

Rashmi Mohan: This is ACM ByteCast, a podcast series from the Association for Computing Machinery, the world's largest education and scientific computing society. We talk to researchers, practitioners, and innovators who are at the intersection of computing research and practice. They share their experiences, the lessons they've learned, and their visions for the future of computing. I'm your host, Rashmi Mohan.

If you had a crystal ball to predict into the future, what would you ask of it? Whatever it might be, our next guest might have the answer. An AI expert, a mathematician, a change maker, a computer scientist, an advocate, Anima Anandkumar wears all these hats with a plum. Anima is the Bren Professor of Computing, while also being the youngest named chair professor at the California Institute of Technology or Caltech. She's also the director of machine learning at Nvidia, where she leads a group developing the next generation of AI algorithms. With a string of awards, including a Guggenheim Fellowship and NSF Career Award, Young Investigator Award from the US Air Force and the Army, and most recently named an ACM Fellow, she is no stranger to the spotlight. Her work has been extensively covered in PBS, Wired magazine, MIT Tech Review, YourStory and Forbes, with the laser focus on using AI for good tech, she's a force to be reckoned with. Anima, welcome to ACM ByteCast.

Anima Anandkuma...: Thank you, Rashmi. That's very kind.

Rashmi Mohan: Wonderful. We are so excited to have you here. And I'd love to lead with a simple question that I ask all my guests, Anima. If you could please introduce yourself and talk about what you currently do and also give us some insight into what drew you into the field of computing.

Anima Anandkuma...: It's a pleasure to be on this podcast and at the moment I'm very excited about how AI intersects all the other domains, especially in scientific domains, how we can integrate AI into existing workflows as well as create nuance? Can we reimagine how we do science, and along with it, bring better inventions and discoveries at a much faster pace? And that's a great place to be, because I'm always learning about new areas and seeing how AI intersects and enables all of these different domains.

Rashmi Mohan: That's amazing. I'd love to get into more of that and understand, it just could go in so many different directions that I'm really looking forward to digging deeper into that. But what was that introduction into computing or computer science, Anima? Was there early exposure that drove that interest?

Anima Anandkuma...: Yeah, I was just really... It was a privilege to be in a household where math, science and computer science was very much part of daily life. Both my parents are engineers. My mother is an engineer, which was not so common in the community. I grew up... In fact, she tells me that she went on a hunger strike for three days to get into engineering because my grandparents were worried that she may not get a suitable husband after studying engineering, but she was a

trailblazer. And that really helped me a lot to think of this as not something an anomaly, it's something that women could be very good at. And I had role model too at home itself. That was great to have.

And it was also great to see that intersection between science, and math, and engineering. My parents ran a small scale industry manufacturing, automotive components, different parts of the manufacturing processes. So I would go see those machines and at some point they decided to make them computerized, what we call CNC, or Computerized Numerically Controlled machines. And so that helped me think about programming as something very tangible and physical because the program code would go and move these huge components, remove material from that. So seeing that, and engage, and seeing computing as not something in its own bubble, but that's some interacts with the physical world in such a real way was really great.

Rashmi Mohan: That sounds amazing. For one, I think just the idea of inspiration starts at home, and I think you had enough inspiration from both your parents, which is incredible. And I think the other thing that you bring up is, I think for many of us computing may have started only with introduction to maybe software programming, but you saw the applications of it and how it could actually solve engineering problems or real life problems. And I can imagine how that could've basically not made this a difficult or a challenging field to get into. It felt very real, it felt like something that you could definitely achieve.

Anima Anandkuma...: Yeah, absolutely. And I would just do these exercises at home for fun. My parents would have these manuals and we would just be looking at how to move the target to a different location. So it was very geometric and real and the program to enable that. All of these different areas came together that way. I also had my grandparents be very passionate about math. My grandfather was a math teacher, so he would always give me problem sets and concepts even before they were taught in school. And he would initially say, "No, no, that's too advanced for you." And that was almost like a primer to me rebelling and saying, "No way. I can do this. Bring it on." And so we had this back and forth. And my grandmother as well would give me all these puzzles. She was very much hands-on showing us different puzzles and embroidery and all these aspects that are very geometric and physical. So seeing both, the abstract and the physical, come together was just a great way to grow up.

