Deloitte.

AI Ignition

Ignite your AI curiosity with Eric Siegel, PhD

Beena Ammanath (Beena): Hi, everyone. Welcome to AI Ignition. My name is Beena Ammanath and I lead our Global Deloitte AI Institute. And today on AI Ignition, we are joined by Eric Siegel. He's a bestselling author and a former Columbia University professor and the founder of Gooder AI.

Eric, you have a fascinating journey. Can you share with my audience a little bit about yourself and your journey so far?

Eric Siegel (Eric): Sure. I've been in the field of machine learning for more than 30 years. I was a professor at Columbia University, where I also got my PhD focused on machine learning. And I've been an independent consultant for 20 years in the field, and now, the co-founder of Gooder AI, as of a year ago. So my trajectory has sort of broadly been like I fell in love with the concept of AI as a kid. It's where I landed, although I'm now really questioning what the definition of that field is. But I did very much fall in love with machine learning about 33 years ago, and I'm still just as excited as ever about that technology. But what I've come to, in more recent years, is a question of, well, is it actually realizing value? We're creating potential value with it by creating models that predict, but are we capturing that value? Well, it turns out that oftentimes these enterprise projects fail to do so and get to that sort of last mile of value capture. So that's been my main focus in recent years.

Beena: Yeah, and you've also been a leading voice in the field of predictive analytics for many years, 33 years. You've written a book on it. Can you start by providing background on predictive analytics, the field itself, and why it has become so crucial in today's world?

Eric: Eric: So predictive analytics is sort of the original enterprise AI. When you're talking about business applications of what people might call AI, it was pretty much predictive for the most part up until a couple of years ago with the unprecedented advent of Generative AI. But predictive analytics is learning from data—that's the core. So you could also call it enterprise machine learning or predictive AI now, to differentiate it from generative. Learning from data to predict, to determine which individual is going to click, buy, lie, or die; commit an act of fraud; turn out to be a bad debtor. Be a good place to drill for oil. Be a satellite that's going to run out of battery or a train wheel that's going to fail—at that level of detail, per individual items or cases—at that level of granularity, in order to drive operations. Because all of our existing large-scale operations are composed of many kind of micro decisions—"micro" relative to the overall enterprise. So determining who's going to click, buy, lie, or die directly informs who to contact, mail, investigate, incarcerate, set up on a date, or medicate. So turning these—this is the most actionable thing you can get from data in terms of enterprise operations. This is the technology you turn to to improve your existing large-scale operations.

Beena: Yeah, yeah. And you know, my exposure to predictive analytics... So I studied AI 30 years ago, a while. I think that's very old, but things have, you know, much of what we did then was more theory, right? It started becoming real in the last 20-odd years. When social media came in and we had access to big data, and my exposure to predictive analytics, that early exposure was on, you know, predicting jet engine failure, predicting how much power would a windmill generate, and so on. But it's so applicable in every part of the business, right? It helps businesses function better and more efficiently.

Now, you've also been in consulting but also in academia. And what are some of the biggest challenges that you see companies face when trying to bridge the technology-business gap in AI deployment?

Eric: Yeah, I mean, that gap is tough. And it really manifests in terms of implementing these types of predictive use cases, these value propositions. And to be clear, if we want to call it predictive AI, to differentiate from generative, we're alluding to two major categories of use cases of value propositions, more than any particular type of machine learning. Not so much the technology, but it's how it's used, right? So in these predictive cases where it's directly informing pretty much all the main large-scale operations that we conduct, the problem is actually getting that change into place.

So the science is good, the data scientists are good, and the machine learning software is good. It develops a model that predicts. That's the cool part, right? That's the rocket science. That's the thing that we fell in love with originally as data scientists, which is the ability to ascertain patterns or formula from some limited—even if it's big, it's relatively limited—number of examples, historical examples, or what have you, or labeled examples, from which to learn, and it derives something that then will hold in general over new previously unencountered cases and situations—and in that sense, has truly learned something; not just discovered a pattern for this data, but that holds in general. Learns something about the world that therefore is very, very potentially useful and valuable, but only do you capture that value if you act on it, if you deploy it, implement it, operationalize it, actually change those operations with the prediction. And by prediction, let's be clear, we're talking about probabilities. We don't have a magic crystal ball. So what is developing, what is generating or calculating, on these per case bases, is a number between 0 and 100 or zero and one. The next best thing to a magic crystal ball. Some kind of a probability saying, well how likely is it for *this* case? And how about for this case? And how about for this case? So that you can decide where do you draw the line, and which of the individual customers who should be contacted that you spend $2 for targeting marketing? The transactions that should be audited or blocked as potentially fraudulent?

