Ignorance Podcast

[00:00:00]

Mayla: Let's welcome our guest scientist for today. Please introduce yourself by name only and your favorite thing to do outside of work.

Larry: my name is Larry Hunter. I have lots of favorite things I like to do outside of work, but I swim in the ocean every day and that's definitely one of them.

Mayla: Nice, thanks. Now that we know that you're kind of a human person, not just this amorphous scientist, let's get down to some more questions. What questions excite you in your work? What questions are you currently pursuing?

Larry: Well, I'm, uh, I'm lucky that I'm getting to start on a new project that's a ambitious one.

it's really about, uh, a question I have been curious about for decades. the question is, where do our desires come from? Where do our wants, our wishes, our fears?

Our desires are not the same as the desires of a, a bug or a bacterium, but I think bugs and bacterias do want things And so how does that work? [00:01:00] And, Are there other kinds of things in the universe that might want something? What would that mean? what's the physics, if you will, of wanting something?

And then you could also look at it from sort of a, an AI perspective.

is it possible for an artificial intelligence thing to want anything? What would that mean? And then there's this sort of more. Personal growth, psychology piece of it.

What do I want out of life? What would really make me happy? What should I be doing? what do I wish I wanted? Or, gee, I wish I didn't want that. How do human beings come to want what they want? And how does it change through our lives and how do we get new ones and lose old ones?

And so there's a lot of angles, uh, to this, to this question. I.

Mayla: Yeah. And I would love to know some of those answers, right from my own perspective of how do I know what I want? Right? Even just personally, I think it'd be interesting to answer that for myself too. And maybe even for yourself also.

and I think it has huge impacts as you're exactly saying, [00:02:00] right? On many, many different fields, um, which is really, really cool. Um, my next question is how do you hope to impact the world um, so that's kind of a nice segue, that this has a lot of impact

and so how do you want to do that?

Larry: You know, one of the problems with sciences is until you figure out, you know, discover or invent. what you intend to, then it's a little hard to know what the impact or consequences are going to be.

And sometimes they surprise us. Um, but I think there's a bunch of places, where you could imagine if we had some answers to those questions. So maybe we could build something that actually wanted. an artificial intelligence that had its own desires and fears or whatever.

maybe we could build physical systems. active droplets is a thing that physicists and material scientists are, are really interested in right now. How could we make, thermodynamic system out of parts that we understand that wanted something that acted like it cared about something? And we [00:03:00] could also, you know, go down the personal growth path.

we're social creatures. We'd live in a social system where there's status and we would, you know, we look to people with high status about what we should want, and we also wanna climb up that status hierarchy.

AndUm, we want love, we want all kinds of things that, um. Come from others. And so some of our desires are socially transmitted, like our other ideas. And so understanding that better, might help us be happier human beings. But the implications of answering these questions are not always positive.

It might also risk, you know, better kinds of disinformation and you know, more convincing advertising

Um, so, the implications are vast as of the questions.

But for a lot of basic science, I think it's important, and this is a very basic science question, I think, uh, not to be too driven by the potential impacts.

Mayla: It's [00:04:00] interesting, right? This idea of impact in general, that you answered it very specific to kind of your questions and your research direction. It could also be answered in relation to you as a person, right? So just to call that out. And also I definitely implanted in your brain the desire to want to do my podcast.

So it's already there, right? Just we that all the time. Yeah, exactly. So I just wanted to call that out as a really fun example. And I think it's really important that we both think long term and short term, right? And this idea of application is a really interesting question that is definitely still talked about right now around research of how, what should we be funding, right? Basic science, applied science, and all these different pieces. So really, really apt question to bring up for sure.

Larry: I'd like to talk about that a little bit because one of the reasons I feel so lucky you, you know, how does this, how does this question impact my own personal life? This is the culmination of, of my entire career in some ways.

So I started off, you know, in, in graduate school and I was a, you know, baby scientist, worrying about, uh. The implications of the [00:05:00] questions. I was interested in goals for knowledge, um, and how do we come to have particular desires.

and I've advanced to the point, um, where now I'm a tenured full professor. And I recently changed institutions and was handed a pretty good startup package, um, to advance my research. And so I'm in an incredibly lucky position that very few scientists ever get to, which is I have a fair amount of money to go study whatever the heck I want.

I feel so much right now that this is my moment and I really want to take full advantage of it and make the best progress I can with the, this opportunity.

