

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

Well, I'm super excited to have an amazing thought leader in the field of cognitive AI, human-centered AI, Dr. Michelle Zhou. By way of introduction, Dr. Michelle Zhou is a co-founder and CEO of Juji, the maker of the world's only accessible, cognitive artificial intelligence assistants, which ultimately enabled the automation of human engagement tasks, empathetically and responsibly, all with no-code.

Prior to starting Juji, Michelle led and managed the research and development of cutting-edge interactive intelligent technologies and solutions at IBM Research Watson, including IBM RealHunter and Watson Personality Insights. Michelle's work has resulted in a dozen widely-used products or solutions, over 100 scientific publications, and 45 issued patents, all in the interdisciplinary field of human-centered AI that intersects AI and human-computer interaction.

Currently, Michelle also serves as the editor in chief for ACM Transactions on Interactive Intelligent Systems. She received a PhD in computer science from Columbia University, and is an ACM Distinguished Scientist. Dr. Michelle Zhou, welcome to ByteCast.

Dr. Michelle Zh...: Thank you, Bruke. Thank you for having me.

Bruke Kifle: Yeah. Well, we're super excited to have you. I'd love to start off by asking, who is Dr. Michelle Zhou? Can you maybe tell us a bit more about your upbringing, your education, your career? I'd love if you could highlight some inflection points over the course of your life that have led you to where you are today and what you do.

Dr. Michelle Zh...: Thank you for asking. I'd love to. By race, I'm a human being. More specifically, an adult female. By training, I'm a computer scientist. I got my PhD in computer science from Columbia University. I have always been working in an area called human-centered AI. It's an interdisciplinary area that intersects artificial intelligence and human-computer interaction. By looking back in my career, there are really perhaps about five inflection points or milestones that led me to where I am today or what I'm doing today.

The first inflection point, that would be, first of all, computer science. It wasn't my first choice to study when I applied for college. I was born and grew up in China. Actually, Southwest China. Both my parents are physicians. In China, you must take a college entrance exam to get into college. My first choice was to actually study biology. Maybe because of my parents' influences.

I wanted to become a biologist to give humans magic powers, like making us fly or making us see through the walls. This is perhaps really for my love of science fiction, again, my parents' profession as well. But unfortunately, I wasn't accepted by the biology department I applied for, so I had to make a last-minute switch. I just randomly, literally, chose computer science, which I literally knew nothing about. I would say that would be my first inflection point, which pushed me to computer science by accident.

My second inflection point is after I graduated from college in China, I really wanted to study my graduate degrees in the US. So I came to the US. I would say two professors ... I went to Michigan State actually for my master's degree. Two professors there really helped me find the direction I wanted to go, which is still my current direction of interest. One is Professor [inaudible 00:04:35]. He actually gave me the opportunity to work on visualization, graphical user interfaces for power management. Because of the project, I really just love it. I said, "Oh my God, I can really create those different types of interfaces that enables humans to better interact with the systems."

A second professor was my AI professor, Professor Carl Page, who taught me literally artificial intelligence and also allowed me to do two AI projects with him. And because of these projects, again I said, oh I have to do AI, but many people perhaps actually do not know the late Dr. Carl Page is the father of Larry Page, one of the co-founders of Google. Actually, of course, I didn't know then, because Google didn't happen then. That's definitely my second inflection point. Because of the two professors, I really wanted to do a PhD study that would be in the area that intersects actually artificial intelligence and computer graphics.

That's how I found my PhD thesis advisor, Professor Steve Feiner, at Columbia University, because that's exactly his area. Since then, I have been working in this human centered AI area by marrying artificial intelligence and human-computer interaction.

The third inflection point I would say is, prior to starting Juji, I worked on several human AI-centered systems at IBM and of course, Columbia University. One of the systems really made me to do more in this area and especially what we're doing at Juji. This system's called System U. Later on it became IBM Watson Personality Insights service. What this system did was, basically you can use the analytics, the right algorithm, to automatically analyze the person's communication text to infer this person's personality traits. Think about it as if you have Twitter tweets.

