



AI Ignition

Ignite your AI curiosity with Greg Brockman

How do we take responsibility for the varied uses of 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 will take a deep dive into the past, present, and future of AI, machine learning, neural networks, and other cutting-edge technologies. Here is your host, Beena.

Beena Ammanath (Beena): Hi, my name is Beena Ammanath and I am the executive director for the Deloitte AI Institute, and today on AI Ignition we have Greg Brockman, the co-founder and CTO of Open AI, a research organization that aims to ensure that artificial general intelligence benefits all of humanity. Welcome to the show, Greg. We are so excited to have you. Greg, I am so happy to have you on AI Ignition today. Thank you so much for joining us.

Greg: Thank you for having me. Excited to be here.

Beena: So, let's start with the current status of AI. AI is such a broad topic, and you hear so much in the news about it, everybody knows what you are doing at Open AI. What do you think is the current status of AI today?

Greg: So, I think AI is definitely a field in transition. I mean, if you look at the history of the field, I think it has gotten a reputation as the field of empty promises. And you look back to the 50s and people were trying to figure out, hey maybe the hardest problem is chess, and if we can solve chess, then we are going to solve all of AI and we'll have an artificial human and it will be great. In the 90s, chess came and went, and it didn't really feel like our lives changed at all. I think that 2012 was really the moment that things started to transition with the creation of AlexNet, which is the first sort of convolutional network on GPU that was actually able to just meaningfully outperform everything else. It's the first time that like all of these, you are a human, you think hard about a domain, you write down all the rules, all those systems just were paled in comparison to what you could do by just having a neural net or just kind of an architecture that can absorb whatever data you show it, and suddenly that's the best way of doing things. In image recognition, and everyone kind of looked at it and said, "Okay, that's a cool trick, but hey it's never going to work for language. It's never going to work for machine translation." And it kind

of has. Basically, the whole past decade has been, okay we are starting to see that for all of these areas that feel like machines have never really been able to touch them, things like language, things like understanding what's in a video. We are starting to be able to crack it, and the way we crack it is it's all with the same piece of technology, it's just a neural net.

So, that's all the background, that's where we have been. And I think GPT-3 is also a moment that shows a sea change in terms of what the tech can do. This is the first time where you take one of these big neural nets. It's basically just an architecture you dump all this data into, and suddenly it's actually economically useful. Like people actually want to use our API to just talk to our model. Some people do it for fun, but actually many, many people are building businesses on top of it. You have various applications ranging from assistants that help you write better copy, whether it's A/B testing and we actually have had a number of users do pretty rigorous tests with GPT-3 versus human copywriters and found that actually GPT-3 does pretty darn well. You have people who use it for games, things like AI Dungeon where you actually have an infinite dungeon like a DND style game that is generated by this model, and there is a huge range of other applications, such as people who are trying to help tenants who receive eviction notices or other legal proceedings in very complex language, trying to summarize it and point out the important points for people who don't have access to their own legal counsel. I think there is this huge variety of different things that people can do once you have a machine that can start to actually interact, meet humans in the language that we all operate in. So, I think that it's this exciting moment where you can actually train one of these networks and they are actually economically useful.

Beena: Yeah, so true, and I am going to date myself here because I studied AI back in the late 80s, early 90s and that time it was mostly theory. We didn't have access to massive amounts of data or the compute power that was needed. So, it was mostly what was written in the books and ideas around it. I am not kidding but back then even personalized marketing, which is so prevalent now, was considered to be too complex for an AI or any kind of software to do.

Greg: But I also think that all of that work laid a really good foundation. There is always this joke that everything was invented by Schmidhuber in the 90s and it's not entirely false. That there are a couple of ideas that are out there now that are kind of totally new and unique, and partly because we saw like when you actually build things in this certain way and we have this particular hardware architecture for GPU that maybe doing massive parallelism in this particular way is useful, and that is where a transformer comes from. But most of it I think is actually, we basically have 70 years' worth of people thinking super hard about the same neural net piece of technology we're using, and so we can just build on the shoulder of giants.

Beena: That is so fascinating because I think in a way we were unencumbered by restrictions. So, it was just ideas that we're spouting on, what if this was real, what are all the possibilities. That's a great point. I never thought of it that way.