Mayla: The goal is really about the question, right? Making progress in those questions and starting to answer them for some impact in some way, whatever that may potentially be, whether it's on you, the world, different things.

Larry: The main impact, if you will, that I'm after, um, is understanding,

how it is that things can matter to us, uh, is one of the most fundamental questions in the universe. It's right up there with the origin of [00:06:00] life and the nature of the universe. It's, I mean, it's a really fundamental question, um, especially about human beings.

Mayla: who has influenced you, Who are people you look up to?

Larry: That's a hard question. It's a, a huge list. there is a ton of work from scientists in many areas, um, that, that are relevant to this.

I've been reading, you know, statistical mechanics to try to understand people who have quantitative theories of, of, uh, how of semantics, of how something can matter to a physical system.

so Dave Walpert and, and, uh, Artie Ksky, wrote this paper and, and, and statistical mechanics that I'm particularly enamored of. And, uh. Changed the way I thought about the world. And, Lisa Barrett Feldman at the other end of the spectrum, um, is uh, one of the first people to really pursue what is now being called affective science.

Looking at emotions, uh, with the scientific [00:07:00] perspective and, uh, her contributions. There are many. are really about the role of introception of our own, uh, nervous system's, perception of what's going on inside us. Um, and that, uh, gets character labeled culturally with words and characterized as, as emotions.

There's something about knowing about how we are doing ourselves, our bodies, um, that it plays a critical role in what we want.

If I can't breathe, um, there's a signal in my body that makes me very desperately want to breathe. And, um. That's, that's kind of an obvious one. Um, but, uh, I think, um. Feldman's characterizations of how emotions work or, uh, a good way to, a good intro to that is her book, how Emotions Are Made

So, you know, Barett Feldman would say, uh, you know, we, the, the classic story is, [00:08:00] you know, you, you get afraid and therefore you get goosebumps in sweat or whatever. And she would turn that around and say, you get goosebumps in sweat, and then your brain says, oh, I'm afraid.That Introception is a, is a huge through line and, and her work.

And now there's a lot of work at NIH is having an interception workshop. It's a, it's a, a hot new topic in neuroscience. Um, so there's a, there's a whole bunch of neuroscientists who are doing this who are really important, inspirations, um. And, uh, I would be remiss, um, to not mention my mother in this, uh, collection of people.

so my mother was always about, do what feels right to you. And she's, she's been gone for many decades, but I still have her in my head as one of the, the, the things that cause me to desire or, you know, look ance at some of my desires and whatever this, this funny meta thing that goes on.

But, um. She would love this.the voice in my head of my mother says, Larry, you are doing the right thing. and you know, my partner really drove me to this in a way because I was having a bit of a, uh, I'd like to call it a [00:09:00] midlife crisis, but that makes me younger than I actually am, or I'm gonna have to live to very old, uh, worried about what I was gonna do next.

Um, and she was like, forget all the noise, just. Think about what, what is it that you really want? Right? And I really sat with it for a while and I, I tried to figure out, okay, what's, you know, I have so much life left.

What, what is it that really does matter to me? Um, and so it. Really with her help got me here.

Mayla: Yeah, it reminds me, uh, similar of like how I kind of move through my career in different ways, as you know. And, um, I've realized that I kind of study what I don't love or what makes me uncomfortable, Um, so I studied the unknowns, right? Being the unknowns that made me uncomfortable. And your unknown in this case was, what do I wanna do next? Right? What's my desire? And so as you said, like on the meta level, it's very cool that we're kind of studying these things that also both help us, right, have this impact on us and hopefully also impact on other things, right? By being understanding or other pieces that you've already mentioned. Something that also can [00:10:00] drive us is success, And so a question I would love to ask you is how do you know if you've been successful at your work? And as a scientist, what is your evaluation criteria?

Larry: Oh, wow. Great question. Different evaluations for different kinds of science. so one of the really hardest things about, uh, pursuing this question around, you know, how things can come to matter is trying to figure out empirical tests of theories. How can we actually get some evidence that would say yay or nay?

Because it's really easy to do armchair theorizing,

I, I should be careful here. It's important. Impossible to do good psychology. Alright? I don't want to, to, to call into question. The way we study people by studying their behaviors and reaction times and the content of their utterances. And there's lots of of good science that can be done that way.

