



AI Ignition

Ignite your AI curiosity with Dr. Anima Anandkumar

How will the job market change with the evolution of AI?

Introduction: From Deloitte's AI Institute, this is AI ignition, a monthly chat about the human side of artificial intelligence with your host, Beena Ammanat. 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): Hello, my name is Beena Ammanath. I am the executive director of the Deloitte AI Institute, and today on AI Ignition, I have the honor of speaking with Anima Anandkumar. She is the director of Machine Learning, leading research at NVIDIA, and brand professor at Caltech. Her research spans both theoretical and practical aspects of large-scale machine learning. Anima recently won the Women in AI Award, and I am so happy to be speaking with Anima today. Welcome, Anima, it's great to have you on today's show. How are you doing?

Anima Anandkumar (Anima): Thanks a lot, Beena. 2020 has been a tough year for so many people, and I'm so grateful that I get to be here today, and I get to continue my research, even if it's virtually both at Caltech and NVIDIA. I'm hopeful for all the developments that are happening and that we can get a hold on the pandemic and further accelerate our research when it comes to drug discovery vaccine development. There's so much there that AI can help, and I'm excited at those possibilities.

Beena: I know you are a well-known researcher in AI and machine learning. Can you tell us some of the interesting areas that you are currently focusing on?

Anima: Absolutely. For me this is an exciting time to be an AI researcher. Deep learning broke through and we were able to see gains in so many domains simultaneously. And the question that I'm very excited about is, where do we go from here? What are the next generation AI algorithms that can now very much solidify the foundation that deep learning has given us and ask in terms of how we can get generalizable AI. So, not just do it on a narrow task, not just fit to the training distribution that we are given, but be able to extrapolate to new scenarios, be able to handle constraints like safety and stability in autonomous systems, such as self-driving, be able to handle distributional shifts, be able to train on simulations and go to the real world in robotics, be able to use AI to make scientific discoveries and make a big impact on different scientific domains, whether it's solving partial differential equations in

such a broad set of applications, whether it's understanding our own brain and getting inspiration in terms of how our brain can do such amazing tasks. Those are all problems that I'm very excited about.

Beena: Yeah, there's so much there to unpack, I don't even know where to start. I know you also lead the AI for Science initiative at Caltech. Can you speak a little bit about that? I know it's focused on bringing together AI researchers with experts from other disciplines to push AI tools into every area of science and engineering. What prompted you to join the initiative and to lead it? What are some of the things that you're doing under the AI for Science initiative?

Anima: Absolutely, Beena. AI for Science is an initiative that I cofounded with Yisong Yue, who's also a professor at Caltech, and what we were seeing was, across campus, so much interest in using AI in different areas. For instance, I'm working with Zach Ross, who works in seismology and wants to understand how we can do early detection of earthquakes. I'm working with Tom Miller, who wants to understand how the energy in the molecules can be calculated so quickly to enable applications like drug discovery at scale. I'm working with Doris Tsao, who works in neuroscience. She has monkeys in the lab to have neural recordings to understand how visual perception works and what is the role that feedback has in our brains. So, you can see such a broad range of applications. And so as an AI researcher, it's very hard to give our expertise and collaborate in so many areas unless there is a nice structure for us to easily integrate AI into all these such different domains. It's also important to note that it's not a straightforward application of an AI algorithm. It's not like I hand over a source code and there, we're done. There is new AI to be invented when we understand the deep challenges in all these domains. We need to do unsupervised learning, we need to be able to robustly learn and handle perturbations, we need to be able to work with limited data. So these are aspects that, to me as an AI researcher, are also great testbeds for us to develop new next-generation algorithms. And what the AI for Science does is to provide that umbrella to enable such activities. We have periodic workshops, we have machine learning researchers who volunteer their time and have people come from across campus to talk to us in terms of the problems where AI can be useful. We are providing them with tools and starting source code for different deep learning methods. So, it's like kind of asking how we can scale up this collaboration in a systematic way to enable AI into all these different applications.

Beena: That's fascinating. You are also one of the most well-known prominent people who straddle between research and industry. So, my question to you is, what does a typical day look like for you? How do you switch between research and industry? Can you walk us through a day in the life of Anima?