Greg: The thing that I find most interesting about this field is that it doesn't comport with the way that we think science should go. And you can really look back to the history of these debates between the connectionists, the people who basically said, "Hey, these neural nets which are kind of this like if you squinted at a model of the information processing of the brain," and then the people who are the symbolic systems people, who have said, "We humans are really smart, we know how systems work, let's just write down all the rules and try to like really engineer a system that operates like we do." And there were these great debates of which one was going to win, and it really looked in the 60s like these

connections people were a bunch of hacks. There was really this whole campaign waged with books like the Perceptron and things like that to discredit all of them and it really worked. And somehow the connectionist school of thought has just been this steady rising tide. We see that there is this smooth continuous curve of increasing progress that has lasted since 2012, but we actually did a study too where we extended it backwards and we found that actually there has been a smooth exponential in terms of the increased performance of these networks over decades now, and I think that points to something really real that's going on, and it's just so fascinating.

Beena: Yes. Now you are looking at core AI. How do you see AI creating new opportunities across industries?

Greg: I think that this is really the most important question. We think about AI, I actually think AI has a pretty bad brand. Like if you ask someone on the street what they think of AI, you are probably going to get some reference to a Hollywood movie or something like worrying about how it will impact jobs or people's lives, and I think all of that is fair. I think it is actually really important to confront the downsides, but I think that the reason we work on it, the reason we are so excited to build it, the reason that we have a whole—that Open AI as a company dedicated to trying to build advanced AI is because we see the promise that it can bring. I think that the real goal is to be able to build systems that can help us solve problems that are otherwise out of reach. You think about healthcare. So, I have a very close friend who has a very rare medical condition and every time he talks to a new specialist, they always bounce him to another specialist and it's a real problem. He is often bounced back to the same person and they are like, "Oh well, there is this thing that they didn't consider." Imagine if you just had one doctor who just understood the whole extent of human medicine. This is not something that feels outside the realm of imagination, but we have been totally unable to achieve it, and I see the direction we are going in so many fields is towards increased specialization. So, if we just had a way to integrate that information across domains, I think that we would be able to build much better, much more capable systems. Think about things like education. I think everyone has got a story of that one teacher who really influenced them who was so good, who really understood them, why can't all of our teachers be like that? And I think that the naive answer, by the way, is to say, "Okay great, that's what we want to build AI to do, that's what we want to really achieve is we just want to build the artificial teacher, build the artificial doctor," but I don't think that's actually quite the right thing, maybe that's the long-term picture, but I think that the immediate term picture is building tools that help our teachers become more effective, building tools that help our doctors become more effective, and building tools that help anyone who is engaged in some creative pursuit or some passion or whatever activity it is that they are doing, give them tools that make them more effective. And I think that's actually where we are right now. You look at something like GPT-3 or we have a new model called DALL-E, which you type in some text and it generates some pretty great images for you. There is the Steve Jobs quote of "Computers are bicycles for your mind." I think that these tools are kind of like rocket ships for your mind. It's almost like think about spellcheck, it's like none of us really have to be spellers anymore. We have removed the mechanics of how do you go from this idea in your head to words on a page, there is part of those mechanics that is just so much more efficient now. I think that where we are going with these certainly generative AI technologies that are actually working today. These aren't sci-fi, these are actually here. You are actually starting to see creative applications become way more effective, way more of those mechanics and the brainstorming and the trying to just have someone there who you are able to bounce ideas off of, I think those things are starting to become something your computer can help with.

Beena: Yes, and there is so much to unpack there, Greg. One of the trends that we have looked at is bespoke for billions. What you are speaking about, whether it's healthcare or education, how do you provide that personalized service but at scale. If you can provide personalized education, personalized healthcare by leveraging AI, it's exactly made for you, customized just for that individual. I think there is so much opportunity there. We have spoken about GPT-3 a few times. Can you explain where we are with GPT-3? What are some of the applications that you are seeing in the real world that might have seemed obvious to you, but some that actually caught you by surprise. Can you share a little bit about the applications of GPT-3?

Greg: First of all, it's a funny thing because I would say that building the GPT-3 API was probably the hardest project I have ever worked on, and part of it was because it was a super hard technical problem. We basically had this research model that we had to productionize and make run really, really fast, but part of it was because we didn't know if it would be useful. There are lots of people who have built APIs for various models in the past, people tried it with AlexNet and none of those businesses worked because I think that the models just weren't capable enough. You have a great image recognition system, I mean, okay, a somewhat great image recognition system that can recognize up to 1,000 categories and be easily confused if you do something that's not quite what ImageNet looks like. No one wants to use that, no one wants to pay for it certainly, it's not really going to move the needle. So with GPT-3, we were worried, maybe this just won't be a thing that's good enough, and we have been really pleasantly blown away by that. I think that in the early days we went around and we were sort of pitching this API to people I think back in January of last year. I was driving around trying to see if anyone would take a look at our thing, and everyone would look at the demo and they would be like, "Oh this looks cool, but maybe it will fit somewhere," and kind of put it on the back burner.