So you're driving these decisions based on those probabilities. That's a pretty big fundamental change that involves, yes, the idea of a probability, which in and of itself isn't that complex, but it's such a big paradigm shift that actually getting it deployed, right, actually getting the organization and stakeholders to greenlight at the end of the project that—it's not the end of it once you've actually deployed it, but sort of the culminating point of the initial project. Let's actually get it deployed. And more often than not, it turns out that with new enterprise machine learning projects, it fails to deploy. The stakeholders who have not gotten their hands dirty, in the sense of getting involved with the semi-technical understanding, get cold feet. So you've got to dirty their hands in order to keep their feet warm so that it will get authorized to actually deploy. Right? So the reframing that we need is to consider these business operations "improvement projects," which does not roll trippingly off the tongue the same as AI, doesn't sound nearly as sexy, but in a sense it really is because it's the value! That's where you're realizing value. That's where you're actually improving the way operations work. So currently, the sort of, if I was going to psychoanalyze ourselves, it's like we're more in love with the rocket science than the launch of the rocket.

Beena: Isn't it always true, though, that [laughs] you know, it's always—

Eric: Well, not in the case of rockets! Right? In general, we're more in love with—it's the launch of the rocket that's the really exciting part, right? We need to transfer that same mentality over to these projects.

Beena: Yeah, but you need that some shiny object, shiny to... almost drive that change, move the technology forward.
If you call, like, AI ... Statistics. For a long time, what was used in the real world was statistics. But if it was *called* "statistics," it wasn't as sexy as AI—

Eric: No, statistics is three syllables, and by the time you're halfway down, the second syllable, everyone's already falling asleep.

Beena: Yeah! So, you know, you almost need some kind of a trigger or some kind of reward that helps you move forward and, you know, successful AI deployment is not just about the technology. It's also about the people, the process, the change management that goes with it.

Now, Eric, you've also been up in academia as a professor. How has that shaped your approach to deploying machine learning in the real business world?

Eric: So my first foray into academia was a lifetime ago; really, turn of the century at Columbia University. But more recently, ending about a year-and-a-half ago, I spent a year as a visiting analytics professor at University of Virginia Darden School of Business. So that kind of represents my character arc moving from technical to the business department, where we care about the actual operationalization, the actual creation of value. And it was during that professorship where I really developed the idea, the paradigm, of what I call "Biz ML"—the business practice for running machine learning projects successfully through to the deployment.

So that was the sort of purpose of me going to Darden was to develop the ideas and put it together in a way that not only speaks to the techies, but to the business people. And that's the whole point, right? I mean, the thing that's often missing more than anything, these projects—if you're listening to us right now and you're not a data scientists, you're what we need.

You're what these projects need more than anything, more than better technology, and data scientists with more training. Even, I'll daresay, more than bigger data. We need people who have ramped up on a semi-technical, totally accessible understanding, in terms of what's predicted, how well, what's done about it—and that trio I'll probably try to repeat again—just that basic understanding that defines the project and its value proposition. Ramp up on that basic idea: It's exciting, interesting, pertinent, and not very technical. So that then you can participate end to end, across the steps. For example, as I formalized it with Biz ML, a six-step practice, end to end, across the project, with deep collaboration with the techies, the tech-biz collaboration, the data scientist and their customer, the stakeholder, the person in charge of the operations meant to be improved with the model.

So let's get that collaboration going. That's the message I'm really trying to get out there.

Beena: It's amazing. And you also have an earlier book, *Predictive Analytics*. Tell me a little bit about what prompted you to write that book, and what are some of the takeaways from that book that came out in 2016?

Eric: The purpose of the book was like, look, there's lots of sort of pop science books out there, *Freakonomics* and the Malcolm Gladwell books. Those are fun, right? Why shouldn't there be one also about my field and our field, machine learning? Which should totally be understandable to my uncle, to anybody who's not a data scientist. And it's of interest. It affects everybody every day. So let's put it out there and make it clear what the value proposition, which for any given application or use case is basically what's predicted and what's done about it.

And let's talk about why it's valuable. So sort of two of the main concepts are what I call the "prediction effect" and the "data effect." And the prediction effect is that a little prediction goes a long way. So we don't have incredible magic crystal balls, but we can predict better than guessing, and it turns out that that's generally more than sufficient to render a dramatic bottom-line improvement to large-scale operations.

So it's putting odds on the numbers game that is business. So I call that the prediction effect. And the data effect is sort of the whole point: It's data is always predictive. When you get your … the big data movement was about, "hey, look, we've got a lot of data"—just as a side effect of conducting business as usual and recording all these transactions and happenings and what have you. Data is a recording of the collective experience of an organization. It's a long list of historical events, in a sense.

So it's experience from which it's possible to learn. And it turns out that, yes, you will find these correlations that are predictive, that are the building blocks for these predictive models. You can sleep well at night knowing that at least the tech side of this stuff is going to work. Once you start doing the analysis, you're going to derive these models that really work well.