Anima: Well, that's a great question, Beena. I'm not sure if I can answer that well, because there isn't something like a typical day, so every day is different and interesting, and now with the pandemic, things going virtual, it's also very different from where it was before because before there was also quite a bit of events and conferences to go to, and there is also so much interest by the public to know more about AI. But now I'm spending more energies in recording podcasts like this and hoping that a broad set of people can get access to this, so I can spend more time on research, so that aspect I'm thankful that there is less travel and there is more time to focus on research projects. And for me, like straddling between industry and academia is somewhat more organic and natural because at NVIDIA, I have the freedom to pursue my research goals, we are open sourcing a lot, we are publishing, we are out there in the community, and that gives me the freedom to really pursue what are the exciting questions which NVIDIA can be part of, what is cutting-edge research we can do. And so that makes it more seamless, and I'm very thankful for that. So, a typical day looks right now as many virtual meetings where we are discussing project ideas, we try to meet as a group, even though this is now virtual, try to make sure everybody is feeling good, they are able to cope with these changes, and debate and brainstorm where are the gaps. Why is AI not able to solve, for instance, an Olympia challenge, which is a really difficult

math challenge at an international level. Or why is AI not able to do language understanding, why it cannot replace human content moderators or tackle misinformation the way it is today. So, I also engage in a lot more brainstorming now at the group level and trying to make sure even in this virtual realm, we can all be connected and we can pursue our goals together.

Beena: Fair enough. I think a general day, even compared to last year and this year, is so drastically different in which we could never have anticipated. You touched a little bit on generalizable AI. What do you think, which industry or which sector will see generalizable AI initially? How do we get to generalizable AI, and which industry might see the first use of general AI?

Anima: When we say generalizable AI, I think it will also be domain specific, because the question is, what is the level of generalization you need? If you're doing face recognition in a very controlled environment, say, the lighting, everything you have under control, it is to let somebody inside a building, and most of them are employees, so you are only letting employees in, so you have a database of the faces that you want to recognize. That is a much more controlled, narrow domain compared to on the other end of the spectrum, we have self-driving where there could be so many scenarios that are never present in training data, you can't possibly cover all of them. And it's also safety critical. There could be many of them where it's a life and death situation, and the behavior is very different from how humans would react in those situations. So, it's such a broad spectrum to say what does generalizability mean, and that's what I think is a big challenge when we are now trying to get AI into the real world. We have to now understand in each domain what is the level of generalization we require. For instance, in the project that I'm doing with Professor Tom Miller at Caltech, we are asking the question for computing quantum energy in molecules. How can we train on reasonable-sized molecules so that it's possible, it's feasible, if you try to train on large molecules, such as proteins with thousands of atoms, it's way too expensive. So, can we limit our training to just, say, molecules of size up to 40 or 50 and then can we generalize it to these large protein molecules with thousands of atoms? So here you can see that this is no longer the classical machine learning paradigm where the training data and the testing data looks similar. So here the level of generalization is to be able to transfer to large molecules. And the way we went about doing this was to use domain-specific knowledge. For instance, there's something known as molecular orbitals that we discovered can transfer very well to these large molecules without needing to train on them. So, I think it depends on the domain, on what level of generalization we require and how do we bring in additional knowledge without just the data and a supervised learning technique. For instance, in autonomous driving, it can be the use of simulations to augment the real data, and we are now extensively studying these techniques of how to combine simulations and real data together. So, in that scenario, the generalization is not to overfit on simulations but go to the real world.

Beena: I know you work across several industries. I've read some of your work, and you look broadly across multiple use cases and also using the same technology in different scenarios. Are there any industries or is there any work that you think AI will not transform, say, in the next 10 years?

Anima: Beena, that's I guess the trillion-dollar question or even bigger. I do think in every industry AI will start the transformation. The question is can we get until the very end point? Probably not. So, some are low-hanging fruits and the others are much harder problems to solve. I've been talking about autonomous driving. Even there, getting to Level 3 is something that we are now seeing a lot more abilities, a lot more car companies doing that and L4 and, ultimately, L5. So, there are stages to what we can do, and I think in every industry it's going to be that. I think we will get to a point at least with a narrowly defined task or with having a human in the loop, we can get some benefits, but the question is can we take a human completely out of the loop and can we have a general AI system that can just adapt, that can be able to handle all scenarios. I think that's a much harder question. So, I would say the

same also for healthcare. I mean, now that we are seeing with the pandemic, it's so important to be able to discover new drugs very quickly, have vaccine development be done in record time. All this, I think, points to the promise of AI to speed up these processes. But there are also other aspects. We do need human trials for vaccines, so probably we cannot shorten beyond a point. And the same with drug discovery. There are so many complexities there, so how do we evaluate whether a drug is a promising candidate even in simulation. There's so much of additional domain knowledge that is to be built in and we're making progress there, but those are some of the harder problems to go all the way to get the benefits.

