AI Ignition

Ignite your AI curiosity with Cade Metz

To develop AI, believe in AI

From Deloitte's AI Institute, this is AI Ignition, a monthly chat about the human side of artificial intelligence with your host, Beena Ammanath. We'll take a deep dive into the past, present, and future of AI, machine learning, neural networks, and other cutting-edge technologies. Here's your host, Beena.

Beena Ammanath (Beena): Hi, my name is Beena Ammanath, and I'm the executive director for the Deloitte AI Institute, and today on AI Ignition, we have Cade Metz. Cade is a technology correspondent with the New York Times, and his new book, *Genius Makers*, releases on March 16, 2021. Welcome to the show, Cade. We are so excited to have you. How have you been? How has life been during the quarantine? I know you've been busy with the book, but how have things been for the past year?

Cade Metz (Cade): Well, like with a lot of people, it's been up and it's been down. As a reporter, I like to get out into the world. All the best stories come from when you talk to people and you visit places, that's where honestly all the best ideas and all the best quotes frankly come from, and so, during this time, it's sometimes hard to be a reporter and you have to be more judicious about when you travel. I have done a little bit of traveling, but for the most part have been at my home in Berkeley, California, with my wife and two daughters. And we're all hanging in there like everybody else.

Beena: Yeah, I know in normal times you could have met in person. I'm in the East Bay as well, and you could have actually had this over coffee, right?

Cade: Exactly.

Beena: When did you start writing the book?

Cade: It's been years now. I'm now with the New York Times, but I really began when I was with WIRED magazine, my previous job, and actually pitched and sold the book to my publisher when I was with

WIRED and then moved to the Times almost at the same moment. I'm not sure it was a great idea to start a book at the moment I was starting the New York Times, but that's the way it happened, and eventually I got to the end of it all. So, I'm happy to be almost at least to the finish line.

Beena: We have to start from the very beginning. What prompted you? What gave you the idea of writing this book? What was that trigger or was it a series of events? How did you come up with this idea?

Cade: Well, you know what it was, it was a person, and that's what I want to convey to people who don't know about this book and may not even know about this industry, the tech industry, and the rise of so-called AI over the past decade or so. But really the heart of the book is a guy named Geoffrey Hinton, who was born in London just after the war, eventually moved to the US and Canada. He became a professor at the University of Toronto, and he held on to this one idea called a neural network. From his days as a graduate student at the University of Edinburgh in 1971 through to the early 2010s, when this idea finally started to work, it took decades. He is a fascinating person for many reasons, and we can talk about that, but one of the most fascinating things is that he held on to this one idea even in the face of enormous skepticism from his colleagues, from the tech industry and the AI community at large, from society as a whole, this idea would sort of ebb and flow in the estimation of even experts over the decades, but he kept working on it, and by 2010, it had started to work in a really big way, and it's an idea that drives so much of the technologies that we now use on a daily basis and that are quickly coming to the fore in so many new areas. That's always the basis for a good story as far as I'm concerned, someone who has a belief and sticks with it even when the rest of the world does not believe it.

Beena: Yeah. So what was that idea?

Cade: A neural network is essentially a mathematical system, and this idea dates back to the fifties and the idea, loosely speaking, was to build a mathematical system in the image of the human brain. Now that's a really loose analogy, but basically a neural network can analyze enormous amounts of data and learn a very particular skill. So, a really good example is image recognition. If you feed a neural network thousands of images of a cat, it can learn to recognize the patterns that represent a cat. So rather than having to hire hundreds or thousands of engineers to spend years or even decades defining, rule by rule, piece of code by piece of code, what a cat looks like, you can feed the system examples and it learns that skill, and that's a key change in the way technology is built now. This is true not just for image recognition, which of course helps self-driving cars or robotics, helps them recognize pedestrians on the road or street signs, but it's what drives Siri. When you speak commands into your cell phone, it recognizes the words that you say using a neural network. It helps drive the Google search engine. And chatbots are now coming to the fore that use the same basic technique, and they're getting better and better at carrying on a conversation.

Beena: You mentioned in your book that Hinton would say, "Old ideas are new," meaning that scientists should never give up on an idea unless it has been proven that it doesn't work, right? Could you tell us a little bit more about what Hinton meant by this?