But for this particular question, I think that's gonna be hard. You know, I think that question isn't separate from doing the science. It's part of doing the science, [00:11:00] how it is that you're gonna figure out whether you're, you know, Larry's crazy pet theory, um, really, uh, has some evidence to support it

And, there's a lot of different angles you have to come at, uh, an empirical question in order to be able to really demonstrate it. Right? So we have our gold standards before I'm going to be allowed to market a drug that I'm gonna claim is gonna cure you of something where you have to run it through a clinical trial.

And that's a very high standard. And, you know, for a lot of of theories. Particularly psychological theories. it's very difficult to, to have a reasonable, uh, double blind, placebo controlled kind of clinical trial not available to us. So what else can we do?

We have to come at it multiple from multiple different angles and evaluate different ways, and hopefully by evaluating enough different ways that are okay, we gain confidence in, in the claims. Um, but being creative about how [00:12:00] to evaluate them is, is an important part of being scientist.

Mayla: Yeah, so it's hard to answer, right? I kind of asked this question intentionally because as you alluded to, it's really a part of the question, right? That whenever we ask a question, we need to also think about how do we know we've been successful at that question? And part of the way that I'm doing this, right and studying scientists is getting to talk to them and see kind of how they answer these questions and how they go through themselves and think through these questions that I think really showcase what science really is about.

Right. Thinking about these types of things, whether that's the question, that's the evaluation, all the different people that we're impacting on all those different pieces. So what has been your most surprising result?

Larry: Yikes. Um, well, so I don't have a ton of results. I'm relatively new at this particular question.

I think the thing that surprised me the most is this introception through line that I mentioned to you that this, um, the, the physics version, of how, uh, uh. Thermodynamic far from equilibrium system can want something. Um, it [00:13:00] depends on, the equations, um, depend on a viability function.

What's the probability that this system will continue in a low entropy state and will not dissolve? Uh, for people who aren't statistical physicists? The low entropy state is organized in functioning, and, uh, continuing in that state means continuing to exist basically, as opposed to turning into.

So even the physical system has a representation of the viability of that physical system. The, you know, theories of human emotions, um, play how interoception and our, our, our body's sensation of itself, um, is, plays a critical role, in how it prioritizes the things it wants or, or doesn't want.

that a sensation of your internal state, um, is fundamental to caring or anything mattering to you at all. and ultimately to your most sophisticated, you know, wishes [00:14:00] and, and fantasies and desires, uh, I think was a really surprising thing to.

Mayla: Yeah. And this, I love this question because I think something that scientists also look for are those surprises, right? We wanna find the things that we may not have expected, and that tends to actually drive us to continue to ask these questions and continue to do that work. Um, so I hope you continue to find some surprising results, That leads you to kind of new angles. Um, that allows you to kind of continue and build out this work further.

Larry: Before you get to your next question. I wanna comment on that. So, I do think, um, that surprise is, is a really strong indicator that you've discovered something, um, even if other people knew it first.

But I don't think surprise is the only way that new questions come. Um, sometimes answering one question, um, actually opens the door to another one. And, uh, there's a, uh, another scientist who's influential much more in my early career, a guy named. Uh, who, uh, I went to grad school with and we taught each other all, all kinds of things.[00:15:00]

Um, and, uh, including how to cook each other's cultural cuisines and stuff. It was hilarious. Um, but one of the things that, uh, his science did, and this is, you know, early AI a long time ago, um, was look at natural language understanding as a kind of question transformation. You come in with a set of questions and

then, uh, you process the text and the result isn't necessarily answered so much as new questions.

and so I've, I've had that model in my head of how scientists work is as I make progress on understanding things, even if it's not terribly surprising, that progress almost always raises new questions so often about how to evaluate it or what it would mean if I translated that into this other area, right?

If I. Understood something about interception and, you know, metabolism, bacteria, does that mean anything about human feelings? Um, that kind of stuff.

Mayla: Yeah. So I agree with you. Surprise is not the only way. It is one way. And kind of my next question actually is that some other ways that this can happen too. [00:16:00] So my next question is, what has your favorite disagreement? Um, 'cause just as surprises, disagreements can also lead us to ask new questions and come up with new theories, um, and new paradigms to think about as well.

