



# AI Ignition

## Ignite your AI curiosity with Vivienne Ming

**Introduction:** 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 will 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, and I’m the executive director for Deloitte AI Institute. And today on AI Ignition, we have Dr. Vivienne Ming, a theoretical neuroscientist, entrepreneur, and author. Dr. Ming cofounded Socos Labs, her fifth company, an independent thinktank exploring the future of human potential. Previously, she was chief scientist at Gild, an innovative start-up that applies machine learning to predict optimal candidates for technology jobs and to bring meritocracy to the job market. I am so excited to speak with Dr. Vivienne Ming.

Welcome, Vivienne. It’s great to have you on today’s show. How are you doing?

**Vivienne Ming (Vivienne):** I am doing wonderfully well. Life is overflowing with good things, even in the midst of the world being a bit of a mess.

**Beena:** Yeah, that’s an understatement. We are going to specifically speak about AI today. You are a visionary in this space. We met a few years ago, and I admire all the work that you do. But I actually want to start with—let’s take a step back—how do you define AI? What does artificial intelligence mean to you, and can you use an example to explain to us?

**Vivienne:** Absolutely. So, for any of the nerds listening, we could get very technical and I could certainly get into the differences between AI and machine learning and deep neural networks, but let’s give a very concrete example. My first-ever machine learning project was over 20 years ago, when I was an undergrad, I worked on a project, actually sponsored by the CIA, to do real-time video analysis of people’s faces to tell whether or not they were lying. Now, needless to say, you can’t really program a bunch of if/then statements, if the person tilts their head to the side, well, how do you even know they tilted their head? So a lot of what we think of as classic AI is not useful, so this new class, at least 20 years ago, it felt newish, using various forms of artificial intelligence, including machine learning called a deep neural network, to be able to

analyze what was going on in the image, all the way down at the level of the pixels in the way that's a little bit analogous to how our brains work.

So, for me, artificial intelligence, I often go with a very nontechnical definition—it's any system that can make decisions under uncertainty. So, when there's no right answer, is this person happy or sad? Well, you don't really know, but given their smile or they were actually crinkling their eyes, then there's probably an answer that's more right and one that's more wrong. And it was so amazing to be thinking I was going to be a neuroscientist, sticking wires into cats, and to suddenly get exposed to this idea of an intelligent system that could recognize whether someone was smiling or frowning or disgusted. It felt so human. Now once you get under the hood in AI, it turns out there's some really big differences between what it's doing and what people are doing, but that's kind of my go-to.

**Beena:** Yeah, I'm a big believer that it's about machines augmenting human intelligence and I've heard in one of my first interactions with you, you shared the story about how you used AI with treating your son's diabetes or tracking it, monitoring it. Can you share for our audience that story of how you applied AI in that context?

**Vivienne:** Absolutely. At that point in time, I'd been an academic for many years and I built AI specifically to focus on the brain to understand and model what was going on all the way from little neurons to whole-scale brain activity. And then I started a company, actually with my wife at the time, in education, so taking it from the level of brains to understand brains learning out in the world. Could we build an AI that could watch children learn from each other, and I guess our crazy ambition, end all high-stakes testing. If you could understand them while they learn, why would you need to give them a test? I think the success of the test-taking industry is an indicator that my first start-up idea didn't end up quite the way I thought it would, but in the middle of all of that, my son one day had what seemed like the flu and by the end of the week, he couldn't stand up. We rushed into the hospital and it turns out he has type 1 diabetes, which is, for those of you who may not know, an autoimmune disease, not the more common type 2, which you get when you're a bit older. This is sometimes called juvenile diabetes. And at the time he would have just turned four, so given what we know about it being an autoimmune disease, he probably had it for six months already just slowly destroying the cells that make insulin. If you don't have insulin then your body can't take sugar from the blood and put it inside the cells, which means in a sense you starve to death, even though you might have a stomach full of food. So, he was slowly losing that ability to make insulin, and by the time it started, it was too late. There is no cure for type 1 diabetes. So, when we finally get to the hospital, needless to say, this is emotionally a very difficult time, but when we leave my wife and I, both scientists, thought this was going to be something that we were going to tackle head-on. So we started collecting all of this data about our son, like everything he ate to the gram, and by everything, I mean carbohydrates to the gram, proteins to the gram, fats to the gram, everything in a giant spreadsheet, we were crashing that spreadsheet over and over again. We would throw in his heart rate, his blood glucose readings, how much sugar is in his blood, we would put in everything, did he have the sniffles, was he grumpy that day. And then we thought the doctors, of course, would love us. Wouldn't you love to have a patient that brings in an inch thick of spreadsheet with all of the data about their symptoms? Well, in

