

Barefoot Innovation Podcast: Live from the NBA Conference with Michael Akinwumi, David Castillo, and Cat Tucker

***Note that transcripts may sometimes contain errors and that transcript timing notations do not match the posted podcast**

Announcer: Please welcome to the stage Jo Ann Barefoot for a live podcast recording on AI with David Castillo, Cat Tucker, and Michael Akinwumi.

Jo Ann Barefoot: Thank you. I'm Jo Ann Barefoot. I am CEO and co-founder of AIR, the Alliance for Innovative Regulation. And we are, as we did last year, recording this segment as an episode of our podcast show, which is called Barefoot Innovation, and you can hear it on any of your favorite podcast platforms. We're here today to talk about one of my favorite subjects, AI. It's ubiquitous. It burst into our lives two years ago next month, I think, and I don't think we've ever seen a tech development that hit everybody's short priority list at the same time in the way that this one did, you know? Within a few months, a couple years ago, just about everybody was trying to figure out what does this mean, how does it work.

How many of you have daily or nearly daily use of ChatGPT tools? There we go. So, everybody's learning, right? We're all learning intensively together. And so, we're going to spend some time talking about what this means in general, what it means to you, and maybe explore the notion that these tools, especially generative AI, have the potential to drive down costs and sort of democratize access to technology in ways that may be revolutionary for small financial institutions that now face so many costs and difficulties in staying competitive in today's marketplace.

So, I think the panel has been introduced pretty much. We're really excited to have this group of people, and we're going to start by talking about the upside. And then we're going to talk about the downside. You know, have we ever seen something like this that is so amazingly high potential? It's going to cure diseases and, you know, end certain types of crime, and all these unimaginable great things. And on the other hand, maybe it'll bring the singularity and the end of mankind. So, you know, which is it? You know, we can't tell from where we sit. But we do know a lot about the near-term and how it's going to be different for you. So, Cat, I thought I might start with you. Set the stage by talking about what you see as the upside potential and why. What are the things about AI, including generative AI, that are especially a fit for the challenges that MDIs face?

Cat Tucker: Yeah. So, you said it in your intro. It really is a leveler. And so, I can talk through, when we think about it at Moody's, one of the things that you spend a lot of time on is understanding your customers, understanding data, getting financial data, doing analysis, all of those things. You're able to, and again, underlying is that you have the access to the platforms and the data and the technology to do that. I think we'll talk about that a little later. But there's a lot of overhead that

can be reduced. There's a lot of insights that can be pulled. There's a lot of data that can be accessed. There's insights that, you know, pieces that can be pulled together that otherwise, you don't see their connectivity through humans, that you can do with gen AI.

So, there's the opportunity to speed up a lot of processes on due diligence, on monitoring, on things like that. So, we see a ton of upside in just eliminating some of those mundane or those tasks that are usually very heavily human-driven, and very expensive, and sometimes wrought with error, because a lot of them are manual, through the use of gen AI. Now, I won't go and immediately flip you to the downside of, you know, bias, bad information, you know, and hallucination. But above, if you can get the right model in place and the right partners, it really should reduce costs over time and your ability to do some of those basic tasks to run your business.

Jo Ann Barefoot: Great. And Dave, what would you add?

David Castillo: Yeah. I've been in AI my entire career, and I am so excited about what's happening in the last few years. The ability to put an AI in front of someone who can actually interact with it in natural language is incredibly powerful. That's the user interface. The things that you can do through natural language are endless, and we're starting to see that. Entire ecosystems have been built up in the last few years that are incredibly well-financed, verticals and horizontal. I'm very excited about what's coming out on the other side, and as a practitioner, you know, there's a lot of building blocks that have led up to this, and three building blocks in particular.

In 2017, there was a transformer, something called a transformer architecture. That's what enabled the ability of words to be measured against how they fit in relationship with other words. And then you had to be able to do that with phrases, and you have these multiple dimensions of being able to simplify that. That's where the auto-encoders came about. So, these AIs have been around for a while, but when they were all put together and then you introduce something called a GAN, a generative adversarial network, or a GAN, it is able to create new things from stuff that it learns. And you put all this together and you get the gen AI outcomes. It's incredibly exciting. I'm looking forward actually to the next few years to see what happens, and I'm just so grateful to be a part of it.

Jo Ann Barefoot: And Michael, you work very extensively with these issues, and I know you're going to talk about the downside risk also, but your focus is not only financial services but also housing with the National Fair Housing Alliance. Talk to us about how you're looking at the upside opportunities here.

Michael Akinwumi: Yeah, sure. Thank you. So, when it comes to the upside, I mean, there is a lot of opportunities that come with AI assistants in housing, in financial services. Right? For example, when it comes to risk management. So, and we're thinking about, okay, how do you take these current aggregate models approach to

understanding consumer risk profiles and breaking it down to a level where you really customize the solutions to them so that when it comes to credit scoring, credit wariness, you have more granular data that will help you to understand customer risk profiles, right? So, that's one way, credit risk management.