Cade: That was really the mantra in his lab at the University of Toronto. He would say, "Old ideas are new," and the neural network is a great example of this. By the late sixties, this idea, the larger AI community felt, had been disproven, that it wouldn't actually work, that there were limitations to it. This was based largely on the work of Marvin Minsky, who was well-known in the AI field, and he wrote a book about essentially neural networks and he talked about this limitation. Even Hinton at one point thought that it had been disproven, this idea. Well, it turns out the proof was wrong, right? His mantra was if you know something is going to work, right? If you have not disproven it, you should keep going, and this certainly bore fruit not only for him, but for the industry at large. By the early 2000s, this idea was pretty much dead, right? There were only a few people who were still working on it and still believed in it. There was a fascinating sort of tension there that I think is indicative of technology in general, and AI in particular. And as it turned out, what we needed was more processing power and more data to make this work, and that's what happened right around 2010, when this really started to work with first speech recognition and then image recognition.

Beena: Right. Cade, yesterday we inaugurated our Deloitte AI Institute in Canada and Geoffrey Hinton was the speaker, and he spoke to some of these aspects, and I love those words of not giving up on an idea as scientists till you prove it cannot be done. But how much of it is dependent on the timing as well because all these other developments needed to happen for the idea to succeed. Just philosophically, at what point do you just pass on that idea through generations of scientists or how do you keep a great idea going?

Cade: You're right. A lot of it's about timing, and even Hinton and the people who did believe in this idea, they didn't realize how much processing power and data we would need. They didn't anticipate the internet. They just really believed in the fundamental way this worked. And it had worked in some ways, like there are great examples of this technology working in very impressive and eye-opening ways in the eighties, but it can never get over the hump when it comes to day-to-day technology, and that's what Hinton really believed in. And even now, though this idea has started to work in so many ways, there are many people who are still skeptical that it will continue to bear fruit. The goals are always further out ahead, and now that we have these areas where this idea really works, there's a lot of speculation over how far we can carry this. If we have more and more processing power and more and more data, if we have enough data to model the whole world, will this same basic idea work? A lot of people are really skeptical about that. Others, like Hinton, think we still have a lot of runway.

Beena: Yeah. Now, in your book you also talk about AGI. Can you explain to our listeners, what is your definition for AGI, and where do you think we are today with AGI, and where do you think we're going to go?

Cade: One of the things I really want to try to do with the book is give people a really good understanding of what that term means. We have two very big labs whose stated mission is to build AGI, artificial general intelligence, they call it, and what that basically means is a machine that can do anything the human brain can do. And when you have people who say that, it's very easy to just take that at face value, and as they say that this technology is on the horizon that we're going to get there. But what people need to realize is that even the people building this technology or saying they're

building it don't quite know how to get there. It's such an enormous task. We're still struggling to get driverless cars on the road. They work well in controlled situations. They can't deal with all the chaos that we can deal with as people on an everyday basis. And a self-driving car is relatively simple compared to a machine that can do anything the human brain can do—learn and reason, and the like. And so that's an ambition more than anything, and my chapter where we go into some of the people building this technology is called Religion because it's almost like a religion, like this idea that you can build AGI, it's inspirational for people. It helps them strive for goals in the near and in the distant future, and the notion that it's possible tends to kind of spread from person to person. It's a fascinating thing because you do need a belief, as we discuss with Hinton and others, to build new technology. You've got to think it's possible, and so that's part of what's going on here. The thing is that that technology is so difficult to build and so complex, it's hard to know when that will arrive, if at all.

Beena: Yeah, but going back to your earlier topic, it's like that idea, and at the right timing, maybe it'll become real. But I hear you because the way I think about it is AGI is great headline topics, right? The clickbait headlines, it makes great headline topic, but in the real world, we are nowhere close to AGI today, and I don't see the technologies or any of the tools that's available for us to even start experimenting with it. So completely agree with you on AGI is more of a belief, a religion. But at the same time, it could be also a timing thing.

Cade: It could. A good example now are these, they call them universal language models. One is called GPT-3. Google has one called BERT. These are systems that essentially learn language skills by analyzing enormous amounts of text, right? Human-generated text, thousands of books, Wikipedia articles, other information for the Internet, and they're remarkable systems, like if you sit down and use them, these systems can now generate tweets, blog posts, articles that look like they were written by human, perfect punctuation, but you've got some people who talk about that as a path to AGI. They even think that this system is sentient because they see certain examples of how well it works, but if you sit down and you use it, you realize that you really have to spin the wheel a few times. Like if you want a blog post about a certain topic, you need to try it about 10 times and maybe five times it'll will work well, the other five it doesn't. It spews something close to nonsense. If you want to take that technology, push it in the real world. it's got to work 10 out of 10 or at least closer to 10 out of 10 than it works now. It's focused on particular tasks. It's not as nimble as our brains. It's not as able as we are to deal with all the chaos in the world that I talked about.

