Hyping AI Hinders Innovation

https://mindmatters.ai/podcast/**ep163**/

Announcer:

Greetings, and welcome to Mind Matters News. Erik Larson, a computer scientist, and founder of two DARPA-funded artificial intelligence startups, has been investigating and testing the boundaries of AI. In his new book, published by Harvard University Press entitled The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do, Larson argues that hyping AI to the point of equating it with human intelligence is extremely detrimental to longterm innovation. Why? Well, our guest host, Andrew McDiarmid, senior fellow at the Discovery Institute and host of his own podcast Simply Scottish, chats with Larson about this and others today on Mind Matters News.

Andrew McDiarmid:

All right. So I'm here with Erik Larson. He is not only a wonderful man and great thinker but he's now the author of The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do, and what's your publisher, is it Harvard?

Erik Larson:

Yeah, Harvard University Press.

Andrew McDiarmid:

Okay, all right. So this book just came out and it's making a bit of a splash. Hopefully we can make the splash bigger, but basically Erik, I just wanted to chat with you, ask you some questions about it. I'll start with some honesty, I haven't read the whole thing yet, just because I've been crazy busy with our friend Dr. [Stephen Meyer 00:01:24] and his new book, but I'm fully planning on reading it. I have some vacation time coming up so I'm really looking forward to it because I think you have hit on something that is very valuable for us to look at right now and that is what path are we on with AI? Are we on the right path? Is it going to lead us to the things that are promising? Or should we go back to the drawing board?

Andrew McDiarmid:

So let me just ask you a few questions and that in fact is my first question to you. You're calling for us with this book to go back to the drawing board with AI research and development. Can you paint a picture first for us of what the AI landscape looks like today and why it's not heading in the right direction?

Erik Larson:

Yeah, I mean, basically starting around the year 2000, AI went from the old way of doing things to a kind of data driven way. Basically we're dealing now with what I call big data AI, which is basically AI works the best when you have massive data sets, and that has made certain things like image recognition, face recognition on Facebook for instance, personalizing newsfeeds and things like this has made those things, put those on steroids, but it's also resulted in AI being done a certain way which in the book I

describe as induction basically. So the old way of doing AI, we would write rules. I actually came out of this camp and then ended up working in the modern AI but I started in the field around the year ... Actually exactly the year 2000, and at that time, it was just taking off. Google was pretty much unknown except for little pockets in California and there was no Facebook, there was no Web 2.0 in the year 2000, and AI was kind of in one of these winter periods where people had lost a little bit of confidence, funding dried up and so on.

Erik Larson:

So at that time, 20 years ago when I started, the field was still trying to basically write rules or use what I call in the book a deductive approach, right? So deduction in philosophy is well-understood, the classic example is all men are mortal, Socrates is a man, so Socrates is mortal. So it's a way of specifying premises and then reaching a conclusion and those inferences are rule-based because you actually write or specify the knowledge and then you reach the conclusion that way.

Erik Larson:

So there was still, the first company I worked at, we were still using that method. In fact, it's a famous AI company in Austin, Texas not far from where I am here. But the web had taken off a few years earlier, I mean I think the first commercial company on the web was 1994 but the traffic really started to accelerate in terms of the growth of webpages on the World Wide Web in the latter part, just 1995 to 2000 basically in that part of the decade, the web had this exponential page growth and so all of a sudden there was all this data and so the old methods, actually a lot of the AI that we use today in modern AI, I call it big data AI, a lot of those methods, machine learning methods, deep learning is one of the current methods that's really popular. Those methods actually have been around for a long time. Deep learning is based on something called the neural network and a neural network actually is one of the older machine learning algorithms in AA. I think the original neural networks were called perceptrons and they actually emerged in the 1950s pretty much at the inception of the field itself.

Erik Larson:

So the actual algorithms, the learning algorithms weren't new going into this next decade in the 2000s, but the amount of data available to feed them was just exponentially growing because web pages are basically text and images, right?

Andrew McDiarmid:

Yeah.

Erik Larson:

So you see ... When you think about AI today, you think about the big companies like Facebook uses AI, they hire some of the best AI people in the world to run their labs. Google uses AI, everybody's using this buzzword AI. What they're using are these machine learning algorithms that are actually very old and they've gotten certain modifications. Like deep learning for instance is basically a stacked neural network, so you take more than one neural network and you stack them on top of each other to create a deep hidden structure for the learning algorithm.

Erik Larson:

So yeah, we ended up with basically these old ... What we used to call empirical techniques or learning techniques with 10x, 100x, 1,000x, 1,000,000xs amount of data and all of a sudden they were doing

things that they wouldn't do before. So there was a period where the AI community kind of slowly and then really all at once dropped the rule-based approaches and adopted the empirical methods or the machine learning approaches. Yeah. So the strength is is that you can do a lot with data. The weakness is that you need a lot of data to do anything basically, right?

Andrew McDiarmid:

Yeah. The switch from deductive to inductive was basically to deal with the amounts of data that we were acquiring?

Erik Larson:

Basically yes. Empirical approaches or the inductive approaches just started to show immense promise because there were huge data sets to feed them.

