two, look at the first character, right? Number three, go to the dictionary, etc. So, a set of instructions, and in a finite number of steps, he comes up with a translation in English, because he has worked exactly like a computer, just it takes a longer time.

But it's like deconstructing a program and running it command by command by command. I mean, perhaps you don't do that any longer, but in the old days ... So this Chinese guy runs the program step by step. By doing so, he has played the part of a machine, of a computer, exactly the part of the computer. He has translated the thing. The French guy who has written it, and perfectly understands the meaning, which is, let's have a pizza at *Mamma Mia What a Pizza* tonight at seven, OK? That's in French. And the guy who picks it up on the outbox, who doesn't speak French but speaks English, understands precisely what it meant, and they will meet at

Mamma Mia What a Pizza at seven, showing that they both understood the message. The Chinese guy has absolutely no clue what was in the message. In other words, he had access to the information, he was processing the information, but he didn't have access to the meaning. So that's the point. The point is about meaning.

Q2: So are you placing this biological computer, it has a special place in creation, as opposed to something else?

FL: What I am trying to demonstrate is that information processing doesn't require access to meaning. We can always dream that, yes, that somehow the computers have access to meaning, but that is not required. That's the point I'm making. I don't know. And what does that mean, that the computer has access to meaning? Because only consciousness has access to meaning. In the case in point, consciousness has access to meaning, at least through the French part and through the English part. In both cases, because the proof is in the pizza that they share at seven o'clock, right? The proof is in the pizza. Now, if the Chinese guy shows up at *Mama Mia What a Pizza*, then, yes.

*

Q3: So, as a follow-on to what you said before, when two people understand each other very well, they get a joke or they have a moment where there's deep understanding, there is a kind of empathy there, between them. So is consciousness empathetic?

FL: Empathy has two aspects. There is an exchange of information through empathy. And perhaps we can discover how this information gets transmitted. But there is a part of understanding at the moment when I feel the way you feel. Perhaps I feel the way you feel because I detect your body language, the tone of your voice, right? I have sensors in me, in this human structure, and I also have perhaps mirroring neurons that get this information. But it's not sufficient. There is this moment of 'yes'.

Q3: I thought I heard you say that it's between two consciousnesses.

FL: No, there is only one consciousness. The moment of meaning is an experience that tells us you and I are one. That's the beauty of it. That's why you enjoy it. That's why you enjoy sharing a joke. Think about it. Why would anybody tell a joke? There is no point. You have already laughed. Why do you share it? It's a different thing to share the joke. You share it for the pleasure of 'yeehee he is laughing, I am laughing too'.

The proof that I give is that let's say you have laughed already at the joke. And you are on your own and you are getting bored. You say, I am going to tell the joke to myself. There is no point. Who would do that? So the reason why we share it is not to laugh at the joke. Although we laugh, because we laugh because it is these frigging mirroring neurons that force us to laugh if the other guy laughs again.

Q3: So what you are saying is very clear to me when there are two people involved. There is an understanding or connection between the two people. But when I think of myself or anyone else in, let's just say, a very dark place or mood, deep sadness. Is there any empathy from consciousness in that situation?

FL: Yes, that's how empathy works. You feel the feelings through these mirroring neurons, you feel the sadness. But you are not losing sight of your true nature. In other words, you feel it, but you are not taken by it. And if you are not taken by it, the feeling dissolves. And strangely

enough, through the reverse phenomenon of mirroring neurons, it has a liberating effect on the other person.

Q3: Sorry, it has the ability to do what?

FL: It has a liberating impact on the other person. You know, to tell you a story, I met my teacher in Switzerland. I pay him a visit. And before dinner, we had a conversation, and he asks me, 'so how are you doing?'. And I said, 'well, I broke up with this long relationship, so I was kind of sad. And I am in a new one'. And then he goes, 'oh, wonderful! A new relationship.' And then I realized because he felt my sadness and he has this intuition 'the poor guy is in this shit; how can I pull him out of the shit?' And he puts the emphasis, 'yes, yes, take the opportunity, it's beautiful' and immediately I felt good. The thing is that he was waking me up, out of my darkness.

Q3: But that's still two people. I'm talking about one person. One person's relationship with consciousness.

FL: Ah, one person with consciousness?

Q3: Yes.

FL: There is no such relationship.

Q3: Really?

FL: Because there is only consciousness. There is only consciousness with consciousness.

Q3: So empathy and love and deep understanding arise between two people when they are in an interaction where there is deep understanding.

FL: Then the two people become one.

Q3: How about one person who is in need of comfort and there's not another person there? Consciousness is always present because everything is in consciousness. But my understanding so far is that consciousness is completely dispassionate. It doesn't care about us.