Beena: Right. Now, you also talk about the public's concerns about weaponization of AI and how things could go wrong, and how AI could potentially be used in warfare. What do you think we as the general public should be aware of or what should we be thinking about when we think about AI in the context of warfare?

Cade: Yeah, it's another complicated situation that I hope to lay out in a clear way in the book. There are a lot of opinions on all sides of this. What has happened over the past few years is if you have technology that can recognize images, that can be useful in so many ways, whether it's a drone that does recognizance or a drone that has a weapon on it. The same technology can drive both of those things. What you've seen recently is that technology is inside tech companies. It's not inside the DOD.

The DOD wants it, but the talent and the data and other things are in commercial companies, private and public companies. And a lot of people working for those companies have been concerned when they suddenly wake up and they realize that the company they work for is building that kind of thing. And the line between a weapon and something that is not a weapon is not definite, right? It's the same technology. So again, there's a spectrum of concern there, but I think what has become clear is that these are real issues now that we need to think about. There are companies building this technology for the DOD, and it's worth thinking about how these technologies can be used and how they should be used. It's not isolated to the US. So many other countries are building similar technology. How do we think about balancing our needs with what's happening abroad? These are hard questions.

Beena: Yeah. I love the fact that you're raising some of the top-of-mind questions and topics that we should all be discussing and having a way to figure it out before we get too deep into the development from idea to deployment. How important is it where the technology is being developed? We hear a lot about China replacing United States as a leader in AI. What are your thoughts on, do you think China will replace US in the development of AI?

Cade: Well, one of the interesting things that has happened because so many academics have moved into industry as this technology came to the fore is that there is this belief in industry now that you openly publish your latest research. It's just how it works. Academics have long worked that way and now the companies in this field tend to publish and even share their technology as open source software. What that means is that the latest advances are available across the globe, right? So the latest advances in the US are available to China and vice versa, and what ultimately that means is that what becomes important is the talent. The talent and the processing power, and the data needed to build this and the will to build it. That's the way to think about these things. So, a lot of people have said, well, we should have export controls on this stuff. Don't let it get to China. Well, it's all open anyway, right? It's all in research papers. It's going to get there anyway. And then you start to think about the flow of talent. Foreign town is so important to the US. So many people in the field were born outside our borders. So, you can't just clamp down on immigration and say, that's the thing to do. Again, these are complicated issues, and you need a real balance and realize that this is a global situation now.

Beena: Yeah. So true. And do you think the policies and regulations are going to be able to keep up with all the advances that're happening in AI?

Cade: It's very hard, right? Because you're not just thinking about policy and regulation here in the US, you're thinking about this globally. Again, this is a global situation, and let's say, even if we focus on weapons, you could ban autonomous weapons here in the US, but what about the rest of the world? So you would have to have a global agreement not to do that. Again, we keep talking about so many complex situations, but that's the reality we live in.

Beena: Do you think because of all the complexity around figuring out the technology and the implications of it, it's actually going to slow down? Do you think we'll actually go backwards till we figure out a safe way to deploy this technology at scale or is there no going back?

Cade: Well, again, these are very, very large public companies that are driving this. They have aims and they very much want to get this technology out for reasons of money, as well as other reasons, and now that some of these technologies are in the industry, it's hard to go backwards, but I think what recent history has shown is that we do need to think long and hard about this and a lot of people are. I think you're going to see a lot of push and pull. We continue to see a lot of push and pull across the industry, across the community, inside companies, outside companies. We're going to see a lot of discussion about all these issues, and it's unclear how it's gonna play out.

Beena: Your book is perfect timing because you've taken an objective view of what has happened so far. So it kind of level sets on the technology itself. So, I like the way you've positioned the book more as a very factual, historical perspective of what has happened so far. One of the technologies that we see a lot of use is in natural language processing. And if you had to think of what a chatbot of today looks like 10 years from now, where do you think that technology is headed?