retrospect, of course, we were foolish, but I naively thought, this is wonderful, think of all this data that you'll be able to use. They were angry with us, angry that we were wasting their time, so that ended up being an even bigger shock to me than the original diagnosis. The fact that we literally had tens of thousands of data points and instead, they asked us to handwrite 15 numbers, five days, morning, afternoon, evening, and then they just eyeballed those numbers and came up with the treatment plan. And I was just stunned, so that night I bought a book on endocrinology. I may be a neuroscientist, but I probably slept through the classes on endocrinology. And actually the very next day, we started working on hacking all of his equipment, sent the data that his own body was producing to my personal server, turns out I broke all sorts of federal laws, and I took a model that actually was from my world, a model of predictive coding in the retina in your eye, and I readapted it to try and predict his blood glucose levels in the future. And needless to say, it's not a cure, but I got to tell you, anyone listening that might be a parent, the one thing I would wish for you in your life is that there is some moment, a moment of tragedy will almost be inevitable in everyone's life, that there is that moment and you are the one person in the world that can do something about it. So, I got to build this system for my son and it actually worked. I actually got to wear a pair of Google Glass to the White House. They were not happy about that, but I ended up meeting the president, taking the only presidential portrait ever taken with Google Glass, these wearable smart glasses. And the reason I was wearing them is because I was in Washington, DC, and I got an update from my own AI system saying that my son's blood glucose levels will go low in a half an hour. And so, I sent a text message to my sister back in San Francisco, who was taking care of him, and it didn't go low. And the best part of it is, we were able to just give it away. So I think that gives some sense of two things. One, again, an application of AI, notice I didn't really go into equations that much or anything, a very human experience, but also what happens, what is the way people should be thinking about utilizing people and AI together. That was just me. In fact, I gave that algorithm away to a bunch of companies with no strings attached, and a couple years later, one of them got the first product onto the market and I tweeted at them a genuine congratulations followed with "and it only took you four years and \$40 million more than it took me." Imagine if you had a workforce full of, forgive the arrogance here, people like me that all by themselves could go out and tackle some of the most challenging problems in the world not because AI replaced them, but because they were being so deeply augmented by them.

**Beena:** I love it. I absolutely love it. And I know you've worked across several industries and you've looked at AI much more broadly. Can you share some other examples from other industries as well?

**Vivienne:** Yeah, so probably my biggest claim to fame other than the project with my son is that I was the chief scientist of a company called Gild. Gild was one of the first companies trying to apply artificial intelligence to take bias out of the hiring process. And this was many years ago. And I will freely share, this is a much harder problem than we ever appreciated. It's a perfect example of just having AI does not magically solve a problem because if the recruiters used our tool in the exact same way they used every tool before, then our magic AI doesn't fix anything. Nonetheless, we built a database of 122 million working professionals, and again one of the first systems that predicted not how well you do in the job market, are you going to get hired,