We had a couple of partners that really latched onto it. AI Dungeon is an example of a real early adopter who they actually had taken GPT-2 when it was available. Actually, I remember when I first talked to them, they said, "Hey, I think that all we really care about is just driving down the cost of inference," and it turned out that actually the new technology we created was so compelling that it was actually fine if it was even more expensive to run, and it just turns out that the capability had just sort of reached a level that even we were surprised that people found so compelling. I think that one other area where we were totally surprised was in No-code. I think this was something that a lot of people have seen because some of these people doing these demos ended up with some very viral tweets because it's very compelling. You just say to the machine, "Hey, I would like the code for a web page that has a red button and you click it and this thing happens and then it actually generates the right code for it." It's really fascinating and part of it is like it's a little navel gazing because obviously we ourselves are coders and a lot of the people who are going to consume this technology are coders, and so we are all excited about, hey, is there a tool that can help me code. But I think it really points to this area where, back to how consumed we all are by the mechanics of whatever field of endeavor we are in. Coding and typing at the keyboard and trying to translate from what's in my head and this natural language that we all think in into a way that the machine can understand, that is a tedious part of the process. You think about the whole history of computing has been trying to make it so that the machines meet us more where we are rather than us trying to type everything out in bytecode. So, I think that was really cool to see people actually start to build companies around it. We have seen lots of copywriting startups, and I think maybe the most interesting and exciting to me is to step back and say people are actually building businesses on this, it's not just like hey you build like a throwaway project at a hackathon. There are lots

of people that have done that, don't worry. But the thing that's really exciting to me is that there are people who actually have lots of customers and are growing really fast and that we are building a platform that is really sustaining an ecosystem. It has also been really cool just to see we have a developer community. It's about 30,000 people now who are in our developer Slack, and that's a lot of people that we are starting to touch across all over the world. So, I think that maybe for me that's the exciting thing is that there are all these applications across this wide variety of different domains, there are some where we just don't work yet. There are some where we are not ready to support them because there are such high stakes and that we kind of need to go in with kid gloves to start, but I think that in all of these core applications, we are just seeing an amazing amount of activity and it's really cool to see and support.

Beena: I am smiling because when I studied computer science, one of the first languages I studied was Pascal and then there was COBOL and assembly language, this heavy syntaxed languages. I get asked a lot by young students who want to study computer science, they are like, "Which computer language should I learn?" And I say, "Just wait for a few years, you won't need to understand the syntax or learn programming language, almost we will hit that point where the machine will write the programming language for you, you just have to articulate your vision and what product, what's your MVP that you want to see." So it's fascinating for me to hear that and how far we have come. I know GPT-3, you developed and launched this past summer and seeing this progress in such short time. But I had one more question on that. Unlike GPT-3's predecessors, GPT-1 or GPT-2, you chose to not open source the model or the training data set, opting instead to make it available through a commercial API. Can you explain why did you choose to do that? What was the reasoning?

Greg: There are a couple of reasons. There is one, which is the simplest reason in some ways, which is that we are trying to build a sustainable business and that we want to make sure that we can support that business through a revenue stream. But I think that the more interesting answer is that we have always had an internal conflict within ourselves trying to figure out how to really achieve the mission, because on the one hand we really believe that AI has all these benefits that are just out there waiting to be captured, and it's really important that those are brought to the world, brought to, in our charter we say, all of humanity. On the other hand, we know that these technologies are dual use. These models are so general, they can be used for positive, they can be used for negative, and it's quite hard to design a model that can only be used for one but not the other. So, you are really stuck in this dilemma. You open source a model, there's no undo button, and we really felt this keenly with GPT-2. That's the first time where we said we have a technology that we are not sure about, and we chose to do a staged release for that one. So, there we sort of released the smallest one and over the course of a year released bigger and bigger ones as we could kind of see how the ecosystem adapted. With GPT-3, I think we have kind of said, I think we are in a world where we do feel like we need to be able to take responsibility for the uses and so we actually have a whole team dedicated to the safety side of this. So every time someone goes live, every time you see a GPT-3 user, that person has gone through our review process that we work pretty closely with all of our developers, and that we do something in some ways analogous to the stage release process with our users and that often will approve them for a certain quota and kind of see how that goes and then scale them up as they figure out their own way of controlling for the safety. I think that the answer is I think that it's not enough as a technologist just to build a technology and toss it into the world and say hey it's the world problem to figure it out. I think that it is also on us to take responsibility for how this is going to impact society. So, I think that an API