Andrew McDiarmid:

Okay, so was the next progression -

Erik Larson:

Yeah, so I mean ... I'll just give you an example, say on Facebook if you want to personalize your newsfeed, that means that basically it's going to be looking at what you're looking at, what news you're clicking on, and so by prior example, by looking at lots of prior examples, it comes up with what's called a model, a predictive model, of what you're likely to want in the future. So that's the AI system that's actually personalizing your news, and what it's doing is inductively it's looking at what you will prefer, what your preference will be in the future. That's what's happening, and that's basically how all of the AI today works on the web. Like you could go to other examples but that's going to be the basic model for how AI works today, yeah.

Erik Larson:

So to answer your question, they switched to ... So the AI community, we scientists, all of us, I did switch to this as well. My field is in natural language processing, information extraction. So I deal with text, how to extract information from text, how to classify text. Everyone sort of went from this old way of doing AI to this data hungry way of doing AI because those methods worked better basically. It was to deal with the amount of text on the web to make sense of it, but it was also just because as an AI scientist, all of a sudden those were the methods that you had that really were showing the best results in terms of accuracy and so on. That happened ... I mean in the last 20 years, that's completely dominated AI, yeah.

Andrew McDiarmid:

Okay, so we have this landscape with largely inductive reasoning and inference. We have tons of data, but we're nowhere close to the general intelligence that we're being promised by futurists. So we'll touch on that as we go. Another basic question, how did you get interested in AI? Did you always look like you were going to study that in school and have that as a job?

Erik Larson:

Yeah, I mean I started out studying philosophy and math actually. Those were my two majors as an undergraduate, and I actually didn't take ... Like really take very much computer science as an undergraduate. I think I took a couple of courses in what are now almost defunct languages like Pascal

or something which kind of ages me I guess. Pascal is such an old, old language. But I took a couple courses as part of my math major but I didn't really ... I didn't study AI or I mean computer science per se and even in graduate school, I started in a PhD program in philosophy. And then I found myself getting interested in the core sort of philosophy of AI issues, like can a machine actually show, simulate or reproduce actual human intelligence? What are the limits of machines? And there's all kinds of interesting explorations and inroads into those questions in math and in computer science and in philosophy.

Erik Larson:

So I started really focusing on the philosophy of AI as a PhD student in philosophy, and then at some point ... I mean frankly, the true answer is is that my wife got pregnant with Brooke, our first child who is now 21, and I was in a philosophy program making pennies teaching as a teaching assistant and so I basically taught myself how to program and ended up getting a job which was a huge break for me in what's really a famous ... It's an old-fashioned AI company, but it's actually still around. It's called Cycorp if anybody knows it, but I ended up getting a job at Cycorp and actually transitioning from an academic into an actual professional computer scientist and so on, so ...

Erik Larson:

Although there are aspects of computer science where you're just basically ... You're an engineer, right? It's sort of like building a bridge, it's not so much theory as it is you're actually just building systems, right so [inaudible 00:11:56] ... I got interested in it through philosophy is the short answer.

Andrew McDiarmid:

So if you hadn't taken the philosophy classes, you might be a millionaire now, because you'd have the math, you'd get into the tech and you'd make millions. But you took the philosophy, you decided to think about it as well. Well in one of your podcasts I looked up your -

Erik Larson:

Well I did start two companies and sell them, so -

Andrew McDiarmid:

Oh. Well there you go. There you go. That's right, you are a tech entrepreneur as well, and that is cool. That's what's cool. One of the cool things about your voice is you're not just, "Hey, we should watch out for this stuff." You're working in it, you're innovating in it, you're not afraid of what's to come, but you want to caution people and you want to make sure we're on the right path and doing it the right way.

Andrew McDiarmid:

I found some of your podcasts on the web and in one of them you say, "I don't know what a mind is, but I know what a machine is." So how does that assertion illuminate where we are on the AI debate?

Erik Larson:

I get this question a lot and it's an understandable question. But people sort of ask, "What's your view about minds and machines?" And those kind of philosophy questions, I deliberately avoided asking those in the book because as a computer scientist, I wanted to write an argument for computer scientists that was also understandable by the broader public. Like when I say I don't know what a mind

is, what I'm saying is is that I can separate that part of my thinking from the problems that we're having in AI and the progress that we're making or not making. So I think the core issue in AI is that something that minds do that we know that they do which is basically hypothesis generation, there's a type of logic called abduction. That type of thinking that we do, this abductive inference, which I can explain if you want later, but remember how I was saying that big data AI is inductive, inductive inferenced? Inferenced is just from what I know and what I see around me, what can I conclude, inferenced. And it has a kind of formal understanding in computer science. But there's this type of inference that minds do, that people do, that we can't figure out how to program.

Erik Larson:

So the book was really an exposition of here are the three known types of inference, deduction, induction and abduction. Here is basically the entire field of AI since its inception in 1956, divided between deductive attempts, old-fashioned AI, good old-fashioned AI, it's been referred to, and inductive attempts which have basically dominated the last 20 years as I was saying, and there's been essentially no progress, other than some token progress using some [inaudible 00:14:52] and so on but really no progress on abduction and abduction is common sense thinking. It's what we do when we have a conversation like we're having now. It's what we do when we navigate in dynamic environments, like we go buy a gallon of milk at the store. That's just completely absent from the work that we do as computer scientists. In fact it's not even that it's absent and we haven't gotten to it yet, it's that ... I could tell you this as someone who builds and designs these systems, nobody even knows. Like nobody has a clue of how to do that, to do abduction.