And another way is, you know, imagine you have a customer and you're really wondering what type of conversations you want to have with them, right? But now, you have a tool that can help you to decide what the next best conversation with that customer should be, because of the data that you have, a lot of data that can be help you to really understand the personality, for example, of that consumer. Right?

And then the next thing is, I read about it with time. There used to be a time when, anytime I went to travel out of the country, I would have to call my bank and say, well, I'm traveling out of the country, please monitor my credit, my credit card and all that. Right? But right now, with AI, right, we can actually understand, based on some data that you have, your location data, right, you don't even have to make that call. So, it really simplifies our process. So, when it comes to housing and financial services, those are some of the opportunities that we see.

Jo Ann Barefoot: Let me ask this. This is an audience that I think is especially attuned to the well-being of your customers and communities. Do any of you think that generative AI in particular has the potential to revolutionize the consumer's ability to run their financial life? To make better choices, to avoid tricky hidden fees, to save their money more easily, to manage their bill paying and budgeting, by having an agent? An AI agent that is going to be able to advise them and can ask a question and get an answer. Is that science fiction or is that on its way?

Cat Tucker: I think it's real today. So, there's one, again, the point you made is that one, you can speak to it. So, the natural language interface of gen AI is the most powerful piece. Like, the most powerful coding language of tomorrow will be English.

Jo Ann Barefoot: Yeah, that's awesome.

Cat Tucker: But when you think about your ability, so today, what you can do, and I'm sure there's a lot of experimentation, especially at JP Morgan Chase, and I spent 18 years there, but when you think about your ability to have a conversation with a document, or have a conversation, so T&Cs and terms and conditions, have a conversation with two documents. So, there's the ability to load documents into certain gen AI models and ask it questions or have them talk to each other in natural language back to you. So, the power and the ability to consume information which otherwise, you'd never read, you'd never get down to the details or couldn't find them, you can ask questions, and they can have a position or a point of view, which is a little scary. But it's possible.

And so, I think managing your personal life, but you can extend that out to anything where you need to understand vast amounts of data and go deep, or even go broad. It's going to give you that ability to have a conversation with that data, which is incredibly powerful.

David Castillo: Yeah. What I'd add to that is that because there's this concept called prompt engineering or prompting, the dialogue that Cat mentioned, that you're having with these documents, the AI is actually, the gen AI actually can start to pick up on the prompt and the style of prompt and actually the structure of the prompt. So, it'll start having a dialogue with you in the context of the communication it's having with you. So, it'll start having a dialogue with you that's familiar, something that you can recognize. It won't be scientific.

I did an experiment just the other day. I was in Las Vegas presenting to a group of accountants, and I asked the gen AI a question. I said, "What is gen AI?" It gave me a very scientific answer about auto-encoders, GANs, and transformer architectures. I said, "What's the definition of gen AI for accountants?" It gave me a definition that was in accountant lingo. I said, "What is the definition of gen AI for 10-year-olds?" It gave me the definition of gen AI for ten-year-olds. But as you start having this dialogue, it's able to adjust the way that it delivers the messaging to you so that it resonates more cleanly with you. And so, it gives you the ability to get smarter about your finances and whatever it is that you're communicating to it about. It's able to make those adjustments.

Michael Akinwumi: I just want to quickly add to Cat's comment. Right? So, again, I know we're going to get to the downside later, right? But as far as, you know, you're using commercially available product, of course, one thing you want to be really intentional about is privacy concern. Right? But with that, given where I think the future is going, anyone right now that actually has apps on their phone, right, you could imagine, say, in a few years to come, all those apps are going to become agents, right? And these are agents that really, you know, I would say bring the computer back to the age divides that you have. So, it's a relief, really. Like, every one of us now, maybe in future, we'll be able to actually leverage those personalized information on our phone or our personal devices without running some of these privacy areas.

Jo Ann Barefoot: Let me remind everyone. If you have questions, put them in the Slido and I can read them here and try to work them in. So, turning to the downside. Let's talk about what we should be worried the most about, and let's keep the focus also on smaller institutions and MDIs, and I'll also just encourage you to talk about whether, as we go through the risks, maybe the answer to risk caused by AI are solutions built through AI. And Michael, I know you've got views on this, so why don't we start with you?

Michael Akinwumi: Oh, thank you. Well, I was hoping I'll go last, but that's fine, too. Yeah. You know, I like to say that the way I think about AI, I actually think of it as a computer-based system that relies on choices made by human to make, for

example, predictions, explanations, interpretations, and decisions that affect real people.