Cade: Well, it's certainly improving at a pretty impressive rate through those universal language models that I talked about. So, what you can do is you can take reams of human conversation, which we have a lot of online now. This is the way we communicate so often is online, turn by turn conversation. These systems can learn to carry on a decent conversation that way, and you're seeing that technology really, really improve. The tricky part is like, as with all these technologies, is that last little bit between pretty good conversation and really good conversation. So if you use these systems now, you could be really impressed for several turns of that conversation, but then once you get deep into it, the systems start to repeat themselves or kind of that real call and response kind of breaks down. The trick is to get to that point where it really, really works well and where it can really pull in all sorts of information that we have in our brains. It's not just about you and I putting words together, it's drawing on memories and drawing on knowledge and being able to inject that in the conversation and remembering what you and I just discussed a few minutes ago and not repeat ourselves. That's a very, very hard thing for a machine to do. And so we have sort of this surface conversation now, but getting that deep conversation, that's going to be harder, and that's a real open question about whether or not, or how soon we can get there.

Beena: Should we even go there? One question that I always try to figure out is, should we even be going there? What's the purpose of it? And there's this notion of digital humans, which is recreating celebrities or people who have passed away, loved ones who have passed away, and trying to keep that legacy and have a conversation with them almost live and it's being enabled by NLP systems, right? How important is that question of, should we even go and develop this technology? What have you learned as you built out this book? Is that something that comes in scientists' minds of, should we even research for it? The classic sentence from *Jurassic Park*, just because your scientists could, they did without thinking whether it should be done. What are your thoughts now having looked at this for a few years?

Cade: It's another key tension in the book, and some scientists certainly think about that. Whereas some people are intent on building AGI, if we go back to Geoff Hinton, Geoff Hinton really questions whether or not that idea makes sense. He says, "I'd love a robot to be able to perform surgery on me, perfectly. Does that robot need AGI? Does it need to know baseball scores? Does it need to know all these things that are unrelated to performing this particular surgery on me?" He says no. He sees the technology

being deployed in specific ways and he asks, “Why do you need something that can do anything the human brain can do?” It’s a scientific goal that is fascinating, and it’s an interesting way to spend your life, but ultimately does society need that? That’s a question that comes up all the time, and it has come up since the fifties in science fiction novels, certainly in movies, that’s always the big question. Do we really need this and where might it go wrong? Do we end up hurting ourselves if we go too far?

Beena: Yeah. Cade, what is the best use of AI that you’ve seen so far? Which is the best use of AI for humanity is good.

Cade: One of the key results over the past 10 years actually just happened this past fall. There’s a lab in London called DeepMind, which is also a big part of the book. They’re now owned by Google. And this past fall, they essentially solved a problem called the protein folding problem. It’s a biological sciences problem that dates back decades. Scientists have been trying to solve this for decades. And it’s an important way of developing drugs, medicines, and repurposing medicines. So, take the situation we’re in now, this pandemic, when that situation comes up, this protein folding problem can help us determine what drugs we already have that could be used to treat a condition, but also how we build new medicines that can treat it. That’s the very situation we’ve been in over the past year. Now, we were able to deal with it relatively quickly, but what if we could deal with it far quicker? And that result is an indication that we can do it far quicker than we did in the past. And so, it provides real hope for future pandemics, that it won’t take so long to produce a vaccine, that it won’t take as long to kind of repurpose the medicines we have and apply them. Again, it’s a complicated situation. You have to go through drug trials and you have to make sure you’re on firm footing before you get a vaccine or repurpose a medicine, but it’s an incredible result because like a lot of the stuff that we’ve seen over the past 10 years, few people thought it would happen. I mean, biological scientists have struggled with this for decades and many people thought it wouldn’t be solved for decades to come, and it has essentially been cracked. Now, it’s not perfect, like I said, and there’s still a long way to go, but it shows the kind of progress that can be made in important areas. That’s an important area.

Beena: What is another area that you are most excited about that AI could potentially solve in the next 5 to 10 years?