Erik Larson:

So there is this very interesting distinction that we can see when you look through the lens of inference that people are doing things cognitively that machines just aren't capable of doing. So I don't know, and I don't tend to ... I'm not the guy who's going to start in saying so we're dualists, so we're this, so we're that, right? Like I got off the boat with that kind of ... Certainly publicly I don't talk about that stuff in terms of what I'm doing with my book, but just even in my personal life, I don't have extremely strong views about what's going on with the mind, whether it's a Cartesian model or it's some kind of other model or something. What I do know is is that our attempts to make the mind ... To basically program the mind in computers have failed and there is a very strong principled reason that they have failed. It does seem like we're dealing with two different things.

Andrew McDiarmid:

Right. Keeping mind and machine separate in our minds is I think one of the first steps to really having the respect to even get to the right questions in this debate. So part one of your book is called The Simplified World, where you explain how AI culture has simplified things about humans, and also maybe overextended ideas about technology and you say it started with Alan Turing and just some of the extension he did on the subject. Has it really come so far up to today? Are we still doing what he did? Like assuming more than we should?

Erik Larson:

So his idea was that the intelligence reduces to problem-solving basically and he ... In the book, I talk about how he seems to have undergone a fairly fundamental change. In his earlier work he talked about the distinction in mathematics between ingenuity and insight, and insight was something he said that mathematicians use that's outside of the formal system that they're working with to decide on what

parts of the system are interesting to think about, so when some mathematician comes up with a new proof or he comes up with some interesting development of some aspect of mathematics, Turing originally was saying that that's a non-mechanical feature of the mathematician. So part of their job involved this insight which was separate from ingenuity, and ingenuity was the actual working out, the sort of nuts and bolts working out of whatever it is, whatever the proof is, or whatever you're trying to do in math.

Erik Larson:

Then later, he basically seemed to just ignore or ... Like he effectively just jettisoned that idea or that distinction from his discussion. So this is circa 1936, you have Turing talking about these curious aspects of seemingly non-mechanical and mechanical aspects of doing mathematics which of course is just a kind of thinking, right? And then by 1950 when he wrote this seminal paper, which gave rise to the Turing test and is the conversation test that everybody knows about, he just had completely abandoned this idea and he had clearly come to the idea, to the view that we could just program, that intelligence is just problem-solving, so if we solved the problem of human intelligence, if we just write it down in computer programs, then there wouldn't be any difference between the mind and the machine.

Erik Larson:

So he had this change of heart clearly in his work and I definitely think that he's kind of the pioneer and the major figure that all of us in AI look back to. He formulated the actual mathematics of what a computer program is. It's a Turing machine which is a kind of abstraction, right? And all of computer science rests effectively on this formalism called after him, the Turing machine obviously. So he really is, he's the key kind of linchpin for the development. Maybe von Neumann, the other guy ... So we look to him as kind of like the father of AI really and his idea I think was fundamentally flawed. That's not to take away from his brilliance as a mathematician, he was a mathematical logician also which is one of my fields that I studied and he really deserves to be in the annals of history as this brilliant guy. But in terms of his view of human intelligence, I think he had a kind of reductionistic view of human intelligence, and that was inherited by the field through the decades on up to the current view and I think that's unfortunate. That's what I mean by the simplified world, a simplified reductionistic view of the human mind, of human thinking. And I think that's unfortunate actually. That has all kinds of downstream problems, yeah.

Andrew McDiarmid:

Yeah. So the first part of your book, you're discussing that, that simplification of things that we really need to be looking closer at, and then you get into the different types of reasoning and you say that we're not really on the right path right now. Are we still focused on inductive and how do we get to abductive, is it possible to get to that in AI? You obviously don't know the answer but you can surmise.

Erik Larson:

Yeah, I know ... So the problem with induction of course is that you can't deal with exceptions. So the big data AI, this is just ... I don't want to confuse the listener, big data AI just is what we mean by AI today. It's the AI that all the tech that's running on your phone, it's Alexa that interprets the signal from your voice into text sequences and Siri and so on when you have voice-activated stuff on your phone. It's movie recommendations on Netflix, it's automatic recognition of your friends' faces on Facebook. Like all of that is under this umbrella of big data AI, and the key point there is is like I was saying, the inductive inference ...

Erik Larson:

So this also is, this comes from philosophy induction, and the classic example is that all swans are white because in England, if you go and you look at swans, it turns out that every swan you observe is white and so from seeing a large enough sample, this also undergirds statistical theory by the way as well. So if you get a large enough sample size of swans with this property of whiteness, then you can generalize to this ... You can make this inductive inference to this rule that all swans are white.

Erik Larson:

So the rule is never completely certain because you could observe ... It suffices to observe one black swan to invalidate the rule, and it turns out if you go to Australia, there is a black swan and so it turns out that that inductive generalization is wrong. This is the same thing as saying that on Netflix. So I'm building a recommendation system for movies right now at my company, my software company. So this is the same thing as saying that every time that Erik or Andrew watches a movie, it has these features in it and it's over say this number on IMDB and so on and so you just develop the same kind of inductive model as the ... It's the same thinking ultimately as this sort of all swans are white problem, so it says, "Okay, these are the movies that Andrew is going to like in the future." Of course if anything about your past movie wasn't complete, your movie preferences, any of your past movie viewing behavior wasn't complete for purposes of this predictive model, then the inductive generalization is going to be wrong, so this is exactly the problem. As an all-encompassing method to generate intelligence, it's not adequate. Very obviously not adequate. It just doesn't handle exceptions. So anything new comes along is going to be a major problem for an inductive system.