Rashmi Mohan: That sounds like such an amazing environment to be exposed to all of these topics. We often talk about how the pipeline to get women interested in computing is so weak, and maybe we should start earlier to be able to expose folks to the areas of both technology and computing. Sounds like you had that perfectly set up for you, so we'll definitely get into that in a little bit. But I know that for your undergraduate education, Anima, you went to IIT Madras. Was there anything there that changed the way you thought about the field, or did it inspire you, or did all the training, whether conscious or unconscious, that you

received through your childhood help you be successful in your undergraduate education?

Anima Anandkuma...: Yeah, first of all, getting into IIT, as you would be aware, is not easy. It's an entrance exam that is highly competitive and you have to be in the top few hundred to get into IIT with your choice of major. To me, I think tackling that exam and doing it in Mysore, where I grew up, where there wasn't a lot of formal preparation classes. I did consult with university professors, but I had to take the initiative. It was me being very proactive and mindful about how do I tackle such a competitive exam and finding the right mentors and supporters to help with that. It gave me a sense of purpose and a sense of almost entrepreneurship because if I needed to do it, this was well beyond what the standard syllabus would cover, but that also really forced me to understand the concepts deeply.

And I think this will connect to how we think about education in the future, especially with ChatGPT and so on. I think the IIT entrance exams were those that a ChatGPT wouldn't easily solve. I'm pretty sure, because they weren't standardized problems, they required some out-of-the-box thinking, they required you to really understand the concepts deeply. But that challenged me, so having that kind of exam gave me that reason to go deep. And then I realized this would be great to do it all my life, to really not just touch the surface but go deeply into concepts.

And once I landed in IIT, being in a cohort of students who are just amazingly smart and also multidimensional, I think that surprised me because I assumed everybody would be just books down nerds, which they're great at, but also, we organized student run conferences where we brought in researchers from across the world. We built bridges made entirely of paper on which we could work on. We did a lot of cool nerdy things. So having that cohort of students who are still my friends was great, and of course, all the professors and the environment that it brought about. I have great memories from that.

Rashmi Mohan: That's amazing. It sounds like you've been a hustler throughout your life, and I love the fact that you were talking about multidimensional exposure. Because if I look at your career and the work that you've done and the work that you're doing, it is how you branched out. It's not the straight and narrow, you've used the computing principles and the skills that you've learned to apply it to solve real world problems, and I'd love to talk about that as well.

Anima Anandkuma...: Absolutely. I think it's important to think about education not as a fixed curriculum or a fixed path. And even the job we land up as one you would do for life, that unfortunately is mostly over for most of us. And the way we keep up with this ever-changing technology and the landscape of jobs is to keep re-skilling, to keep learning. And I think inculcating that early on is so important. And to me, I think having my parents be entrepreneurs who are always bringing on new technology, they're not just satisfied with something that's already

working, but taking big risks in terms of capital investment and bringing new technologies to ground was something that I grew up as an everyday thing.

I also take inspiration from my great-great-grandfather, Dr. Shamasastri, who, from what I hear, the story goes that he left his village as a teenager, came to Mysore, ultimately did his PhD in English and Sanskrit, rediscovered the Arthashastra from 300 BC. So having this really needle in a haystack, it's this ability to understand that there is something deep in these palm leaves that's a language that is unknown. And being able to crack the code and connecting it way back to history and to such a landmark text gave me the inspiration that if you follow the path of curiosity, great things will come.

Rashmi Mohan: It's incredible, what you said about your great-great-grandfather in itself is mind-blowing. But also I wonder if that inspired your journey, because you seek out challenges. From what I've read, from what I've heard of you, that's something that is core to who you are and the way you approach your work. I was wondering if you could maybe take us through some of the early work that you did. I know during your PhD you were working with low power, lightweight communication methods and IoT devices, and that was many years before maybe it became as popular as it is now. What were some of those breakthrough moments? How did you pick those areas of interest? And where did you feel like you really developed your skill as a researcher?