And when we talk about impact on real people, you also want to think about, well, you need data, right? But the moment you think about data, those of have that data science background, part of our training is that data is objective truth, objective representation. But you won't look far in housing, in lending, before you realize that most of the history of discrimination, redlining, you know, even segregation, right? They are present in the data. And it requires some intentionality for us to actually see how we can use AI or algorithms responsibly to take care of those bias that you have in the systems.

You know, one of the things that also comes through the training is that, well, when you develop your systems, you have to think about the performance, right? Or the accuracy. And when we're talking about AI systems, the question is who really defines what should be predicted? Who are the people that you have at the table to decide the type of problem you're trying to solve, or how to even measure the performance, right? Now, if all those are taken care of, the good thing is that, at least in the last 10 years or so, there are solutions that will be used to make sure that those, for example, who are not represented in the data, they actually have the opportunities to also be part of the solution.

I can give you an example of some of the stories that we've done at the National Fair Housing Alliance, in partnership with a FinTech that is based in L.A., Fair Play AI, where we looked at, you know, historical housing data over the last 10 years, and we looked at the parties, looking at some demographics information, we realized that we either use AI to fix some of the biases that we see in the system, or we wave our hand at it. But we know there's a lot of opportunities, and what we did was that we used a technique called distribution matching technique. This is a technique that doesn't only look at accuracy or the performance, which is really the mainstream metric that institutions use to evaluate AI systems. But we built fairness into the technique itself so that, as you're optimizing on the performance of your system that you're developing, you're also optimizing for fairness.

And some of the amazing results we got was that if you're looking at on the right and disparities, when you use this distribution matching technique or the distribution matching technique in particular, you are able to improve those disparities by up to 13% for both African-Americans and also Hispanics. And when we'd also look at pricing, right? So, pricing disparities. We saw that you could improve that by up to 20%. Right? So, again, when it comes to using AI system, addressing or fixing the bias, the math is there. The techniques are there. And it can be used to make sure that you are jointly optimizing for performance and also fairness.

Jo Ann Barefoot:

Dave, why don't we go to you? What would you add on the downside risk top list?

David Castillo: Well, first of all, when you look at gen AIs, gen AIs really don't have the ability to do reason like a human reasons. What they do is they understand how to correlate tokens and where they exist next to some other token, like in language. What does this word mean when it's next to this word? And so, what happens is that if you're training one of these gen AIs with data, and that data has the biases that Michael mentioned, it's going to start to measure the relationships amongst these tokens, and it's going to form these inherent biases. There are bias detection techniques, and Michael mentioned the ones that he uses. There are a few others. And most people, especially the people who are bankers who are involved in making sure that there's fair lending, you know, that there isn't discrimination in the types of banking products that are being put out there, they implement those types of bias services, because they have to. There's controls around. Most banks will actually insist that there's controls around it.

But the problem will continue to exist as long as we're building and training these models on such huge corpuses of data, and all of us know that the internet isn't biased, correct? And so, what people tend to do to counteract some of that bias is to introduce their own data. And actually, I heard Cat talking about this in the room earlier, so I'm going to let her talk about some of the bringing your own data in something called a RAG, or a retrievable augmented generation. And that's an ability to not only stop, to counter the bias, but also to reduce the what we call hallucinations.

Jo Ann Barefoot: Cat, what would you add?

Cat Tucker: So, I think everyone hit kind of the downsides. The other, I mean, security, of course, of your own data and what environments you use it in is one. Bias is another. And then making sure that the results you get are actually accurate. So, one of the things RAG is a means to. You know, LLMs are, you know, it's everything in the internet. So, like you said, it's inherently biased. We can't change history of everything we've written and everything we've said and everything we've consumed. So, that bias is in that LLM. So, it's a matter of marrying, you know, data that you trust, calculations or perspective that you trust, which is through RAG, with that model that you use. So, again, but that's a very expensive thing to do for a lot of larger organizations. So, a lot of larger organizations can do that fairly easily.

Some of the things that we're doing, not a shameless plug but just explaining, but some of the things that we're doing is we're embedding that with some of our services and trying to offer those through our core program to MDIs and CDFIs, so that they have access to RAG models as well, and leveraging LLMs. But when you think about, so another thing is also learning. We use non-learning models, so you don't want it to learn from your prompt or from the responses, in some cases. Some of that is, you know, confidential and proprietary information. So, there's a lot of different ways that you can, you know, counter the bias.

So, another way, when you think about marrying the right data. So, we are a data company at our core, data and technology. And you know, in the ratings business, we understand financials, macroeconomics, environmental capabilities. But if you can marry that LLM with data that you know is secure so you have that traceability, so that when you search something up, you can see the reference point. That's how you're going to be able to deploy that in your business and trust it. So, there's a couple different ways of non-learning models, protecting your data from escaping, marrying it with data and logic that you understand to make sure that it's using the power of the LLM against something that you trust, and then making sure that the data.