Then actually that's exactly the demo we did. You can then use the System U to ingest the tweets with the Twitter data, and automatically infer this person's personality profile. And because of this work, I found that hard. It just really opens up tremendous opportunities for computers to gain a deep understanding of each individual. Not just about their behavior, it's about their unique characteristics, for example, how open-minded they are, how thoughtful

they are, and how they handle life's challenges. That's one of the reasons I actually left IBM is, I wanted to really further the research in this area, hopefully with more freedom. Because if I'm doing a startup outside, I have a fewer constraints and more freedom to do more research because there's so many challenges still to be addressed. I would say that's my really inflection point number three.

Inflection point number four and five is really what I learned in Juji during the past five and a half years at Juji. Because we wanted to really build the smart computer systems to better understand each individual, we found out that very few people, in fact, have the data, have enough data, sufficient data, to be analyzed, to be used in a very intelligent way. Then we decided to create conversational systems. Once we create the conversational systems, then, I would say inflection point number five, nobody would want to use it if they have to painstakingly customize their assistants. If they have to train their assistant from scratch, if they have to put in every intent they can think of that this AI assistant needs to understand, no way for them to use it.

So we have to make it a really no-code and reusable AI to actually promote, to encourage, the adoption of our AI assistants, so that you can see that's how we get to this point. It is why I'm so passionate about no-code, reusable AI, and especially also the cognitive AI part of that, in terms of really enabling AI to know each user and each person in a much deeper way.

Bruke Kifle: Wow. That's such wonderful stuff to hear and I think your career and track record certainly speaks for itself. I think you also did justice highlighting the role of mentors and advisors and guiding you towards the field that now you have become such a pioneer in. I think at the end of the day, you made a good choice by choosing computer science. It's all for the best.

Dr. Michelle Zh...: Thank you.

Bruke Kifle: You touched on the topic of, no-code AI is sort of the last inflection point, which is sort of the key to democratizing AI, making sure it's easily accessible to all those who maybe don't have the necessary technical background or programming capabilities. I think broadly there's been a longstanding movement around no-code platforms for app development or website development. How do you describe no-code reusable AI and how do you see it actually bridging the AI divide?

Dr. Michelle Zh...: Okay, thank you for asking that, Bruke. This is one of my really favorite topics. Let me actually first break down what is AI. Even though people may have different definitions, I would say one of the common definitions of AI is, a machine with certain human skills. For example, with the human perceptual skills, to see, to understand images and videos, or having a human's natural language processing skills to be able to interpret sentiment, interpret meanings in natural language text. Because we want to teach machines human skills, it's

really not something trivial. It requires AI expertise. It requires sophisticated software engineering skills, not to mention the large amounts of training data or intensive computation resources required.

As you can see, not every organization, certainly not individuals, can afford to have all these required elements, which means that in order to enable reusable AI, first of all, we have to acquire AI and acquire AI is non-trivial. What the reusable AI means is, I always use analogy, everybody is very familiar with advances SpaceX has made. They have to make the rocket first. Making a rocket is not easy. Similarly, making an AI is not easy. Making the rocket reusable for the next trip, that's even harder. Similarly, making AI is already hard enough. You want to make AI reusable, which means that it is literally about transferring intelligence from a one AI to another AI. It's like when you have teaching, let's say you have taught a child certain skills, but you want this child to be actually able to use these skills in a totally different context, a different environment. That's called a transferable AI. More, you want to bring up another kid. You just want to completely transfer the first kid's intelligence to the second kid. You don't want to retrain the second kid from scratch. That's about reusable AI.

As you can see, if you can enable reusable AI, it saves a tremendous amount of effort and expertise required to create AI in the first place. And more, because of reusable AI, the time to value has been tremendously increased. Now, let's do a no-code, because no-code is really built upon the usable AI. Why people need to do a no-code AI in this case is, I have a really special meaning. The meaning is, because each AI requires a little bit of customization. For example, to be able to speak the language in that domain, to be able to communicate with users for a particular task, you always require some of the customization.

No-code means, how can you customize this AI without writing a line of the code? Which means that you can directly reuse the intelligence when we customize the intelligence. That means about the no-code. No-code AI is not about the building AI from scratch. That would not be actually the purpose of it. The purpose is... So, as you can see, in this world I remember there is a statistic showing that it's still a very small fraction of the population know how to code.