Anima Anandkuma...: That credit there goes absolutely to my advisor Lang Tong, because he was working on sensor networks and connecting them wirelessly way before it became this edge, AI and IoT or whatever you may call it now. And the core mathematical challenges are still the same. If you have these noisy sensors with limited battery power, limited abilities to communicate, how much do you process on the sensors and how much do you send it out to the cloud? Because ultimately the goal is not to record all the raw data from the sensors, but to really make inferences about, say, if it's a chemical plant, is this working normally or anomalous? If it's a temperature field, what is maybe the average temperature and other such measurements? And so thinking about co-design of communication, learning, quantization, everything together, was something that I learned very early on. And all of those lessons still hold up.

Rashmi Mohan: That sounds amazing. And so through that work, Anima, did you gain your interest in looking at these large data sets and starting to think about using machine learning to make sense of them?

Anima Anandkuma...: Yeah, absolutely. This went on as a path of curiosity, because in order to solve this highly interdisciplinary problem, it wasn't just virus communication, it was statistical signal processing. And that certainly shares all the foundations with the machine learning. But as I started to think about modeling the correlated measurements in these different sensors, what would be a good way to model high dimensional distributions? And that's what got me into the area of

probabilistic graphical models. And I connected with Alan Willsky at MIT and spent the last year of my PhD as well as a year of post-doctoral research there.

And so that got me almost seamlessly into machine learning because as I dealt more into the probabilistic modeling questions, it became clear that there is just such a wealth of problems to be solved. There is theoretical foundations that are yet to be resolved, there's, of course, huge practical implications. And so, it was great to see all that even before deep learning took cough, because it was clear that we do have now lots and lots more data available, and we need methods that are able to take advantage of that. And classical statistics wasn't the answer.

Rashmi Mohan: That's amazing. And also, it sounds like your really strong background in mathematics set you up so well to be able to think about solutions to these problems. You're often referred to as, I guess, the inventor of use of tensor algorithms to process large data sets. How did that come about? How did you think about that application?

Anima Anandkuma...: First of all, I'd like to emphasize that math is a very important foundation for thinking about these problems deeply. And in high school, the first time I took probability courses and saw that there was this foundation to deal with uncertainty, because until then you look at Newton's loss of physics, geometry, everything deterministic, where the answer there can be only one, but the world is certainly not deterministic. So seeing those tools of probability that are so powerful for us to make sense of this uncertain world around us, I think changed, it had a deep change in me in high school. And so I was always trying to take courses like that. I was doing statistical signal processing before that wasn't a compulsory course, so I was always going towards that. So in that sense, I was preparing for machine learning even before it was considered as a mainstream topic, or there were many courses offered when I was in school.

And so I'm very thankful to have that foundation, but I also want to emphasize that that shouldn't be a limiting factor. I can go into that in detail later, but as deep learning started taking off, there are certain things we won't be able to analyze in a mathematically rigorous manner, but that doesn't mean they don't work. Sometimes the theory and practice are not happening at the same time, and I follow both those journeys and see where it takes us, and a strongly performed empirical work where we are very careful about trying to understand the phenomenon, trying to test it rigorously is as strong as any theoretical work, and we need to do both. So with that in mind, coming back to your question about tensors and how I came about that, as I was working on probabilistic graphical models, natural question was, how do we model these high dimensional distributions with many, many variables?

We do have more data to learn, but still not enough to learn an arbitrary distribution over them. By which I mean, these are from, say, some real world scenarios where typically the actual distribution is lower dimensional. So think

of, for instance, if you want to model all the topics of conversation on some social media, let's say Twitter, these topics are constantly evolving, even the acronyms and the words are new, so there's no way we can just label a few topics and decide that this is all we need to categorize. So we need what we call an unsupervised learning method, so we cannot just go by labels of pre-categorized topics. And how we went about modeling these kind of scenarios is to think about hidden variables, the topics as hidden variables that relate to the observed distribution of words. And by looking at core currents of words together, we can ask, what are the underlying topics?

And so these kind of, what we call latent variable models, where we have hidden variables in addition to what we observe. So in this example, the hidden variables are the topics, the observed variables are the words we see in our Twitter feeds for instance, modeling that relationship allows us to then ask about efficient algorithms to extract information from them and to learn the underlying topics. And so we went about thinking about what are efficient ways to extract information. One simple way that people have been doing for a long time is to think about the occurrence of words in different conversations and trying to do a low rank decomposition, meaning somehow assume that if words tend to be spoken similarly for a topic, you can hope to extract the topics. But it turns out there's some deep algebraic reasons why matrix decomposition methods do not always tell you the underlying topics. What we call fundamental identifiability.