actually lets us have kind of the best of all worlds and to really get it out there and get people to be able to use it.

Beena: Greg, I absolutely love what you said. I am often reminded of that remark from that movie *Jurassic Park*, just because your scientists could, they did, and when they didn't pause and think about it, and what you just shared is so powerful. It's not just about technologists to build technology and put it out there and let people choose to use in whichever way. And there will always be bad actors. There will always be the positive things that you can do, but there will be negative things that can be done, and it's so important for us to make sure to put those guardrails in place upfront and make that conscious decision. That for me is so powerful. I know Open AI is set up as a nonprofit. Could you explain why Open AI's status as a nonprofit is so significant?

Greg: So, I think this one is also a very interesting internal conflict that we have had of really trying to figure out how do you build what you think is going to be the most transformative technology and you are really excited to build it, but you also feel like there's a lot of ways that it could go wrong, right? You want to make sure that you build it in the right way. And before Open AI, I was one of the founding employees at Stripe and I was there when we were four people, and I built that to 250 people and it's a very successful company now, and I really believe in the capitalist system. I think that there's a lot of really amazing things that happen by giving people the incentive to build successful companies, to build successful innovations. So I think that there's a lot of power in just saying hey, just be a normal company and do things in the normal way. And I think that the great thing about a nonprofit is that you can focus entirely on mission impact, but the downside is, I think, not that much happens in nonprofits, and I think the problem with for-profits is that when you have this infinite incentive to grow, I think that's where the problems begin, if you're already at scale and you've just got to figure out how do I grow 2X this year, I think that at some point you've got to start sacrificing focus on your mission. And so we really didn't feel like there was a good option available to us, and so we ended up doing something that we designed kind of custom for our problem, but I think actually may be a good pattern for other transformative technologies or other kinds of companies that may grow to be very big if they're successful. And what we have is we have a nonprofit entity that is Open AI that sort of governs everything, and then we have a what we call a capped profit company that is a subsidiary of that. And so the cap profit, it has investors, it has employees who get equity, so you can incentivize people in the traditional way to actually make this thing succeed, but there's a limit, there's a cap, and beyond that point further proceeds don't go back to your shareholders, they go to the nonprofit to benefit the world. I think it's sort of a middle ground that I'm hopeful will both practically mean that we can build this successful system, we can build the system and actually deploy it and have customers and get it out into the world, but it also at some point means that we don't have this infinite incentive to keep growing, which means that if we're already at scale and actually really deeply integrated into society, I think we really can focus entirely on mission impact. And I think that is something that if you just look at the recent history of technology companies, I think if there's kind of one core thing that I think went wrong, I think that incentive to grow is just really overwhelming, and it's not just at the top. It really is every single employee feeling like that's the thing that they have to do, worry about that stock price, that I think that it becomes really hard to focus on anything else.

Beena: This is so well thought out and I agree with you. I think this is potentially how you will see more companies where there is a for good part and then there is a capped for profit. It's such a great model for companies to follow and actually stay true to providing ethical technology in a way which benefits all

of humanity, and I know it's part of Open AI's plan to distribute the benefits of AGI to all of humanity. Could you elaborate a little bit on what this mass distribution might look like, what prompted you to come up with this notion of distribution to all humanity?

Greg: Yeah, I think the mechanics will be, I think, a really interesting question. If we really do succeed and there's a lot of value that we're creating, I think figuring out what's the right way of getting this to everyone is going to be a really important and hard public conversation. But the thing that I think we can see very clearly is like the thing that we actually want to produce is, we want to build an amazing AGI service. We want AGI services that empower every individual, and whether that's the low-cost doctors or tools for doctors, whether it's the low cost, you know, we talked about education as well. I think that if these services being easily accessible, available to everyone, like that's the kind of goal that we're shooting for. I think there's definitely the potential for if it's—I think we're very interested in things like universal basic income and the idea of like if there's literally a bunch of dollars that are accruing to the nonprofit, maybe that's something that's literally given, distributed to the world or maybe it's more about really focusing on delivering the kinds of services I just described, but whatever it is, I think that figuring that out is going to really be integrated with the technology and the actual kinds of services we're able to deliver.