but the quality of your work. So, we did this by analyzing code samples from over a million software developers and then pairing that against everything we knew about the rest of their lives. We did it by analyzing the sales history of hundreds of thousands of salespeople and again comparing that against what they knew about the rest of them, their work lives. And it was so much fun, we ended up on the cover of the *New York Times* and *The Atlantic*. It was one of those crazy, head-spinning runs. What AI can do is fundamentally change the economics of that solution. So, for us, for example, instead of being a tool to people, you wait for someone to apply, then maybe you give them a test or an interview, we were a sourcing tool. We could crawl out through the Internet and just find people because we could do the equivalent of an executive search report on everybody. Were they perfect? No, but actually that's one of the, let's call it, ethical AI choices we made was only say positive things. So, our opinion was everyone is right for some job. We should just sell everyone for a job, instead of giving people reason, for example, to not hire someone when at the heart what we really found was predictive, and again this won't sound like hard numbers in AI, but this is the sweet spot in many ways. We found what was predictive was if we could tell if someone was resilient. If we could tell whether they had good perspective-taking skills. We actually found social skills were equally predictive for the quality of work by software developers and salespeople. And we would use these AIs to just begin to surface whether we saw that. For example, by analyzing Q&A sites, like Stack Overflow, where people go to try to learn new things about coding. We actually found that people with the highest social scores, the most upvotes, the most badges, not predictive at all, but if we actually analyze the questions and the answers and, again, using AI, because you could hire an army of people to do this. Again, a good example of something where you have an intuition that this would work, but you would never have the money to just hire enough people to go do it. But we could build an AI that could then do that for you. Then a truly talented recruiter could take all of that information getting surfaced by the AI, find the right people that were the right fit for the company, and we took it one step further. We even templated a letter that said, "Here's why you should come work for my company." And then unfortunately, this is the lesson learned. The vast majority of recruiters didn't use it. They would open up a profile in our system and they'd look at your name, your school, and your last job, exactly what recruiters have been doing for literally 100 years. It's been well-studied. So, they—two-thirds of recruiters in our system—used our incredibly expensive and very fancy AI to do traditional recruiting with all of its inherent bias and everything else. We saw those indicators of where truly creative people did exceptional things and the results were amazing, but again it tells you this is fundamentally a human system. This is humans and machines interacting, and the AI just can't magically make these problems go away if the humans aren't invested in that as well.

**Beena:** Yeah, so true. So, I led one of these projects where we built this best predictive model for a factory floor machinery. It was on the critical path for production and it was 99.999 percent accurate, extremely great AI solution from a data scientist perspective, but none of the workers wanted to use it. And at that point, that project is a failure because if your users don't adopt it, how do you drive that cultural change to actually change that behavior to be able to use these powerful tools? I love this example that you gave. Did you try to change the human behavior? How did you tackle that?

**Vivienne:** We really tried to invest in that idea. I've learned—I've now had the pleasure since my days as an academic, I've started six companies. I've been the chief scientist at three others. So, I've had my turn trying to figure out how to get people. I sometimes describe myself—I invent broccoli. Well, it turns out nobody actually wants to use or eat broccoli. They want to eat brownies. So, then I have to spend a lot of time figuring out how to bake my broccoli into a brownie. My system—these are things I truly believe in—but it doesn't matter if nobody uses it. So, for example, here I mentioned that we have these examples, these companies. One of the best was Gild itself. We did a voluntary project for a nonprofit called Girls Who Code. So, many of you've probably heard of them before, they're this big one teaching young girls to get into coding and STEM in general. I've actually done some work with them and it's been a blast. So, they asked if we could help. They wanted to get out of New York and San Francisco where they were heavily represented, and they gave us a challenge—find us a set of teachers in Detroit to teach six weeks of coding for free to a bunch of teenage girls and as many as you can because they wanted to run as many classes as they can. Well, needless to say, that sounded a bit like this Herculean task, like it was given to us intentionally to fail. So, we asked for one thing. Here's the very human side of it. We asked all of the board members to sign a letter. So, you had all these famous CTOs and tech celebrities signing letters, and we used our system. We found something like 27 people in Detroit that met our criteria. Then, we used our own templating system to write them a personalized letter. We added the signatures on the end. I believe if I'm remembering my numbers from years ago, we got 20 out of 27 responses and something like 30 to 40 percent of the people we targeted actually took the completely unpaid job. So, we used it to build a relationship, again using AI combined with the best of what people can do, can be transformative. We would never have found the right 27 people without the AI. We wouldn't have known that they'd have been right for this job, that they have the skills, but the AI, at least as it exists today, is incapable of understanding someone so well. I would argue even GPT-3, for the nerds listening right now, is not capable of understanding someone so well that they could send a truly personalized note that would actually convince them to take six weeks of their lives and donate it to this good cause. So, there was a point in Gild's history where there was a real tension. Do we build what the customer is asking for? Isn't that the ultimate lesson of the market? Or do we build broccoli, but we invest in training our customers?