So doing it as matrix operations is just not enough to extract all the underlying topics that may be present in these Twitter data. What we instead did was to ask, "Matrices are not enough. What is the next object we can consider?" And that's where we happened about tensors, which are extensions of matrices to higher dimensions. And so here we are now thinking about co-occurrence of triplets of words or higher order concurrences, so it's not just looking at pairs of words, but higher order concurrences. And that intuition was deep because it said, "Now with tensor algebra, we are able to extract topics that wouldn't be possible with just matrix algebra."

Rashmi Mohan: I think that's amazing, that you were able to sort of extrapolate or hit upon that idea. I was wondering, Anima, would you be able to give us an example of how that might work or how you have used it?

Anima Anandkuma...: Yeah, going back to topic modeling, think about now looking at all the words that occurred together in a tweet. So let's say I tell you that in a tweet the word apple occurs. Just if I tell you only this information, what can you say about the topic? You can probably say that it's either the fruit or the company, but you don't know which of the two. So what if, now I told you pair of words, like there's apple and orange, but orange is also a company, so there is still some ambiguity. But what if I told you the words, apple, orange, and banana occur together in this tweet? You are now more certain that the topic that is talked about is a fruit, some conversation around fruits. So this is just a very simple toy

example to drive home the point that if you look at concurrence of multiple words in these different tweets, you will have much more information than looking at just pairs of words, concurring in these tweets. And that way you can extract the underlying topics much more effectively.

Speaker 3: ACM ByteCast is available on Apple Podcasts, Google Podcast, Podbean, Spotify, Stitcher, and TuneIn. If you're enjoying this episode, please subscribe and leave us a review on your favorite platform.

Rashmi Mohan: That's perfect. For a lay person like me, I think it's perfect for me to understand the concept that you were driving. I've also seen the video that you often share, Anima, about the one where the robot is doing the back flip and that you compare it with the dog, and you talk about reinforcement learning and how that's been maybe a cornerstone of some of your work. I was wondering, could you maybe explain that further and how you've used it?

Anima Anandkuma...: That's a great question, Rashmi. And things are developing so fast in that area. So maybe I should just go to what's happening now, especially with foundation models and our ability to come up with zero shot learning of high level concepts, and with that be able to do a variety of tasks together. So when I showed the video before a few years ago, I was commenting on how even if Boston Dynamics robot can do back flips very impressively, it's not as agile or as general as a dog. Dog can take our commands, learn new skills, generalize to new tasks very quickly, but that story is fast changing. In a paper that will have appearing at ICML conference over the summer, but that's already available online, called [foreign language 00:24:14] a visual motor learning, what we show is that now you can have models that can take in text and image commands.

So for instance, if you want to instruct your robot manipulator to rearrange the chess pieces in a certain pattern, it can go and do that. You don't tell that as a very laborious text command, but you also show the image and say, "Rearrange it this way." And so that's much more convenient. So this ability to give what we call multimodal prompts and interact with the robot in a very natural way is something we show is possible. And this way it can just generalize to a variety of different tasks. You can ask it to rearrange, you can ask it to put one blocks over another, you can do these different tasks in what we call zero shot manner, meaning you can easily come up with good performance on new tasks without having to explicitly train on them. And so that's really exciting, to see how far we've come along in such a short amount of time.

Rashmi Mohan: That is very impressive and definitely is a great view into where our future is headed. I'd love to understand from you, Anima, in terms of the areas that you apply, both machine learning as well as moving into AI and the tools that you've built around it. I know you've done some work around healthcare, you've obviously done work in the robotics world, and I know that there is recent interest also, there's a lot of work that you've done around climate change, any

of those topics would love to hear more, but also how do you determine where to go next? What drives your interest?