Beena Ammanath: How much do you think this, you know, from your purview, how much has this field evolved, changed? How much of businesses changed, evolved in the last 10 years, you would say?

Eric: Well, I think tremendously. I think it's much more pervasive that there are predictive analytics projects. What's happening now with generative completely changes the conversation as of a couple years ago. I would say the jury is still out about whether that improves the viability and leveraging of the predictive use cases, which are older but not old school.

Most of its value is generally untapped. And I've been out there saying, look, we should probably—most companies, it depends on the company—should probably be investing and investigating into predictive use cases at least as much as generative, because it may not be as amazing, but it's oftentimes more valuable. It's how you improve existing large-scale operations.

They're really apples and oranges, but generative tends to suck the oxygen out of the room and it is taking away from predictive budgets probably a bit too much, in my opinion.

But more generally, you know, because it's a longer time frame you asked about—sort of a 10-year span—things have changed dramatically because now all companies—even before generative, it's like, well, of course we're going to do something predictive if we're a large enough organization or at least have large enough operations worthy of this type of endeavor.

And so, even though surveys and industry research are both that I was involved in, for example, while at Darden and others, are showing that more of these projects fail to reach that deployment than succeed, even when they were intended for deployment, even though—even if it's 20% that succeed, 20% of a lot of projects is a lot.

The bread and butter of the conference series I've been running since 2009, Machine Learning Week, is about those proven, value-driven use cases across companies and all the different presenters at the conference. There's a lot of success stories out there, even if there are maybe even more failures so far, because we haven't quite made that sort of paradigm adjustment.

Beena: I was going to ask you what prompted you to start the machine learning conference. You alluded a little bit on, you know, how do you make it accessible to everybody and especially the business folks? What prompted you to start these events and conferences?

Eric: Well, the vast majority of AI companies, or at least AI startups, are verticals. So it's like, here's a particular use for a predictive model or for Generative AI, a large language model, or what have you. We're horizontal. So instead of being like, oh, we're dealing with AI challenges, the whole inception of the company is to deal with those challenges.

The challenges we've already been talking about today, but a little more specific. It's about the metrics. It's about determining how good is AI, how good is this model, not just in terms of arcane technical metrics, which are generally the only thing the data scientists deal with during the development of the predictive model, but also really straightforward business metrics like profit and savings. And by turning that corner, which is gravely needed, a fundamental shift that's been outstanding now for decades, we can create the business user console of like, OK, I'm trying to target this fraud detection. Well, how do I use this model? And, depending on how I use it, how much money is it going to save, in terms of more fraud captured and/or fewer, disruptive holds of transactions, or fewer audit staff, people spending time investigating?

So there's all these trade-offs they need. It's just the arithmetic. It's not the rocket science, but it's very particular arithmetic that requires a specialized visualization so that it makes sense. It's been presented in terms of business metrics. And we've made that much-needed move from these technical metrics to provide the visibility so business stakeholders can actually get their hands dirty in that sense.

See something concrete, see it in terms of the straightforward KPIs that are the lingua franca of business. And then it becomes plausible to greenlight the deployment.

Beena: Yeah, you know, with the rise of AI, obviously the topic around ethics of AI has also gained a lot of prominence in the past few years. What is your approach to addressing the ethical considerations of machine learning, when looking at making it more trustworthy?

Eric: So here's the issue. When you're doing these types of projects that are going to drive large-scale operations, you're not just streamlining business and optimizing the number crunching. You're governing, in a sense. Lots of the time, for a lot of these projects the decisions that are human-facing, whether in health care, credit assignment, or even legal applications, are very—they drive consequential decisions. So literally, are you going to get a loan? Are you going to be released early on parole from prison, depending on a prediction of recidivism, which is will you be incarcerated again upon release. So the way these models manifest, the ethical considerations come in, not so much whether they get it wrong, because they will get it wrong. Again, there's no magic crystal ball. And if we had no computers and it were humans making these decisions all the time, they'd also get it wrong.

So by using data, hopefully we're getting it wrong less often and right more often, and demonstrably so in many cases. But there's a secondary issue, which is how often do these costly unfair decisions, where someone's wrongly incarcerated for additional time or wrongly denied a credit card application, or housing, or any of these really consequential decisions. How often are those what are called "false positives" taking place more often for one group of individuals than another? So it comes down to distinguishing by protected groups, like race, national origin, and gender. And looking at the false-positive rate, and that's called machine bias. That term kind of comes from an article of the same name in *ProPublica*, which is the most often cited with regard to that type of issue.

Beena: Eric, thank you so much for joining the show today. And thanks to our audience for tuning in to *AI Ignition.*

Eric: Thank you, Beena. It's been great talking to you.

Beena: Be sure to stay connected with the Deloitte AI Institute for more AI research and insights**.**

Visit the AI Ignition Podcast Episode Library
Deloitte.com/us/ai-ignition