**Beena:** Yeah, I love the analogy of broccoli to brownie as well. For the manufacturing plant, actually the way we made it work was we did gamification. So, incentivizing for eating the broccoli even though they wanted brownie. So, we gamified it so that every time you use an AI solution, there were some points given, which you could then redeem for gifts and so on. So, I think AI adoption and driving that change is an ongoing challenge. So, Vivienne, we did a survey. We recently released a survey, and in that we found that we were looking at companies who are just getting started in AI, the ones who are invested in AI, and the ones who are far advanced. So, there's a large population that's just still dabbling in AI in the enterprise industries that we serve. What suggestions do you have for CEOs or board members that are fairly early in their AI journey for exploring the implications of AI for their business? How should they think about AI? If you are a very large, traditional, old enterprise, how do you think about AI in the context of your business? Any guidance?

**Vivienne:** So, you have all of this structure in place because you're trying to manage a very large organization, you're trying to mitigate risk, and then someone like me comes along and says, "AI is all about a creative workforce." What's the point of hiring brilliant creatives if you then tell them what to do? You should tell them why to do it, get them aligned around the company vision, but it's interesting, we're here in the middle of lockdown in our global pandemic, and as perverse as it is, it's been a blast for me because it's allowed me to run highly controlled experiments on the social networks of companies because all of our conversations are happening through this logged and recorded medium. So, we are able to look at what's going on, but what's the one thing we're afraid of disappearing when no one can meet with each other in person? Well, one of the main ones that comes to mind is innovation. How do you have innovation if everyone thinks of it as this random chance, this smart people crashing against one another and new ideas coming out of it when no one gets to crash into each other. Everything's scheduled and structured and you have these terrible camera lenses to talk through. So, we're able to actually look and see, is innovation still taking place and what sort of organizational structure really supports it? By the way, if it's not clear where I'm going with this, I think this is fundamentally also how we need to be thinking about AI and people leveraging AI. So, what we find is if you have small teams with very flat hierarchies and a certain degree of isolation, so they occasionally talk to your other small teams, but they are very focused on themselves, we still see innovation happening there. In fact, we've been able to do some experiments where we could create this artificially without people even knowing by slowing down communications coming in and going out from that team and seeing, wow, actually it increases the rate at which they develop and enculturate a new idea and then there's a right moment to spread that to another team and then let it sort of percolate a little bit and then finally spread it out to everyone else. So, in other words, it's like the opposite of this complex old-school hierarchy, very local decisions, very flat, lots of small teams, not how a lot of organizations work today. And yet, if you think of some of the examples I gave already of developing the treatment for my son's diabetes for the AI system we built. I built a system to reunite orphan refugees with their extended family members based on facial analysis. I developed another one because I have this curiosity about gender wage gap, and it was able to crawl through the websites of 60,000 companies and collect data, including all the way down to facial analysis, to look for variables that were highly correlated with wage gap. I've done scores and scores of projects like this, and what's most consistent out of all of them is—that gender one, it took me two days all by myself. Could you imagine even 10 years ago, maybe even 5 years ago, doing an extensive economic analysis of the business behavior of 60,000 companies.

**Beena:** So, Vivienne, having seen this whole AI space evolve and where companies are heading, where some of the challenges are, what do you think AI or the world of AI will look like 20 years from now if things go on the trajectory it is at right now?

**Vivienne:** If it continues the way it is, then I don't say this to inspire some political passions in anyone, but the way AI goes right now, it does drive a lot of economic inequality. Again, this idea that we really split apart. We create this tool, which is an amazing driver of productivity for an elite few, and then everyone else has very routine jobs and they're kind of a pair of legs to walk in AI around the room. That to me is not healthy. It's not economically healthy because it doesn't actually tap the capacity of any country. We lose out so much if we just drift forward