Anima Anandkuma...: I think that's a great question. To me, it's a mix of exploration and exploitation, to use the reinforcement learning terminology, because I am curious and I want to explore enough about new topics, especially in ones that I'm certainly not the expert, when it comes to climate modeling, or climate change, or even healthcare, working with surgeons, there is the deep domain expertise that I can never hope to replicate unless I put all my focus and energy on that. But the idea is to explore enough to learn abstractions. How do I go from something that is so complicated to enough of an abstraction to see what can machine learning do here?

And it could be something very straightforward. There's already an existing model, you just plug it in, or it could be something very deep in terms of requiring us to invent new algorithms, new frameworks, new architectures. And I've seen both, but you can only do that with enough exploration to learn new areas where machine learning is still a big challenge. And that's what I've done in all of these areas that you mentioned, and also trying to see how that unifies. We can go to each of these areas, try to solve the problems, but along with the abstraction comes the benefit that once you invent these algorithms, we can swarm many, many other areas as a result of that.

Rashmi Mohan: Absolutely. Is that your day job, Anima? Where you're basically thinking about, "Here is a problem that is unique and could be a great application for ML, and so let's seek out potentially new algorithms that can then be applied to other areas"?

Anima Anandkuma...: I wish it were so systematic, but most of the time it starts with conversations. And this is what I think, especially the value of in-person and chance meetings are, whether it's at Caltech, whether it's Nvidia or other conferences and events. It's learning new things and extracting knowledge from all that. And many times also it's scientists reaching out to me to ask how machine learning could make a difference. So it's a combination of all this, but what I always try to do is to learn enough about it and keep my curiosity open that even if it's not something I can help today, maybe in the future, this is something that can be solved. When it comes to specific topics, you mentioned in healthcare, we've been working on a number of areas. In particular, I've been collaborating with Dr. Andrew Hung from USC, and now he's recently joined Cedars Sinai here at LA who works on robotic surgery.

And so he's a robotic surgeon. He has wealth of data and expertise in terms of collecting videos, both of life surgery as well as students being trained to become surgeons. And our goal was to ask, can AI understand these gestures? And this is such a safety critical application, although one day we may see AI completely take over human surgeons, I don't think it'll be anytime soon, and it requires very careful testing. So instead we wanted to ask, how can AI help



today the current surgeons, as well as future surgeons who are training? And so this is where we are developing now systems that provide real-time feedback. As the surgery is happening, can you provide feedback? And we've started testing this on trainees for doing it on 3D printed tissues. And with safety critical applications it's also a great lesson on how we want to start in cases where safety is not critical and build our way up.

And it's great to see so many interdisciplinary areas come together. Of course, computer vision and video understanding is important, there's also kinematics and robot control, but as well as the human interface aspect, when you are having surgeons performing these careful maneuvers, how do you ensure that you give them the right feedback? Don't distract them. And we are also personalizing that because some people prefer audio feedback, some may prefer a quick vibration. So this is something that I think is such a nascent area. And so by working closely, we are bringing back this lesson of how do we bring AI to safety critical applications? How do we think about the human interaction aspect as we're delivering real time feedback?

Rashmi Mohan: I wonder, Anima, all of these areas sound absolutely amazing, they sound like great deep areas of research. Perhaps you dream for somebody who is in computing research, do you feel like these opportunities come your way because you've been able to... You have the credibility of having been at very strong academic institutions, as well as having understanding of how these things work in industry, and what does it mean to actually build a product out of some of these concepts?

Anima Anandkuma...: I think, both Caltech and Nvidia have been just great places for me to connect with amazing people. And having, of course, the history of working with interdisciplinary areas, I think will encourage people to approach me more. That's probably the case. But on the other hand, I'd also think of the conversations these days and the connections these days can happen much more easily. Before Twitter made big changes, we used to have so many technical conversations on Twitter, and so we need forums like that where it becomes much more democratized and easy to connect researchers across different domains, across different age groups, different institutions much more seamlessly. So we need a replacement, to be honest.

Rashmi Mohan: For sure. And I'm sure the pandemic had its own share of challenges thrown in with us not being able to meet. I'm sure there were many conferences that's probably a great hub to meet like-minded folks, or folks that are working on interesting areas that you'd like to collaborate with, and you don't have the ability to do that. Hopefully, now with things returning back to in-person, there'll be more of that.