Andrew McDiarmid:

And that would seem to be why different organizations that use this technology are constantly wanting feedback from you. They want that thumbs up or thumbs down. They want you to rate the last thing that you bought on Amazon or the last movie you watched because that will update their ability to account for exceptions or just a larger data field. Yeah, because I have noticed that. Like if I listen to a song and I don't really like it but I'm not afraid to try new things and new artists, but then they'll start recommending that artist and artists like that artist and then we're off on things that don't really relate anymore.

Erik Larson:

Our preferences in the future do likely have some connection to our preferences in the past. But it's very un-useful for things that require any kind of insight, going back to Turing's original idea. So if there's some new movie that comes along that's not in the dataset, then it just is going to ignore it. So induction is a blunt instrument that looks at the overall scope of prior cases to determine future cases, and if you try to use that as a way to generate a true AI or like a general intelligence, it's just absolutely hopeless. Even this conversation we're having now is just full of exceptions, and full of requirements for you to make slight hypotheses and leaps and guesses about meaning and so on. Like in other words, you can't have just ingested a million prior conversations to figure out what to say next. That just won't work. You actually have to understand what's being said now and induction just doesn't give us that kind of capability.

Andrew McDiarmid:

And that's interesting because you see Netflix pouring money, lots of money into new content, and they actually are using big data to figure out what to make next. So again, that can be tricky too. Because if

people just watch it because it's there, like are they watching it because it's there or is it there because they want it? There's that give and take there that can cause a lot of issues, and I guess maybe the futurists will say, "Well, that will even itself out over time," as we get closer to stuff. But it would seem like it would just put a dampener on humanity and our ability to not only think outside the box and to not only new things, but to learn from our mistakes, to adjust. It kind of gets in the way of that, I guess.

Erik Larson:

Yeah, I mean so ... Actually I was, a couple of years ago or I think maybe, a couple years ago now, Hollywood, some of the major studios in California bought these ... I don't remember the companies now, but some major platforms to predict blockbusters. I think this kind of thinking is just ... Like there's this point that I don't make in the book but I feel like it's a good point to make with you because I know I kind of ... From prior discussion with you, I know we share a lot of the same views about technology and there's some curious fact about ... There's this constant attempt in culture to replace actual strong contributions by actual people, like novel movies and new ideas and just all kinds of cultural contributions that happen in the free play of a free society, with statistics from the past, and putting the AI label on just endless statistical analysis of prior successes is kind of like trying to paint the future without anything interesting happening.

Erik Larson:

So it is in this larger sort of [inaudible 00:28:06] sense, it is almost like this machine is sort of gradually taking over our inclinations to innovate and to do and to try new things, but too much of this AI technology just inserted and just deployed everywhere is really, really pulling against that happening. So one just obvious example is that the AI systems that the studios are using are not going to recognize anything that wasn't a formula. So we just know we're getting formulaic movies if they're using these AI systems, right? But we can draw a direct line through this, right? From the fact that the AI used today is inductive in its central commitment, and induction relies on prior examples, from that fact, you draw all of these negative consequences about we are now using formulas. Like what worked in the past is going to be the most frequent blockbusters, the ones with ... Just on a bell curve, the ones that worked in the past. So all of the surprise hits are going to fall out of that, they won't be as likely.

Erik Larson:

So it's effectively repainting Hollywood as being more Hollywood, like more formula and less really interesting stuff and same with publishers in fiction and so on. It's troubling. I actually don't deal with this in this book but I'm glad that you took the ... I'm glad that you asked this question and we went in this direction because I don't deal with that in the book per se. I talk about how using computers to try to do core science is a foolish game and that's a related issue with the question of innovation.

Andrew McDiarmid:

It is a good question to explore. But yeah, just this whole thing of innovation, and before we even started recording, I asked some questions and here's a good question. Is there lessons about the ethics of innovation from the past that would be useful to us today? Can you think of anything they learned about innovation in the past that we could really learn from as we're innovating today?

Erik Larson:

Well yeah. I mean you could go all the way back to some of the core discoveries in science were definitely not big data driven discoveries. So I actually, this is one of my favorite examples is Copernicus,

who of course gave us the heliocentric view of the solar system. At that time they thought it was pretty much the cosmos, but the sun in the center. There was a theory before that that persisted for 800, 900 years called the Ptolemaic model where the earth was at the center of the cosmos and Copernicus obviously flipped that and the original model that he constructed, this sun-centered model, was actually not as predictive. It solved if it could be completed a couple of difficult problems in astronomy, but they had been ... The point is is that they had been accumulating data to patch up this problem for hundreds of years and so they had effectively a pretty accurate view of the movements of celestial objects and so when Copernicus actually proposed his theory, it went against all of this massive accumulation of data.