Even for the AI customization, you want the people who have domain knowledge, but without programming skills, to be able to customize the AI in a way they want it. This actually opens up the opportunity for AI to learn better. Think about it. Because IT people may not necessarily have all the knowledge. Let's say you wanted to create an AI assistant to help recruiting, or maybe create an AI assistant for healthcare. You want the domain experts, the subject matter experts, to infuse the knowledge into the AI, not the IT people. The no-code AI and reusable AI really opens up the door for AI to be adopted, first faster, and also for AI to be improved much more rapidly, which lowers the threshold, the barrier, to entry, to enter into the AI field, to adopt AI, to use AI.

Bruke Kifle: I see. I think you summarized it well, especially on the reusable AI side. We're seeing a growth of foundational models that are being trained on large corpuses and ultimately being fine tuned for different specific downstream tasks. I think we've seen a lot of performance gains, but it's also minimizing some of the environmental impact that comes with training large scale models, because if you're able to apply them to many fine tuned downstream scenarios, then hopefully that's able to reduce the environmental impact.

On the last point, you mentioned the idea of using no-code AI to enable users to ultimately be able to customize the AI for scenarios that they're particularly looking to use. Now, critics might argue that while you might take away the code away from programming with no-code platforms, you don't necessarily take away the core logic behind the algorithm design. This idea of conditionals, loops, and oftentimes beyond just writing the code itself, solving and architecting a solution for a given problem, performing the test, deploying it, is really where a majority of the challenge lies as well. How do no-code platforms help address this problem? As we segue into Juji, how do you think about this as a challenge?

Dr. Michelle Zh...: Thank you, Bruke, for asking this great question. Actually, you are asking this question because you have knowledge about computer programming. You know about IT, you're working for Microsoft. And from our clients, most of the people who are subject matter experts don't have IT background. They don't care about algorithm designs. They don't care about conditions loops. They don't even know they exist. What they care about the most is, what this tool helps me to do. How can I use this tool to achieve my goals? That's what Juji is really actually working hard toward is, how can we explain the AI's function in the way the domain experts, the subject matter experts, can understand, can customize it, without getting into the nitty gritty weight of understanding underlying algorithm, condition loops, or maybe even recurring functions. Because first of all, they don't care, second, it doesn't mean anything to them.

I just actually last week gave a talk about the challenges and opportunities of enabling no-code AI, reusable AI, for subject matter experts. You can also call it... New York Times called it, "for the masses." There are three challenges, which are very different challenges than programmers have been facing. First one, AI design. Because AI is powerful, but yet it's that not that powerful. How could you teach people to understand AI's power, but in the meantime, realizing AI's limitations. That's a huge one for us, because if you don't teach people that, and people can actually completely rely on AI, and it turned out it won't work. But then if they realize the AI limitations, they then refuse to use AI. Let me just give an example, which is very... I observe them, our clients are doing this one, but once they get more information on the one, they completely change how they design.

Some clients have some experience working with conversational AI. They knew that they have this perception that AI doesn't work very well. In their

conversation, they always use, I call it the button. They will tell you, would you like to continue, versus let the user say what they want. They will give you two buttons, yes or no. Because they didn't know how powerful the AI is. Another one is, they really want to ask the open-ended questions. For example, "What kind of medical conditions did you have when you were a child?" But you want to elicit really open-ended. But because they didn't know what AI can do, they basically then limited the answer to say, condition one, condition two, then it turned out maybe it doesn't cover all the possible conditions.

This is a one-way extreme, which means that they don't really trust AI very much. They wanted to restrict human-AI interactions. Another extreme, people haven't had a lot of experience working with AI, and they would completely trust AI. They will put in very open-ended questions like, "So what is the biggest challenge in your life right now?" And if they don't know, if the AI has a lot of limitations, they're not prepared, I call it, if they do not anticipate. And the users may come in to say, "Why do you want to know?" That's very personal. Users may respond with, "What's your biggest challenge?" That's why, when we actually give our tutorial, we always say it as, when you do conversational AI design, try to fill in the gap.

G means understands your goals. A means anticipate. P means personalize it. Know your users, personalize it. This case actually, we found that people gradually tended to trust AI. That's one of the challenges that you can see. This is very different than teaching people how to program. It is still programming, but teaching them about the limitations, the scope of AI, if you will. So that's one.