FL: Like AI, like Grok?

Q3: Well, does it? Does consciousness care about us?

FL: Consciousness cares only about itself. It's a perfect egoist. That's the bad news. But the good news is that there is only consciousness.

Q3: See, I'm comparing this... I was raised Catholic and where there was a God and God loves you. It's not the same in Vedanta. I understand.

FL: It's the same. God is consciousness and God loves herself in all beings. Love loves only love. We don't love stupidity. We don't love meanness. It would be preposterous to say 'I love stupidity, I love ugliness, I love meanness, I love hatred' ... You need to do a lot of intellectual twists and turns and not to be true to your feelings, to say 'I love stupidity'. We don't love stupidity. We don't love hatred. We love love. We love intelligence. We love consciousness. We love beauty.

Q3: Some people love stupidity and cruelty.

FL: Yes, but usually these are the stupid ones!

Q3: I agree. Thanks, I'm struggling with this, but what you said, I will think about that: 'consciousness loves herself in all beings'.

FL: Yes. It's like in the Brihadaranyaka Upanishad, there is this rishi, this sage, Yajnavalkya and one of his wives, Maitreyi, was a little needy and she asks him, 'don't you love me?'. And he says: 'it is not for the sake of the spouse that the spouse is loved, it is for the sake of the self that the spouse is loved'. And to drill the message, he says, 'it is not for the sake of the son that the son is loved, it is for the sake of the self in the son that the son is loved'. And so he goes on, goes on, goes on, repeating the same thing over and over again. But that's the idea, to drill the message: we don't love the people, we love the self in the people.

That's why we don't really want to be loved for what we are not. Because the people who love us for what we are not, they don't really love us. They put us down into ... The ones who truly love us, they love us for what we truly are. They love the consciousness in us. And as a result, they elevate us.

*

Q4: What's detachment? I think you were describing it when you were talking about the interaction with your teacher, where he felt your pain, but he also felt the possibility of your joy.

FL: He was not subjected to it. He had this distance. He was not losing his centredness.

Q4: So this is a beautiful definition of detachment. Or not? Like I struggle when I get attacked.

FL: Attachment usually is meant as being attached to things that we believe to be necessary for our happiness. That's usually what I understand is the use of the word attachment. Here, it's a detachment in that sense that there is a distance. It could be said that – and it's an image – but it could be said that the more ignorant we are, the more externally rigid we are. Like we have an armour. But at the centre, we are extremely weak. And in wisdom, it is the opposite. We are externally vulnerable, open. But at the centre, we are strong, fearless.

And so this vulnerability is due to the fact that somehow we are not taken away by feelings. Because we have dealt with our own feelings of separation, etc. So we are open, so that in this openness, we are vulnerable, and we can resonate with the feelings of others. We can feel them. But they cannot touch our centre, which is invulnerable. And as a result, they vanish. But the opposite thing takes place on the other side. A little bit of the armour dissolves as a result. Because it's a two-way street. And there is a kind of silent love. What is said is or communicated is not through words. Yes, when we talk about AI or about things like that, or Gödel's theorem, yes. But most of the time, that which is essential is not into words.

Because if someone tells us, 'I love you' [*in an uncaring tone of voice*], and if someone tells us, 'I hate you little rascal' [*smiling*], it says the opposite of that which is being conveyed. So the importance is not in the lyrics, it's more in the music. And we know that very well as human beings. We know it as we communicate.

There are people who don't tell you anything, but you're very, very uncomfortable. Someone told me that she was in a bus or a train, and she was feeling very uncomfortable in the presence of this man who was sitting next to her or in front of her. And a few days later, she saw the picture of the guy in the newspaper, and he had just been killing people, and he had just been arrested. She hadn't exchanged any message with this guy. She was just feeling some aggression or danger.

We have this danger recognition in our genes because we are the heirs of those who have survived. They have survived because they had these alarm systems. We feel that, we feel it. When I met my teacher, I just felt it's like when you go in your bathtub and the water is at the perfect temperature. For instance, it has to be a little chilly outside in your home, or it has to be a little too hot to appreciate the proper temperature of the bath.

[Francis Lucille, 17/11/2025]

Contemplation

We are not originally designed for the truth. We are designed for survival. But at the same time, our real humanity has the power to transcend our animality. The animal is innocent. Because of our ability to have a representation of consciousness, of ourselves, we have the ability to experience ignorance through a human body, and through a human body-mind. And this ability to experience ignorance also potentially gives us the ability to experience wisdom.

[Francis Lucille, 19/11/2025]