thoughtlessly with, “Can’t I build a chatbot to replace that person?” On the other hand, and this is my thing, I’m not a futurist, so I see some choices we can make. If the choice is augmented intelligence, if the choice is there’s more investment in training, there’s more investment and time put into organizing our management structure and in the technology itself, but the end result is something amazing. I’m not a utopianist. I don’t mean to say that everyone tomorrow becomes a super-doctor or super-analyst. But I do look at an economy that comes much closer to tapping the real potential that people have. Some kid in a favela in Rio, in their head is the cure to my son’s disease. Better than mine. Not my treatment—the cure to diabetes, the cure to COVID—and the likelihood of that kid actually growing up to bring that cure into the world, having the necessary education, the employment opportunities to do that is small, but in a world where we don’t need so much of that rote labor anymore, in a world in which we invest in human capital because that’s what makes AI valuable—human capital. Then we turn that around and that kid or a kid in Bangalore or down the street from me in Oakland, they do something amazing, they do something world-changing, young. In fact, when I think about my kids, I think they’re going to have an amazing life almost no matter what, despite my son’s challenges, despite the infliction of having me as a mom, my kids are pretty much set for a wonderful life. The thing I can do that would help them the most is to help every other kid. And I know that sounds like a complete non-answer to your question of what comes next with AI, but again you cannot separate out the human capacity question from the technological question. The value of AI, the true complementary productivity value in economic terms, is the way it can complement creatives.

**Beena:** Love it. So, how can people stay connected with you, learn more about all the amazing work that you’re doing? Can you share?

**Vivienne:** Well, I would love to. So, after being an academic and then starting a bunch of companies, I found myself able to sort of decide whatever I wanted to do in my life, and I thought what I really loved about both of those things was solving problems. I never really cared whether my companies made money. Don’t tell any of my investors. I never really loved writing grants, by that I mean I loathed it. So, I thought what if I built an organization, we’re called Socos Labs, and people brought us problems, and if I thought my team could make a meaningful and unique difference, then I would pay for everything and we try and come up with a solution to the problem, and if we did, we just give it away for free. So, we don’t exclusively work in AI, but I’ve got to admit, I may not be a one-trick pony, but I’m good at this. So, we use a lot of artificial intelligence in our work and it’s all over the place. Dr. Ming, my daughter has 500 seizures a day. Please save her life. Or my son can’t enter REM sleep, which is, boy, that is unusual, to say the least. Or my company doesn’t know how to retain black women. Can you help us figure out what’s going on? Or our country doesn’t know what to do about education policy. So, we get to work on this wild range of problems. Turns out there is no shortage of people in this world that will take me up on the offer to help solve their problems for free, but it means I get to work on so many different things, my team gets to work on so many different things that we just end up with a lot of wonderful stories, stories that hopefully are helpful to people making tough decisions, whether inside companies or just in their lives. And so, we decided every time we finish a project, as you can see these posters here behind me, every time we finish a project, we put a little movie poster together and we write a story

about doing the actual work. We call them episodes. There's the episode of my son's diabetes, and the episode in which we developed a system to predict manic episodes in bipolar sufferers, and the episode about the work in wage gap, and so on and so on. If you want to see these, they're all there for free. Just visit [socos.org](http://socos.org), and you can read about my crazy projects in so many different spaces and the little mad science incubator projects we run and everything else that we get involved with here. And I guess the one nice thing about that is by telling not just what we learned, but how we learned it, then it's not me just arrogantly telling people how they should be living their lives; rather, it's an honest acknowledgment that when we started the project, we had no idea what we're doing and we tried this and failed and I did that and it made things worse. Oh my goodness! How often it seems that a seemingly good idea does the exact opposite of what you intended. So, if you visit [socos.org](http://socos.org), you can read about all these projects and see the kind of work that we do. And some of these things we even get to offer for free right from our own site. Some we give away to partners that are able to make use of it in other ways. All we really care about is that more kids like that kid in Rio have a chance to have an amazing life so that years from now, your company can pay them an obscene amount of money that they entirely deserve to solve the kinds of problems that are going to confront you. So, it's wonderful. It's the worst business idea I've ever had in my entire life. It turns out that revenue is generally a good idea for a business plan, and yet I hope I get to do it for the rest of my life.

**Beena:** I hope so too. This is amazing work. Vivienne, thank you so much for being on the show, and to our audience, thank you for tuning in to AI Ignition. Follow us at Deloitte AI Institute for more AI. Thank you.

**Conclusion:** Thanks for tuning in to 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](http://deloitte.com/about).