Beena: What does the future of retail or shopping look like? We're doing things very differently in the last several months. What would be your take on the long-term impact of AI on retail?

Anima: Absolutely, Beena. I think when it comes to retail, so much of AI was already getting used, but I think people hadn't completely embraced it like the way we did now. So that will provide even more efficiencies in the system because if it can predict the shopping behavior of people and how often they're buying and how to optimize the supply chain, how to deliver it to the customers in an efficient way. So, I think there are aspects of it that I don't think require fundamentally new AI research or innovations. There is already with the existing algorithms and tools, we can make a lot more progress because of the current situation where people are embracing online shopping much more, and I think a lot of us will not want to go back. Maybe I may want to go to a mall one day just to have like a fun day, but most of the time I might just do online shopping, because I have a lot more choices and it's much more convenient. So I do think that the scale of it will really explode and the options will also be there. The question is how AI can then benefit from the scale, and we'll see a lot of that.

Beena: Yeah, so true. I tend to think about AI as such a broad umbrella. I tend to think about three parallel streams that are speeding along at different paces. One is the core research, that is developing the core technology needed to enable AI, and then the second one is really the applications of AI, whether it's in different industries or different functions, there's the application, and the third one is really the risk and consequences. Because all of this is so new, we don't know all the risks associated with it, whether it's ethical considerations, the consequences around, not necessarily all the laws and rules have been figured out, the use of this technology because each one of these parallel streams is growing at different paces. Changing gears a little bit, Anima, I'm curious about this as you have this unique lens from a research and industry perspective. We hear a lot about AI taking away jobs, but also AI creating new jobs, new roles. My question is, have any of these new jobs or new roles been a surprise to you? Many of them you probably saw them coming. What role that AI created in the past 10 years was actually a bit of a surprise to you.

Anima: I'm happy to see there is such a broad set of roles being created, from data scientist, to data engineer, to somebody who understands the domain and knows what kind of data should be used, to a data consultant, and, of course, the infrastructure side of it, how do we build good infrastructure for AI. And also, I think maybe the most probably surprising one is an AI artist. So, somebody who may not know very well how to paint, I mean, their own hands may not have the skills to do that, but then their creativity is there and they want to maybe use the NVIDIA GauGAN and create amazing landscapes. Same with producing music. So, we'll see people being creative in these new ways, even if they're not having the base skills to do it in a much more direct way, like they cannot directly paint, maybe if they're somebody like me, I can't sing very well, but maybe I can have AI to realize something that I cannot directly do.

Beena: I'm a big believer that AI is going to make human lives better and some parts of it is helping us focus and develop skills which we may not have had time to think about. Staying on the same track of jobs, being a pioneer in this field, what do you think are some of the new jobs that we'll see evolving five years from now, and what's your prediction of what will be the most in-demand job in five years?

Anima: Beena, I think we're seeing things all much more quickly because the pandemic has created now different conditions. So first of all, now remote work is so widely accepted. Thankfully, at NVIDIA, this was accepted even before the pandemic. I know so many amazing colleagues who are in Europe, Asia, even New Zealand. We've had the pleasure to work on projects that have worked so well, and we can adapt to these time differences. But I think much more broadly, people have realized that even amid such a challenging pandemic, when kids are at home, they are having to be helping with all the schooling and everything that is happening, still the productivity has many times been better than what was in office. So I think leaders are realizing that. Forcing people to commute long distances and spend all their energy of the day in that versus staying home and still being able to connect remotely and be productive and maybe once in a while there is an in-person meeting to keep the connection. I think we'll see a lot more of that even after the pandemic is over, so that will be a big change. And with that means even more globalization and democratization, so there will be much more openness to considering candidates across the world and hopefully with much more diverse experiences, so not just looking at candidates from a top school, but now if you're opening up the opportunities and also asking, "Okay, now how do we judge people if they haven't just gone through the traditional academic route, how do we find talent across the world?" I think we'll see also new ways to handle this workforce that is now available for this knowledge economy, and with that, we'll see a new ecosystem of jobs getting created.