Erik Larson:

It was actually a human mind insight that really broke from this kind of Bell curve inductive thinking. In the book I think I say ... Like all the data was fit to the wrong curve, how could AI have helped? It would have actually just further entrenched that conclusion and tried to optimize the geocentric model. It took a person actually, an innovator in this case to discover, and it was in spite of, not because of, all the data. So you do see in the history of innovation and the history of scientific discovery, you see these moments where people have a shift in thinking that really can't be accounted for by any kind of mechanistic analysis of what was going on before. One of the lessons is is that if we don't ... It was Norbert Wiener who said, "If we don't invest in human intelligence in our society, we're unlikely to have a lot of it." So one lesson of innovation is to take Wiener's point seriously and to take the lessons of the major scientific discoveries in the Western world seriously, we need to invest in people. So I think in a large part, AI actually ... It certainly does in science and we can talk about it, I can talk about specific cases in science where I think AI is directly threatening scientific innovation to the extent that people are focusing on the systems rather than the ideas behind the systems.

Andrew McDiarmid:

Well the more tech we have around us, the more we start to kind of get the feeling that tech is the answer to everything. We get that reinforced by big tech and their big yearly or even semi-yearly announcements of new products and updates and if we're not careful, we'd just think tech is the answer to everything. But you have said that there are some things in society, some problems that we have, that are just fundamentally non-technical. Can you embellish on that a little bit? Why isn't everything solvable with tech do you think? Is that just because of the fundamental difference between machine and mind?

Erik Larson:

I mean I could give the case of neuroscience research say. So there was something like a billion Euros invested in trying to understand how the brain works and really in the case of the European project, it was called the Human Brain Project I think, the HBP, in that case, it was expressly ... The goal of that project was to reproduce the human brain on a computer, on a supercomputer basically, and the idea there was that if we can just get to a sufficient level of granularity and neurons and systems of neurons and so on, that we can actually code those connections in an artificial neural network, we can just build an actual human brain. This was an actual project, it sounds like science fiction but it was actually the European Union put huge amounts of money in this.

Erik Larson:

Of course it was a total failure and there were these IBM blue jean supercomputers that sat in the middle of the project and it's a different way of answering your question but tech didn't solve that

problem in science, the program was really a fantastic failure. The guy who started it actually ended up getting fired for a variety of reasons but tech didn't solve that problem in science because focusing on technology rather than the actual natural world turns out to have not been a good idea. It's almost like inserting an artificial ... It was like inserting an artificial layer, and so trying to convert basic research in neuroscience into a software development project just means you're going to end up with software ideas and ideas that are programmable on a computer and your scientists are going to be working with existing theories because those are the ones you can actually write and code and they're not going to be looking for gaps in our existing theoretical knowledge in the brain.

Erik Larson:

So I think in that case, and this is a case that I developed in the book, in that case, it was really clear that the introduction of technology as the kind of driving force for success in that project was a really terrible, terrible idea and it was very obvious at least ... I mean certainly obvious in retrospect, but it was obvious to me at the time, I remember thinking that's never going to work. So there's just 10 billion Euros or something wasted. The United States had a similar project under former president Obama which has been a little bit more congenial to actual human research and so it's met with a little bit more success, but the point is is that the idea that you can replace supercomputers with human thinking and human science and human insight and the hard work of scientific investigation and discovery is just a really bad idea. I mean I'm tempted to say it's kind of a stupid idea, frankly. Like why would anybody believe that that's going to work?

Erik Larson:

Computation in general and engineering in general is a downstream kind of idea. It's not a direct connection to nature and to the world around us and so when you're trying to take these downstream ideas and make them central to investigation, you're never going to get to the ground truth. You're never going to grab the root of the problem and that certainly happens, so all these tech connections are not direct connections, including the one we're having right now, right? It's not a direct connection, I can't reach out and shake your hand and say hi and so on.

Erik Larson:

But a certain amount of that just greases the wheels, the modern wheels, right? A certain amount of that is kind of how we keep however many billion people on the planet connected and a lot of that is necessary to drive business and other aspects of the modern world. But there always is going to be this tension. This is getting a little bit away from stuff that I talk about in the book but I'm more than happy to talk about it. I mean it's interesting in its own right, but yeah.

Andrew McDiarmid:

Yeah, yeah, yeah. No that's okay, bring us back to the topic of your book. Well in part three of your book, you finish by kind of warning us of the consequences of carrying the myth of AI into the future. What will happen if we don't get on the right path with AI? And what will happen if we do in your opinion?

Erik Larson:

Yeah, I mean that second one is an interesting question. So yeah, so I ended the book with a question from the investor Peter Thiel, he had asked this question and I just stumbled upon it watching a YouTube interview that he was giving a couple years ago and he was asking, "Is innovation dried up?" In

other words, did we pick all the low hanging fruit and that's why we see a stagnation in innovation today? By the way, I agree with his assessment of the world circa 2021, you don't see a lot of innovation in AI as a field and just in general, right? Like the AI itself has been in the same mode, I would argue for 20 years but deep learning was roughly 2012 so we're just about a decade into this same way of thinking in AI and there's just nothing new coming out of AI science anymore.