Anima Anandkuma...: I do think certainly the in-person brainstorming has its qualities, but at the same time, the pandemic showed what can be done remotely, and especially people with families, or having long commutes, disabilities, all these aspects that may

prevent them to be always there in person, we showed that we can also be very productive. And this was also the time where we were able to do a lot of interdisciplinary collaborations online and communicate across different areas, because we had started developing this framework of what I call neural operators, which are fundamentally a new way of expanding upon standard neural networks to ones that can handle different resolutions, different shapes of grids. And this is very applicable for solving partial differential equations and other phenomenon continuous domains. You mentioned climate models, weather forecasting. Weather is a phenomenon that happens all across the globe, we may discretize and have measurements only at a few points, but the underlying phenomenon is continuous.

And so we've developed frameworks that not only take in data at one resolution, that can ingest trading data at different resolutions and more importantly can tell you predictions at any location, at any different resolution. And I think this is fundamental because this allows us to now think about, both a theoretical framework where we show we can accurately capture these kind of complex multi-scale phenomena, but we can also see practical gains where we've seen orders of magnitude speed up. In the case of weather forecasting, it was about 45,000 times faster. In the case of climate change mitigation, where we modeled how carbon dioxide would be pumped underground into wells with water, we were able to do this multi-phase flow modeling very accurately and very fast, about 700,000 times faster than current numerical simulations. And so seeing both this theory get developed, but taking that to practice and working with so many amazing researchers and collaborators with deep domain expertise has been very fulfilling.

Rashmi Mohan: Those are some stupendous results. And to see that kind of results come from the work that you're doing, I'm sure is very gratifying. What would you say, Anima, are the biggest challenges then that you face today from your perspective in this field of work?

Anima Anandkuma...: I think AI for science is something that we're only getting started now, in terms of bringing all these developments that we've seen in generative AI to the scientific domains. When we think about text models like ChatGPT, or image models, like image diffusion models, we see such impressive results. But right now, they're still in the creative realm, they don't have guarantees of being physically valid or correct. You may generate a very fashionable looking shoe, but it's not clear in the real world it would be comfortable or healthy on your feet. And to do that, we need to incorporate the physics, we need to incorporate what I call multi-scale structure. So you really need to accurately model the physical phenomenon, or chemical phenomenon, or biological phenomenon. And that could be equations, like I mentioned, like partial differential equations.

It could be symmetries, for instance the weather is happening on a sphere, on the globe, not on a flat surface, so those kind of geometries and symmetries we

need to incorporate. So we need to think about all these domain specific constraints and knowledge and how do we combine that with learning frameworks like neural operators that I mentioned. And that brings about some deep optimization challenges. One thing that has happened in deep learning is to think of optimization as very simple, stochastic gradient descent, Adam, all of these optimizers are readily available in the software packages, and they mostly work well. You have enough data, they work well. But when it comes to scientific modeling, we need to be very mindful of the constraints because maybe there isn't enough data. In fact, in the case of climate modeling, we have no data of the future. That's impossible.

But we have equations and we have data of the past. What was the weather over all these decades since we started recording it? And so how do we take that and extrapolate it under climate change, under now a new level of carbon dioxide? That extrapolation requires us to use the knowledge of physics. And so my one line answer is, we cannot escape physics, which means we cannot escape optimization challenges. And my earlier work when I worked on tensor algorithms was to deeply study the optimization landscapes of these non-convex surfaces and landscapes. So we need to go back to thinking about those fundamental tools and how we can improve upon it. So in short, there is now new challenges that standard deep learning doesn't address.

Rashmi Mohan: Thank you for that really detailed answer. The other thing I also think about is, as you are trying to answer some of these questions that none of us have thought about, and obviously many will rely on in terms of guidance, in terms of decisions, I'm guessing governments and agencies will be interested in the work that you're putting out. No conversation of AI can go by without talking about just ethics and AI. And I know that that's a very deep area of interest for you as well. I would love to hear your thoughts around that. Why is that everyone's responsibilities? How do we collectively take ownership of defining and solving problems around ethics and AI?