Another part of it is AI supervision. When you adopt the AI, it's almost like you adopt a child, adopt a junior assistant. If you tend to ignore them, abandon them, and definitely your users who abandon your AI assistants, because the knowledge is not being updated. It's almost like if you ignore your child, ignore your assistant, they're not going to learn new things. They're not going to advance their knowledge and skills. That's why we are also trying to actually inform our clients that when you adopt AI, be prepared that you are taking on a responsibility actually, seriously. You want to keep it updated and you want to really improve it over time.

And you can see, it's very different. In some ways there are some similarities, almost like you're programming. AI supervision's almost like you have to monitor your program, debug it. But in the world of no-code, people may not understand what do you mean debugging me. But if you tell them it's about supervising your junior assistant, supervising your intern, supervising your child, they do understand, and then they will do that.

Bruke Kifle: So no-code doesn't necessarily mean no responsibility. Exactly, you're still responsible for maintaining the AI. But I think a big part that you mentioned is the importance of educating folks on the scope of AI, the capabilities, the

limitations. Because I think oftentimes there's a misunderstanding or a misconception of what can or cannot be achieved by AI.

Dr. Michelle Zh...: Correct. May I also, what I call it is, add one point to this one. We were talking about the programmers, developers, and the machine relationship, is really what I call it an operator-machine relationship, because they program to tell computers exactly what computers should do. But now with the no-code AI, it really transforms the relationship between human and computer. It's from the operator-machine relationship to what I call supervisor-assistant relationship.

You will teach your assistant to do certain things, but you don't need to be very nitty gritty to the very detail because they don't need it. They already have a certain level of intelligence already. That's why we call them AI. That's why it's also the biggest difference between the programmers learning some machine language, machine instructions, versus subject or matter experts are learning how to teach AI using no-code. It's a totally different level of abstraction, different levels of learning.

Bruke Kifle: Certainly. Focusing on the areas that matter most to the application area. ACM ByteCast is available on Apple Podcasts, Google Podcasts, Podbean, Spotify, Stitcher, and TuneIn. If you're enjoying this episode, please subscribe and leave us review on your favorite platform.

I'd love to learn a bit more about... I know we've sort of established some of the foundations of no-code AI. But I'd love to learn a little bit more about what you're doing at Juji with cognitive assistants and Juji as a platform for your clients and your users. To start off, I would love to learn a bit more about what led you from a career of research to now a field of entrepreneurship as the co-founder and CEO, and what exactly is it that you do at Juji?

Dr. Michelle Zh...: First of all, as I said earlier, and because at IBM, my co-founder and I... Actually, my co-founder was also a very key contributor to the projector called the System U. Later on, known as the Watson Personality Insights. Because of that project, we realized that it's such a big space. There's so many things we could do, so many challenges we could solve. We decided to be entrepreneurs, trying to basically get more freedom to do what we believe would impact the world. This actually leads to today's Juji. At Juji, we have our mission, we call it, unifying machine and human intelligence to advance humanity. I'm a very big believer about this, what do we call it, not we call it, actually J.C.R. Licklider, your MIT professor, called it a human computer symbiosis. So we believe computers will always be humans' assistants.

Not going to replace human, instead, augment humans. At Juji, what we do is we create... We have actually developed the new generation of AI assistants. We call them cognitive device assistants. What does that mean? That means that those AI assistants have certain advanced human soft skills. For example, one of the such soft skills, we call it active listening. Actually, we published an

article about this in ACM conference, the CHI conference. Which means that, during the conversation... In the human-human conversation, in order for the conversation to be more effective, people, we're taught to actively listen to your partners. To repeat it, to show your empathy, to show your understanding. That's what we also have taught our AI assistant to do. Second, when you talk to psychologists, psychologists are not only to understand what you are saying, it's also trying to figure out what you are like. What you are not saying.

So we also have taught our AI, we call it a soft skill, because reading between the lines, which means it is by dynamically analyzing a person's conversational text, and trying to figure out what this person is like, what is this person's unique characteristics. And then you can use that insight to better help this person. This is naturally going into the applications of the cognitive AI assistants, which we found the sweet spot is any type of a human engagement, especially long term continuous engagements that are required. Such engagements are often emotionally charged and also require quite a bit of social interactions, and more, in such interactions, individuality matters. As you can see, the thing about it is, quite a few use cases in this area, healthcare. Thinking about it. When somebody's recovering from an injury, recovering from a disease, it's always long-term, continuous engagement. It's always emotionally charged and the social interaction's always desired.