Erik Larson:

So like innovation has dried up within AI and it's also, it's sort of ... I feel anyway that it's kind of dried up as Thiel was saying, that it's kind of dried up sort of generally. We don't see a lot of new fantastically interesting things coming out of culture, we just see Twitter fights and we just ... It's the same stuff, there's nothing new happening. So did we pick all the low hanging fruit from the scientific discoveries of the last century on up through the advent of the web, or do we have some kind of perversion of culture itself so that we can't find new ideas. So I pose this question, which one is it? Are we out of ideas or are we not well-positioned to find them? And I hope that we're not out of ideas and I actually still earnestly ... Like I actually am working in the field of AI, I actually work for an AI company, my title is research scientist and I actually want to continue to find innovative ways of doing AI. For instance, wouldn't it be nice. I mean there's a practical aspect here, wouldn't it be nice if our systems that recommended movies actually could capture more of our preferences, right? I mean I don't think that's going to particularly hurt human society if we just had better functioning tech. Since we're not getting rid of it, we might as well try to make it work better.

Erik Larson:

So if we had something more than these recommendation systems that use collaborative filters and tons and tons of data, if we had something more on offer, we might actually have a better tech experience, and we have these independent questions as to whether, how much we want to have that be part of our lives. But I'd certainly like to see the AI improve, and I don't think we're out of ideas. I would like to see though, I would like to believe that we didn't pick all the low hanging fruit and we have a kind of stagnating culture today. We have a kind of stagnating situation and we need new ideas. Yeah, and I think that big data AI has become the enemy effectively. Now everybody in the field of AI, if you have a hammer, everything looks like a nail. Everybody wants to throw data at problems, and we've reached the limit. I mean if you look at self-driving cars, there's no more data that we're going to use for the systems that you have these visual object recognition systems and autonomous navigation systems. Like there's just no more stuff that we're going to include, so yeah.

Andrew McDiarmid:

Yeah. So almost too much data for our own good. Or we need to regroup to figure out what to do with it and what might be important to collect versus what may not be. In Google's view, no data point is useless. That's what I think their thinking is. No matter what we harvest, we can use it somehow. Whether it's in advertising or building new services, and that might be true, but you can't reduce humans to a bunch of data points either. That just undermines creativity and originality, and also spontaneity. So I think we'll be wrestling with this for a while, and back to your point about movies, like sometimes you'll hear folks say, "Oh, I stumbled on this new TV show and I just love it." Like are we getting to the point where we can't ... We're not really stumbling on things anymore because it's being brought to us by algorithms we don't really see and don't really sense? Like I want to still be able to say that I accidentally found you or I just so happened to be in the area. Well, it's just that constant struggle to keep humanity going without technology taking over almost.

Erik Larson:

That's a really good point. I would say especially with the question of big tech, so there's something nefarious going on beyond just the philosophical idea of treating people like data points which certainly is the underlying worldview. Like that's what's happening, that's the view of the person is a bunch of trackable data points in a kind of Cartesian coordinate system. Like that's the idea, but one of the big threats to that kind of life experience is the tech companies are actually not just trying to predict, but it's more effective to predict what you're going to do next if they can control to some degree what you're going to do next. Like I think the big tech companies, there's a fantastic book out by Shoshana Zuboff called Surveillance Capitalism where she explains that it's not just that they're collecting data, they're actively trying to manipulate your choices and manipulating choices involves reducing them. So if you're doing unexpected things, they're making less ad revenue. It sounds like right out of some kind of sci-fi world scenario or that can't be possibly be happening. It's a perverse business model and the connection back to AI incidentally is that if you take away the big data AI, you don't have the number crunching capacity to track and manipulate two billion people on the internet. You need the AI to be able to crunch the numbers.

Erik Larson:

So there is the connection but there but yeah, I'm very concerned about that. I've actually thought about writing specifically on that topic. It's not just that they have this reductive view of the human person. It's also that a consequence of that reductive view is that we ought to be able to be manipulated. So there's an underlying behavioristic assumption as well. There's the idea that people are just material things. I mean if you're a data point, you're not having much of a robust, spiritual life. But on top of just the reduction of the ontological question of what a person is, there's the question of sort of how can you control and predict human behavior to create a good society and the behaviorism is this idea that it's just inputs or outputs, or it's stimulus and response. So that all there is to a person's behavior is a set of stimuli that produces a set of trackable responses. That's effectively ... It's not as if this view has been consciously drawn up in the labs of ... In the halls of Google or something I don't think, although who knows. But the point is is that they're implicitly assuming this kind of reductionistic model and they have a huge financial incentive to manipulate. Not just to predict, but to predict by reducing choices.

Andrew McDiarmid:

It's one reason I don't use Google anymore to search. Number one, I don't want to line their pockets with my data points of what I'm searching for, and number two, I don't want to be influenced in what I find by them.

Erik Larson:

Yeah. Yeah.

Andrew McDiarmid:

So it's like where do I turn? Well golly, I can get onto DuckDuckGo which aids with privacy. The answer I think is quit searching online and go to a library and find books on the subject and talk to people on the subject, look up podcasts, listen to what people are saying about the issue. Don't just quickly search and boom, you found your answer.

Erik Larson:

There's a guy, he's a technologist and he writes on these topics and he's fairly well-known, Jaron Lanier, you've probably heard of him.

Andrew McDiarmid:

Uh-huh (affirmative).

Erik Larson:

He's got a book out that repays reading, it's a quick read, it's called like 10 Reasons to Delete All Your Social Media Accounts.

Andrew McDiarmid:

That's right. Yeah.