Cade: Another place that we haven’t discussed where this is really showing progress is in the robotics field. We talked about self-driving cars. I think that’s a key area. It’s a hard problem, but we’ve got progress there. The other places though where you’re seeing this really start to work is in warehouses. Another problem we’re dealing with at the moment is that Amazon, for instance, helps supply us with so many things and they operate through these giant distribution centers, warehouses, they have to sort through all this stuff that we order on a daily basis and ship it out. It comes into the warehouse, it gets sorted, and it gets shipped out. Companies like Amazon don’t have the people power to do that. They’re having trouble hiring enough people to do it. Ideally, they want machines to do it, to sort through that bin of random stuff and get it where it needs to go. That’s a very hard problem for a machine, a robot to handle, but these same ideas we’ve been talking about have shown huge progress in this area, and you now have machines that can take a bin of random stuff and pick it up and place these random objects where they need to go. There are several companies that are working on this, it’s been deployed in the

real world, and that's just one example. Autonomous drones is another example where there's been huge progress, and those can be used in ways that scare a lot of people, but also in positive ways. I did a piece recently on a police department in Southern California, which uses these essentially self-piloting drones to respond to 911 calls. There are a lot of civil liberties questions there, but there are also many reasons to make that happen, and you see this time and again with robotics, and I think that's a field where, again, progress is accelerating.

Beena: Yeah. So true. I remember this was at an aviation company, we were chartered with figuring out what should we invest in for the future of aviation. Till now we've been building jet engines and massive aircrafts. What does this look like in the future? And it started with, the question was more from the lens of AI to say, okay, today we can predict when an engine might fail, so you proactively send an engineer to go fix it, but what is it going to be 25 years from now? We can automate the engine to self-heal, put in some kind of 3D printing to, if there is a scratch on a blade, it can fix itself by putting the software and the hardware capabilities within the engine. What we did when we went through that path of identifying what it should be, we actually looked at advances that were happening in drone technology and we said in the future, you're going to have these self-contained drones that take you from point A to point B. So you can avoid mass air travel. Now, this was about eight years ago, but the pandemic has actually, I'm like, I've been thinking more and more about it because if you look at all the technologies that're coming together, it may not be far when you move even beyond self-driving cars to have these self-contained drones, which take you from point A to point B to just avoid that whole mass air travel problem.

Cade: It's a great point. I think that gets under-covered in the mainstream press and even the tech press because a self-flying drone in many ways is an easier problem to tackle than a self-driving car. It's the same basic technology. It's the same type of sensors, and it's the same mathematical technology in a lot of ways, the neural network idea we've been talking about. But you don't have to deal with the uncertainty, a lot of the uncertainty that you deal with on the road. What you do have to deal with is regulation. So how do you deal with all the planes and drones flying around and that sort of thing, and you have to deal with the privacy issue and the like. This is by no means easy, none of this stuff is easy. But in a lot of ways, it is easier than self-driving cars and there are a lot of possibilities there that I think get under-covered.

Beena: Yeah. So true. I mean, when you start thinking of solving for the larger, more complex problems, there are creative ways at trying all these technologies together can help solve those problems that are not as obvious. Cade, this has been great, and I know your book is releasing on March 16, 2021.

Cade: You're right. I wish it would go ahead and happen.

Beena: I'm sure. It's exciting. I am looking forward to it. And where can our readers buy your book, get the book?

Cade: Any place that sells books. Amazon, Barnes & Noble, Target, Wal-Mart, independent bookstores, which I'm partial to. It's in hardcover on the 16th, Kindle versions or digital, and an audio book, which I

haven't heard yet, but I'm excited to hear. So, however you read or listen to books, you'll be able to find one.

Beena: Did you read your book for the audible version?

Cade: I was going to, but then, like one of my dreams for the book was that it would be read by this British actor named John Lee. He's incredible with audiobooks and, in particular, he does the books of one of my favorite writers, a guy named Ben Macintyre. And two of my main characters are British, and part of what we discussed with the book was giving that sort of British or European flavor to the book and John Lee has read the book. I'm really excited just myself to listen to it.

Beena: That's awesome. So, Cade, how can our viewers stay in touch with you? What are your social media handles, email? What do you want to share as your contact information?

Cade: Well, my stream of stories at the New York Times is one way, but I'm @CadeMetz on Twitter and usually post most of my stories and some other stuff there as well.

Beena: Cade, thanks again for being with us on the show today, and I want to thank our audience for tuning into AI ignition. Be sure to stay connected with the Deloitte AI Institute for more AI research and insights. Thank you.

Thanks for tuning into another episode. Check out our AI Ignition page on the Deloitte AI Institute Website for full video and podcast episodes, and tune in next time for more thought-provoking conversations with AI leaders around the world.