Beena: Yeah, can you share some of the other projects that you're working on? I think GPT-3 is very well known. I read up on DALL-E and I love what you're doing there. Can you maybe start with DALL-E but share some of the other projects that you're working on?

Greg: For sure. So DALL-E I think is also quite interesting. The funny thing about basically all of these projects is that like everything is a neural net. It's totally random, you pour in some data and out comes some good stuff. The theme that we're at now is degenerative neural net. So you are pouring some data and then you're able to get the neural net to generate things like the data you put in, and it's not just memorizing stuff for necessary or even just interpolating between things that's already seen. It really kind of comes up with some pretty novel combinations and creations that you just can't really find in the data set. And so DALL-E is a neural net where we trained it on text image pairs, so you're then able to type in a prompt and out comes an image that represents that prompt. Now of course it's not perfect, but it's quite good and there's a lot of these samples that are like, that we asked it to generate a picture of a baby Daikon radish walking a dog in a tutu and it actually generated a bunch of really awesome pictures like that, and they're certainly much better than I can do. One thing that to me is really exciting about DALL-E, if you step back and think about why do people want to use these generative models at all? I think there's basically three reasons. One is that they can give lower latency answers that a human could. And this is something like AI Dungeon where you can't pay a human to sit over your shoulder and to like be a dungeon master for you as you're coming up with your fantasy game. Two is that maybe they can generate things at bigger scale than humans could. That sounds a lot like web spam or other applications that we really don't want to support. Those are the kinds of things that we don't venture into with our API. And then there's a third thing, which is that maybe the AI can help with an answer or an artifact that a human couldn't, that is actually meaningfully new and different. And I think that DALL-E is the first time where we've produced a technology that has that flavor. There's also AlphaFold is a technology out of Deep Mind that's starting to solve the protein folding problem, and that is again a problem that is just like you're actually kind of cool with the machine going off for a day and thinking hard about a particular problem and coming back with an answer for you. And so I think that that's the real kind of like amazing milestone that we should all be shooting for now is AI that can actually generate novel artifacts or new answers to problems. So that's DALL-E. We have another project that I think kind of has an interesting flavor of this, which is GPT-F. So this is our theorem proving system, and this is basically taking the GPT technology and just applying it to mathematical theorem proving and

there we actually have been working very closely with the formal theorem proving community and have come up with a bunch of proofs that are shorter than the humans have come up with. And so if you actually look there's this GitHub repo where people accumulate these formalized proofs and you can see they're credited to all these different authors and these people who spend a lot of their time trying to figure how to shorten the proofs, and now GPT-F is a credited author in this database as well. And I think that's really exciting, because the thing that neural nets are supposed to be bad at is reasoning. That's the thing that we haven't really figured out. You ask GPT-3 a very logical puzzle, and it's probably pretty hit or miss whether it'll get it right. And I think GPT-F starts to really get at the point of how can we get neural nets, how can we get this technology that is able to really kind of absorb all this common sense. How do we get it to use that common sense for a structured reasoning problem, and we're actually able to make some progress there, and I think that's really, really exciting.

Beena: That's awesome. I live very much in the applied AI space, and one of the challenges that the world that we live in today, within the context of AI, there is a lot of research happening in the AI space, like what Open AI is doing, what's going on in academia, and then in the real world using technology that is still being developed and is not fully matured. What are some of the options or what are some of the ways that the enterprise world can get more connected with the work that you're doing? Or overall, what's the best way for research and applied AI to stay aligned so that for AI to thrive, I mean, we have to figure that path out. There's a lot of one-off partnerships that happen, but what have you seen in your experience that works?