Larry: Yeah, that's funny. Uh, so I, I find myself, uh, I'd like to learn from people disagreeing with me, but I haven't made enough progress in this particular research yet to get very many people to engage with me, let alone disagree with me in any interesting way

I, I do have disagreements with people in the field, and that sharpens, um, some of my thinking about things. So, for example, a really important piece of artificial intelligence is reinforcement learning.

And Reinforcement learning is entirely driven by reward. And my argument with that is not. Against reward, but it, how that is ducking the question. Alright. It's not answering which rewards.and, uh, where do they come from? they're always exogenous. There always comes from [00:17:00] outside the system. Um, so reinforcement learning requires, you know, a, a set of actions that you could take and a reward function and a set of senses that you have to.

They what the state of the world is with respect to your actions and your rewards. but making the reward function outside of the theory to me begs the important question of where do those desires come from and how are they dynamic?

Right? A reward function is a fixed thing, and so one of the things that's obvious. Uh, if you start thinking about this even a little bit, is our desires and the things that matter to us, what we care about, changes over time.

so we can't just imagine there's this fixed reward function out there. Um, we're sort of discovering what to want.

Mayla: Yeah. A book, um, that reminds me of is the neuroscience of You, I think Chantel Pratt.

Um, where I was curious on the updated neuroscience, right? That a lot of our AI comes from neural theory. she also specifically talks about. Reward systems. And it is really interesting thinking about AI versus the human right. That the human can update them through [00:18:00] these desires.

And how do we potentially, if we wanna continue to create AI that has more desires and does the things we potentially desire, right? More, how do we kind of. Change that reinforcement learning and that reward system and make it dynamic as exactly your thinking. Um, and so these are questions I've been thinking about actually a lot too, around how do we influence each other, um, around collaborations and science and all these different aspects that kind of exactly you're asking too. Uh, which is fun. Uh, if you could do anything, not limited by grants, other factors in your space, what would you be doing? What is your vision for the world?

Larry: I am, I am one of the luckiest scientists I know. Okay. I really, I have that and you know, you're right. What I've been telling you for the last half hour is exactly what I would do if I could do anything I wanted. And I didn't have to worry about getting promoted or tenure, and I didn't have to worry about bringing in money for it. I didn't have to worry about anything. I just had the opportunity to, to pursue that question. And this is the question I picked.

Mayla: Yeah, it's great. That kind of ends the [00:19:00] time, and so I'd love for you to reintroduce yourself in an out introduction, giving us your job title in the field, and anything else you'd like to say.

Larry: My name is Larry Hunter. Um, and I was recently appointed a professor of pediatrics at the University of Chicago, but uh, in the, um, there's a section in pediatrics on biomedical informatics, and that's really what I've been doing, in a lot of my career has been,

worrying about how to apply computational tools to advance human health and thriving. Um, a lot of my career has been, uh, around trying to find, uh, artificial

intelligence approaches to those questions, particularly around, uh, hypothesizing mechanisms of action.

My original training was in computer science, so I got a PhD in artificial intelligence in 19 88, 89, which is right when there was a big AI bust.

We called it AI winter in those days. And so, uh, my newly minted degree was really useless.

And so I ended up turning to, um, NIH, um, the [00:20:00] National Institutes of Health, where the Human Genome Project was just getting off the ground. Um, and so I started something called machine learning project at the National Library of Medicine.

Um, and that launched me down this path into, into biomedical applications. But I've always been interested in, in, you know, sort of fundamental theories and, and methods for ai. And so, uh, this is in a way this question about. Desires and mattering and so forth is, uh, getting back to, uh, all of our work.

So, um, you know, some early work I did with Ashwin Rahm result in a book called Goal-Driven Learning. A very good example of what we're about now. I wanna know where the goals come from, what, how did those happen?

Mayla: Uh, thank you for talking to us about your questions and I'll also just add to your introduction that Larry is my thesis advisor for my PhD.

And a lot of these things really allowed me to, um, explore my questions that I wanted to, right, which was about the questions. Where do the questions come from, um, that we ask, and how do we find good questions and ask good questions? And how do we connect [00:21:00] questions at large? So this podcast really is also influenced by a lot of that.

So I'll say thank you Larry, uh, for all of those influences as well. And thank you for, um, agreeing to do my podcast and get interviewed today. I really appreciate it.

Larry: It's been a real pleasure, uh, both doing the podcast and seeing the amazing work that you've done. Uh, in looking at the question scientists ask.