Beena: So, Anima, can you share with us what does it mean to you when you say to democratize AI? What does AI democratization mean for you?

Anima: Thanks, Beena. I think that's something I'm very passionate about, and I love the opportunity for more people to think about it and contribute towards it. To me, democratization comes with two facets. One of them is representation. They should be represented in these AI models. So, when we're talking about face recognition that is being served in this country, all demographics, age groups, genders should be represented well. So they have a right to representation. The other is accessibility, whether it's in terms of AI education, they want to learn about these AI models, computer resources, they want to be able to train their own AI models and their own data, so the data privacy and the right to decide how their data will be used. All this to me comes under the realm of democratization, because ultimately people should have the power to decide how they want to participate in this AI revolution, but also how they can benefit back from it and how they are represented well.

Beena: Yeah, that's fascinating. I do hope we get to that part of the AI journey soon where we can truly democratize AI. We work across several large enterprises, large companies, and we've seen there are companies who are very early in their journey with AI having a few POCs, have a few AI solutions in production. There are some that are much more advanced and use AI across their entire organization with their products and with all their functions. And then there is a big group that's in the middle of this who have AI sprinkled throughout the organization. I have a two-part question. What suggestions do you have for CEOs and board members who are fairly early in their AI journey? How should they be thinking about AI for their business? What would your advice be to the early-stage companies who are just exploring AI?

Anima: I think that's a great question, Beena, because there is indeed a wide spectrum of AI adoption, and to me I think the industries that are struggling to adopt, the primary barrier is data. So, they don't

have existing data infrastructure, they don't have pre-collected datasets that are available, they don't have benchmarks and tasks well specified. I think focusing on that part of it is so critical. And that has to come from their domain expertise. It's not like somebody from outside, if you hire an AI firm, may do a good job, especially if it's a very specialized case, it's not something that a standard AI expert would be able to understand. I think understanding what kind of data is available, what is the quality of it, what can be the ambiguities in this data, that is so critical because that's the starting point. And so that's why we need this model similar to what we've done with AI for Science at Caltech. You need AI experts who'll have the skill set to closely engage with the domain experts and same on the other side, domain experts who can talk across their expertise and across their domains and be able to understand what are the constraints here in this domain, what can I utilize and get that going, so it has to be a very integrated close connection and a collaboration to make this happen.

Beena: And what about the CEOs who are in that last bucket, who are more advanced in their AI journey and have AI being used pretty much across their entire organization, right from their core product to their functions across marketing, finance, HR? What would your advice be? What should they be thinking about?

Anima: I think they are in such a great place, because they're leaders in this AI journey, so they should be looking to invest in research to further keep that leadership role and essentially be the ones to get us all the way, because there are so many hard problems, like I mentioned, that research is still getting done, especially when it comes to generalizable AI. They can look at what are the failure modes in the current AI system and whether new research can fix that. They can also look towards more integration, maybe now they're doing AI in silos, maybe the marketing team is doing AI model on their own, the HR team is doing AI on their own, and the question is, can you build those synergies? Is there now better knowledge and better insights to be had when you bring these different departments together and different models together? I think that's the other exciting piece when you can have finally a very integrated AI strategy for the whole company.

Beena: That's great input. I love it. So one is, really invest in research and they could be the pioneers in their industry on what is possible with AI. Anima, thank you so much. How can people stay connected with you? I know you're very active on social media. I follow you. What's your recommendation for our listeners to stay connected with you?

Anima: Thanks, Beena. I'm on Twitter. AnimaAnandkumar is my handle. There, I talk not only about my research, but also the important societal issues, whether it's diversity and inclusion in tech or now the danger of misinformation. I think these are all aspects that all of us should be concerned about, no matter what our background or no matter what our skill sets are. So I hope you can join me there.

Beena: Anima, thanks again for being with us on the show, 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, Anima.

Recorded Voice: 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.

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to deloitte.com/about.