Erik Larson:

He's got some other stuff out. The first book I think he wrote is one of his best, it's called You Are Not a Gadget, and it's more of a deeper dive philosophically but in this 10 Reasons to Drop Your Social Media Account, he says like never ... I mean there's some principle in there, some statement that ... Never let passively algorithms determine what you're doing next. In other words, actively search for the movie and do the research like you said, right? So it's just a way of stopping this seemingly inevitable slide into just being these passive sort of data points for profit making companies in California. Yeah.

Andrew McDiarmid:

Yeah, and that's honestly something I want to help people do is just hey, here's what's happening, and here's what we can do about it. Don't be passive with your tech and with big tech and with these algorithms, be active. Don't be afraid to step outside the model. Make mistakes, be accidental, yeah. That whole free thinking and thinking outside the box, we need a lot of that. Yeah, I was chuckling the other day, I saw Facebook experimenting with a select group where they'll facilitate prayer requests. Have you heard of this?

Erik Larson:

No.

Andrew McDiarmid:

Yeah.

Erik Larson:

No I haven't.

Andrew McDiarmid:

Facebook will actually come up and ... Obviously if they've gathered the data that you might be a faith believer or religious in some sort, they'll say, "Hey, do you have prayer requests? We'll help you share them." Which is sort of funny because obviously they only care about us connecting so that they can make money with advertising. But what's ironic is that is how prayer works. Not only are we bringing the prayer to the creator, we're also sharing with the community and they can lift up that in prayer as well and so it's just kind of a funny example of how AI is being used in the religious sense.

Erik Larson:

Yeah, I mean it's ... I mean I would even go so far like artificial intelligence, especially the way it's used today, is really just a convenient marketing term that makes it sound like very progressive and we're creating these intelligent entities and so we're on the forefront of science and so on, and it gives it this really interesting, exciting futuristic sheen, but it's really functioning as just big data crunching computers everywhere in your life. But there is no intelligence in inductive AI. Like that's not what's going on. It's really number crunching, and so we can get into examples of that. But it does get a little creepy, there's a system, GPT3 I think that came out from Open AI which is this company that produces a lot of opensource stuff, AI stuff, but GPT3 is a language model. So it's trained on just millions and millions of example techs and so on. It's just this really complicated language model that uses deep learning [inaudible 00:51:25] pretty much everything ... They're using it to generate text which means to write. They're using it to actually create readable and understandable, comprehensible prose.

Erik Larson:

So there's this big push, like how coherent and interesting and meaningful is the prose that this mindless language model that was just trained on billions of documents, right? With I don't know how many, I can't remember, 75 million parameters or something like ... I don't remember now the specs, I probably got that wrong actually, it doesn't matter. But it's sort of like ... There's something creepy about that where in that discussion, nobody is saying, "So we're just using giant supercomputers and number crunching to mimic writing, the exercise of writing, right?"

Andrew McDiarmid:

Yeah.

Erik Larson:

Yeah, and people are saying that if it gets good enough, we're going to ... There's all these problems that happen if it gets just good enough. So it's never going to write a novel and it doesn't understand what it's doing at all, like that's just not what's going on. But it might be able to create all kinds of divisive tweets, and there's no person behind it. It's just good enough to generate a text that starts a war or destabilizes a market or something. So there are these ... I'm really worried that the tech companies just kind of reaching this just good enough phase where it can just be put to all kinds of nefarious uses, but there's no understanding going on behind it and so yeah, there's kind of two intertwined thoughts there. One is so what if a language model running on a supercomputer can spit out a sequence of tokens of words that every now and then we say it's got a whole paragraph and didn't produce gibberish. It's sort of like so what, it doesn't know ... It's not writing, right? Like there's this weird question, like doesn't it bother anybody that that's what all this research dollars and intellectual effort is going into creating an illusion? Using computation.

Erik Larson:

But there's this other really I think, the other strand of that comment is is that that stuff can be just good enough to really screw up the human world that we're living in. Really destabilize stuff, right? Yeah.

Andrew McDiarmid:

Well, in all of this, the thing I think about is the value of just turning it off. Just pressing that power button and walking away and living your life, whether you're the AI researcher or the [inaudible 00:54:12] who just got a new iPhone or the person Netflixing and chilling. Just turn the bloody thing off

and do your thing as a human. That's one of the things I would say to all this is just ... And that's why my ... It's not a formal organization, but my writing on tech is called Authentic Technology loosely, and my little logo for it, at least on socials and stuff, is the universal symbol for the power button. I'm basically just reminding people of the power of just turning it off and living. Sure we can innovate, sure we can create. I mean I believe God gave us the ability to work with our hands and our minds to do awesome things with technology. That's one of the things I've been communicating lately with the word technology, it basically means the art and skill that we have. But sometimes we just got to turn off the devices and just get back to what's important. Family, connection with others, doing good, making the world a better place before you leave it.