Greg: Well, I think it's a great question and I think that we put in a lot of effort to communicate our results, and if you look at our blog, I think it's become one of the blogs that a lot of people in the field to read. And the funny thing about the blog is that I think it was a pretty big innovation when we started it. The way that people communicated was you put a paper in archive and I think that's great for communicating to the narrow set of people who are directly working on what you do, but I think if you want to really communicate more broadly and to really make these things more accessible to others then you really got to put in additional effort to really communicate the meat of your work. Like one thing I think is really interesting is that my co-founder Ilya will, whenever he reads a paper, he'll always say the same thing, "Oh, it's a really simple idea," and he'll sum up the paper in one sentence. I think that's basically how this field works is that the ideas are not deep and there's nothing wrong with that, but I think that due to a lot of the incentives of the academic cycle and things like that, which we generally try to just—that's not what we focus on, but I think a lot of people do, you sort of try to make your paper seem as scholarly and deep as possible. So I think that the number one is really focusing on that communication and focusing on the right goals for that communication. I think that the goal should be to get this thing into people's hands and to make it accessible and make it so they can understand what's going on.

Beena: Yeah, and that's from the research side, right? What should companies do? How do companies who are using AI and want to stay current? What do companies and the CEOs or the board members need to do to stay more aligned?

Greg: Yeah, well, so I think that staying up to date. First of all, just to step back, I think that kind of what we're going through right now, it's a little bit like thinking about the internet in the 90s. There's something that's coming, it's going to affect all of our businesses, right? I think every business is going to have AI in a real way, like I think that people kind of thought this was happening in the 2010s and it kind of just didn't happen. But I think that this decade, we are going to see a real transition and I think that after that more and more rather than less and less. And so the question is, you think about in the 90s, how could you have evaluated what the internet could do for your business? And how could you be on

top of it? And I think that the answer is, today, there's way more information available out there than there was where you literally can go online and just look at all the different things people are doing. I think that what we find with GPT-3 is that, I think every business has a GPT-3 application within it. Every business is a language business. And it may not be that GPT-3 is ready to actually solve that problem or to make enough progress on it that it's worth deploying, but I think that what we found is that people are basically always able to find some area in their business where we can add value, and AI right now is like a rising tide. It's kind of covering all of the problems at about the same rate and so I think that if you find one area, you're probably going to find ten more where these things can be applied. So I think that maybe the answer is, think about it like there's a sea change that's coming that you want to be probably on the leading edge of that rather than on the trailing edge, and one nice thing that's changing with things like GPT-3 is that you don't need to go out and hire a giant team of specialists, that you can actually use an API that doesn't require you to figure out how to run a bunch of GPUs and how to optimize them to figure out different learning rate schedules or whatever other fanciness might be required, that you can use it like you would any other piece of technology. And so I think that it's going to be different for every business, but I think that the basic ingredient of trying to find the smallest problem that you think you can get value from and try it out and really get your hands dirty. I think that's probably the safest device.

Beena: Yeah, and in addition to that, the technologies like GPT-3 obviously there will be new roles that get created, there will be job descriptions that might change. Do you have a sense of what are some of the new roles, we spoke about the new businesses, but what are some of the new roles or what will, working with the powerful technology like GPT-3, what will that look like?

Greg: Yeah, I actually have a personal story here. So my older brother, he was in the army for six years, he came back home and started to do journalism research and he needed a better tool to help him in that research to go and analyze all the news stories out there, and so he actually tried out GPT-3 for this and he just got so fascinated by it that he just like spent all of his time working with the technology and he actually kind of forgot about the journalism research and just spent all his time trying to get GPT-3 to do new things, and he actually was able to get better performance on a bunch of tasks than we did, and now he actually is doing this full time. He's employed at a company that's built on GPT-3 and that he is doing this all day. And I think to me points to something really interesting. We've got this new technology that's coming. There's going to be new skills required to extract the best performance from it and really understanding what it's like, where its strengths are, how to coax that intelligence that's in there in the machine, like that's going to be something that we just haven't seen before. And so I think there's going to be this really important role for people who are excited about that, who really want to work closely with the technology and figure out what it can do, and I think that that is really cool, because I think most technologies, it's kind of the other way around. We know exactly what it can do, and you've got to figure out where it might fit within the business, and I think that here simply figuring out what intelligence is in the machine is itself a really important endeavor.

Beena: Yes, that's a great story. Thank you for sharing, Greg. There is so much happening in the world of AI research and you are kind of spearheading a big part of it. How can people stay connected with you? Where can they follow you? What's the best way to stay engaged with the work that you're doing?

Greg: So definitely the Open AI blog is a great place. I tweet about things on Twitter as well, and I think that those are the main communication mechanisms that we have that we really try to push everything out through.

Beena: Greg, thanks again for being with us on the show today and I wanted 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.

Automated Voice: 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.