Anima Anandkuma...: That's absolutely right, Rashmi. It's much more urgent now as these models are now in the wild. So we have ChatGPT not always giving you the right answers, and the ease with which misinformation can flow and is already taking over large parts of internet is disturbing. And I work with Mike Alvarez here at Caltech, where we recently launched a new Center for Science Society and Public Policy. And that's very timely because we want to understand how we can test these generative models at scale. How would you even test them for bias, or toxicity, or misinformation? Because it's not just giving one input or one dimension. So it's highly complicated to test this across a distribution of different inputs that others may be using. That's one aspect. The other one is of course, how this interacts with how humans perceive this knowledge, how humans interact.

Mike Alvarez works on social media conversations, and we worked with him on using the earlier tensor algorithms for extracting topics at scale very quickly.

And now as a next stage of that, it is to understand how misinformation at scale can be tackled, and how do we provide everybody on this planet with the right tools to glean what is real and what's not? So I think this is something we need to figure out as a community. That's one important aspect.

The other is, of course, as I think deeply about new algorithms, that ideally come with some kind of guarantees or some kind of principled approaches, the more we build into it, the more we can trust the outputs, whether it's robotic surgery, where if you bring that to life surgery, that is very critical. You need strong guarantees, it's not just a few test cases that will convince anybody to take it there. So we need to architect it in a way where we can prove that it is going to do the things we expect it to do. And same with self-driving cars, drones, all these different areas. By working on those areas, I think it also forces us to tackle these algorithmic issues head on, and that'll help us bring towards us trustworthy AI in all domains.

Rashmi Mohan: And these are difficult, difficult problems to solve. They inherently make your job so much more challenging. And that's one thing I observed through all of your conversations, Anima, through this interview as well as previous ones, is you have this aura of fearless leadership. You champion a thought or an idea and you stand by it. I would love to understand how do you build that sort of a muscle.

Anima Anandkuma...: Thank you. That's very kind. I think, to me, it's all the support I have. As I said, growing up, it's my parents who encouraged me to take risks, think about preparing for this entrance exam on my own, but also now my husband, Benedict, is always just there for me as I think about very hard challenges and encourage me to go for it.

Rashmi Mohan: That's wonderful. Really happy to hear that. And thank you, thank you for sharing that. It certainly is an inspiration to me, and I'm sure to many of our listeners, because oftentimes some of these problems, again, if you've never encountered them, they're going into domains that are completely unexplored, can seem pretty daunting. And to, I think, believe in your skills and believe in the ability to break down a problem and just go for it, like you say, is probably the best way to move forward.

Anima Anandkuma...: I think of that as an adventure. And I also, for researchers, for anger and who may want to take less of the risks, you can think of it as a portfolio approach. So you could try something that is very new as an adventure, but you could also have another project where you know things will mostly work out. And so that way you hedge and if the new adventure pays off, that's great, otherwise you still learn something.

Rashmi Mohan: Absolutely. For our final byte, Anima, I would love to understand what are you most excited about in the field of AI over the next five years?

Anima Anandkuma...: So I'm very excited about the future in this era of generative AI and bringing that to scientific domains. Can we think of one universal foundation model that could understand different areas of sciences? What would that look like? What would be the kind of data that would be available for that? How do we bring in the physics knowledge and all of these domain constraints? How do we solve the optimization challenges? How do we incorporate stochasticity at such complex domains and the multi-scale nature of many of these datasets? So to me, I think we are just getting started here.

Rashmi Mohan: Wonderful. Lots of mathematics ahead of us. This has been such a-

Anima Anandkuma...: Maths, engineering, science, everything, social sciences for human interaction, everything.

Rashmi Mohan: Perfect.

Anima Anandkuma...: We need all of it.

Rashmi Mohan: Wonderful. This has been such an inspiring conversation. Thank you so much for taking the time to speak with us at ACM ByteCast.

Anima Anandkuma...: Thank you, Rashmi. It's been such a pleasure.

Speaker 3: ACM ByteCast is a production of the Association for Computing Machinery's Practitioners Board. To learn more about ACM and its activities, visit [acm.org](http://acm.org). For more information about this and other episodes, please visit our website at [learning.acm.org/B-Y-T-E-C-A-S-T](http://learning.acm.org/B-Y-T-E-C-A-S-T). That's [learning.acm.org/bytecast](http://learning.acm.org/bytecast).