Erik Larson:

Well yeah. Obviously I agree. The AI question is tricky because professionally I want to make it better, but you also have to keep in mind though that if we get away from the data, the really data-driven AI, if we get away from that, we may be able to make movie recommendations without manipulation. Like if I don't have to collect all the data about you to actually provide you some valuable service online, then we have an advance in AI and we also move away from this creepy kind of manipulative data collection model. So the data-driven AI, we have all kinds of reasons as scientists or just lay people, the public to get away from it. We have practical reasons and we're at the limit in terms of just the power of it. Within AI science itself, I think we sort of reached, we saturated, we reached the limits of its ability to perform on all these different tasks we want and so on and yeah. So like we see with self-driving cars, in 2016 we thought they were right around the corner, if you go back and read the stuff written about self-driving cars. I saw that Elon Musk just the other day came out and said ... Basically backpedaled and said, "Well yeah, it turns out that these problems were harder."

Erik Larson:

He's just discovering that there are problems ... In this case in navigation, but they're really cognition problems or thinking problems, right? There are problems that just aren't captured with AI. I'm not a fan of self-driving cars by the way. I follow your point about just turn it off. Like I'd rather just drive the car or somebody else drive the car. I grew up around my cars, my dad used to race cars. It's just like ... I just have a cultural ... No, I'm just not a fan of that. But we're not going to see them anytime soon because the tech just isn't as smart as people think it is and it has limits that I tried to really explain this clearly in the book, it has limits that we don't know how to overcome. But like I said, I do see merit in improving AI, especially away from this data-driven model. I don't see a world where we're not stuck with what we call artificial intelligence. For one, China invests billions and trillions of dollars in it and whether we like it or not, we're in a kind of de facto cold war with China. We don't really want to lose that war, so we kind of have to keep developing advanced computational technologies which we call AI.

Erik Larson:

We're sort of stuck in the modern world developing systems that are capable of doing autonomous things. The drone program and everything else. So there's kind of no way out of that that I can see in so far as we have enemies, but -

Andrew McDiarmid:

That's okay. I mean we don't all have to move to Montana and live in a cabin off-grid in order to appreciate life.

Erik Larson:

Yeah. Yeah.

Andrew McDiarmid:

But we do have to draw the right boundaries and have the right ethics around those things that we're creating.

Erik Larson:

Yeah, I completely agree. Yeah, I'm very sympathetic to your project. I hope that's really successful for you.

Andrew McDiarmid:

Yeah. Well it just ... It seems like a human exceptionalism kind of idea. Let's be respectful of human beings and who they are and how they've been designed and let's by all means create things with our hands and minds but let's be careful about it. So yeah, this has been a great conversation, Erik. I appreciate you taking more time than I thought I'd take from you today about this but this is all good stuff. I started tuning into AI and this futuristic talk when I saw that robot dog and it would get to an obstacle and it jumped over that obstacle or it turned around and went around it, and watching that video, I think it was from MIT or some other lab, it kind of woke me up to, "Huh. Maybe it's becoming smarter." But what's behind that, just out of curiosity, what is behind that dog robot being able to jump over something? Is that just knowing what's ahead of it and knowing what to do about it or is there something else going on?

Erik Larson:

Yeah, I mean I don't work in robotics per se but ... Part of the big challenge with robotics is just the system has to actually know what's in front of it. It's actually a hard problem. It's called visual object recognition for the obvious reasons. Like recognizing objects but part of the problem with robots is just that unless they're in very constrained environments like say a manufacturing arm that just has to see a box and move it to some other position, it's just very difficult for them to actually tell what's in front of them. It seems so trivial to us having actual biological eyesight, but it's very difficult for automated systems to do that.

Erik Larson:

Gary Marcus is an AI researcher and he remarked once that if you want to show how stupid AI is, or if you're worried about killer robots that is, then all you have to do is close and lock your door because they have ... They have a notoriously difficult time actually just locating the doorknob and correctly manipulating it to open the door. They may stand there for hours, just trying to figure out how to open a door.

Andrew McDiarmid:

Yeah, I heard about that.

Erik Larson:

So we might not have to worry about the Terminators yet. But yeah, I mean what's behind that is just advances in visual object recognition and there's no magic to that. They're just training on larger and

larger data sets where they detect edges and surfaces and so on in a plane or a visual space. And then in terms of the mechanics I don't know because I work in language processing. So I don't know how the actual mechanics of the legs and stuff and so on and all that stuff is working, but yeah. The robotic stuff is proceeding at a snail's pace. It's not clear, they're going to hit these limits that I talk about with these inductive systems. At some point, they're going to need something that's more insightful going on, or they're just creating kind of some canned behavior.

Erik Larson:

And by the way, this example you mentioned, they probably did that in the background 50 times before they actually filmed it. Or 150 times, right? So yeah. I mean keep in mind that you're seeing, if you see it at all, it's marketing. Like it's -

Andrew McDiarmid:

Yeah, that's true. That's a good point. All right Erik. Well thanks for your time. Appreciate it. The book is The Myth of Artificial Intelligence and it's out now and I know I'm going to read it over my vacation in full but I do recommend the book to those who are listening. It's a great topic and you have some great insight into it, so thanks again for joining us.

Erik Larson:

Thanks a lot Andrew.

Andrew McDiarmid:

Yeah.

Erik Larson:

I really appreciate it. Thanks for your time.

Andrew McDiarmid:

Yeah, you're welcome.

Announcer:

This has been Mind Matters News. Explore more at mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and edited by Austin Egbert. The opinions expressed on this program are solely those of the speakers. Mind Matters News is produced and copyrighted by the Walter Bradley Center for Natural and Artificial Intelligence at Discovery Institute.