So, one of the, you know, we've got a pretty public partnership with Microsoft, and one with Google also, where we're grounding information. So, when you search, and you pull something back, to be able to trust that you trust the source, is called grounding, is grounding that data. So, that's another way. As you leverage it, understand where the data that it relies on is grounded. So, again, all those cons, but there are counters. None of it's perfect, but there are those kind of countermeasures to make sure you can get the power of that tool, but still trust how you deploy it within your business.

Jo Ann Barefoot: So, we have less than five minutes left. Let's bring this down to the practical level. We've got a question in the Slido where, if you're a small bank, where should you start? Also questions on whether these tools can be useful for fraud. The answer to which is, I would say, yes, for sure. It's causing fraud to get worse, but also creating new opportunities. But what's the practical advice? What should the people in this room, you know, go home and do when they sit back down at their desk? You want to start, Dave?

David Castillo: Yeah. I think that everyone can start getting smart on gen AI, and they should. Anyone can pick up chat. Anyone can start to interact with chat. And just getting yourself familiar. Now, in terms of bringing it in to an MDI, there's a lot of low-hanging fruit. Like for example, you can take some of these. You can actually build a bot at a very low cost that just deals with maybe these are the hours that you're available, just some frequently asked questions that the consumer can interact with. Very simple things that don't hurt your brand, don't hurt, necessarily expose you to certain risks. You can get started on those fairly inexpensively and fairly quickly. And then you can progress to other things, like for example, sentiment analysis.

Jo Ann Barefoot: Is that a matter of like, downloading an app? Or where do you start with that?

David Castillo: You can download an app. You can sign up for the services on any one of the cloud providers. There's so many different ways that you can get access to these tools. The other thing I was going to say is you can start with sentiment analysis, about what your people, I'm sorry, your consumers, are saying about you through social media. And then if you get to a point where you actually have a little bit of trust and you can protect your PII, you can actually start to get

sentiment analysis around your emails. That one's a little more tricky, because you have to have the right controls on that. But certainly, there's a lot of non-invasive ways that you can get started. And that's dipping the toe in the water. That's not necessarily going all in. I think at some point, you want to go all in, but you should be educated when you go all in.

Jo Ann Barefoot: Cat.

Cat Tucker: You hit a lot of them. I would say, and I'm just double-clicking on this one, is like really know it yourself. Like, the architecture of work is changing with gen AI. It is going to change everything that everyone does, and I look at it in very good ways. I know what it frees me up from doing. Writing emails and things like that. But I would say, understand it yourself so you can think about how it's going to change your job and what you do every day, and free you up to do higher tasks. When it comes into your business, I think those are really just great solutions, and thinking about customer service is an easy way. Customer sentiment, employee sentiment. But it's going to allow you, just any way that you can take a lot of data and just try to glean out themes or sentiment. It's going to be a really good use of that.

Now, I think practical use and tools that you use in like, your workflow and business, because of the high cost and hurdle to those things, it is going to be seeking partners that, and gen AI is also not this thing that's sitting over here anymore. Like, when we first started talking about it, it was on the, you know, just on the side. It needs to be embedded in almost everything that you do. So, looking at companies that you can partner with that are also embedding that in their solutions, so that it facilitates your existing workflow, or tools that you use today, and how they're weaving that through. But start to do that investigation, because that's also going to speed you up, and it's kind of the unseen way, that you don't have to like, buy AI. It should start to be embedded in some of the capabilities you use today.

Jo Ann Barefoot: Michael? Got one minute.

Michael Akinwumi: Yeah. So, yes to everything that Cat and David said. But at the enterprise or business or institutional level, what I would like to add is that I think most of the issues that we're talking about when it comes to like, AI risk, right, is really because we have low representation when it comes to the demographics that we care about. And I think here, we have MDIs we serve, you know, minorities, that's really the primary market that we focus on, right?

I think there is an opportunity here to actually pool resources together, and you talked about partnership, right? Really identify these tech companies that are actually using AI for good, and you know, you can make your decisions. Maybe you want to bring them in to build assistants that will address your market, or you also want to like, up chase, from some of these commercial entities, right? Or maybe you just want to have like some contracting agreements with them,

right? So, that's one way that I think, in terms of like practical solutions or steps that can be taken at an institutional level.

Jo Ann Barefoot:

If you're looking for a book on this subject, I'm reading *The Coming Wave* by Mustafa Suleiman, the founder of Deep Mind, and it is really, really thought-provoking. It's incredible. My advice that I would add is this is a race against time. This is happening so fast. We're not wired to deal with change that's coming this fast. So, we all need to put a really high priority on it. Please join me in thanking our panel.