Another one is their individuality, their personality matters to the effectiveness, to the outcomes of the conversation. Because, can they stay on track with their treatment, whether they drop out. If their AI can help them to stay on track, hey, that's a winner. Similarly, in the learning, think about the students who take an online program. Normally, typically two to three years, sometimes two to five years, long-term engagement. It's also emotionally charged. They have to overcome a lot of challenges. Again, individuality matters because different students have different needs. Different students have different learning styles.

You can already see this kind of, what I call it, the sweet spots for cognitive AI. Workplace companionship is another one. I just kind of pointed out a few of them, which we found, we have discovered as the sweet spot, if I say the killer apps for cognitive AI assistance, as we have created. Of course, we enable this no-code reusable, which means that it's subject matter experts, like the healthcare providers, the learning coordinators, or the HR professionals to customize the AI assistants on their own and feed them the domain knowledge they have, which maybe IT people may not have, most likely won't have.

Bruke Kifle:

I see. You talked a lot about equipping AI with important capabilities, for instance, like active listening. But I'm thinking from a user point of view, there's also a change in norms in how you communicate with a human versus how you communicate with a chat bot. At least based off my personal experiences, whenever I am on a website, of course it's not an advanced cognitive AI assistant, but the many AI personal assistants that I've interacted with, you think, Alexa, you think Siri, there's some mode of interaction. It's very

transactional. You issue some query or some question or some ask, and you get a response.

So what's fundamentally different in the social norms or the rules that govern how humans interact with other humans, compared to how traditionally humans have been interacting with machines or chat bots. Do you see that changing, and in the meanwhile, how does that influence the design of chat bots currently today, or with platforms like Juji?

Dr. Michelle Zh...:

Again, a great question, Bruke. Thanks for asking that. There is a book by a psychology professor at Stanford University, already actually talked about, when humans interact with computers, they tend to follow the very similar social norms as how they interact with human beings. We actually use the same principles to guide the design of our cognitive AI assistants. For example, you mentioned just moment ago, the existing commercial AI assistants, like Siri, like Amazon Alexa, or Google Homes, they are very transactional. They are impersonal, I would say, because they don't really care about who you are.

What we have done fundamentally different is, to make the really two week conversations. As I said, active listening, that's coming from a human to human communication theory. Reading between the lines, again, it's from a human to human conversation, especially from a computational psychology, from a psychology point of view. About the psychology of the human to human engagement point of view.

That's how we use those principles from the human to human conversations to guide the design of our human-AI conversations. I think that will drive, not just us, for other companies, for other designers, to do the same. Because people actually use the same social norms to interact with the machines. But we did find two things, which is very interesting. One, we always actually educate our clients to make your AI humble, very humble. Because in this case, people tend to be more forgiving. Remember, AI, it's not perfect, it's far from perfect.

Second, being transparent. Tell your users what it knows, what it doesn't know, again, is to gain that sympathy, to gain that forgiveness from your users. That's a very important one. Again, you see why we use this principle. It's very similar in human to human conversations. So if you talk to somebody who is humble, who has the humility, who is very transparent, you are much more willing to open up. You're much probably more willing to talk to that person. It's very similar in the human to AI conversation as well. Everybody always likes to talk to the people who care about you. We call it, who can think in your shoes. That's similar. We train the AI to do the same. If the AI can understand what your unspoken needs and wants, of course it can help you more.

Bruke Kifle:

Where do you think Juji is sort of on this journey of achieving the emotional, empathetic cognitive AI? Is it something that's here? Is it something that's in the

near future? And more broadly, what are some of the biggest challenges that you're seeing in the space, or opportunities for improvement?

Dr. Michelle Zh...: Actually, it's here. It's today. Our clients have been using it and they have seen the outcomes and effects. What's a challenge here is, as I kind of alluded to already is, explainable AI, maybe it's a practice. How can we help the domain experts? Again, they're not programmers, they're not IT people, they don't have a lot of IT background, to discover the magic of AI and to best to use the magic of AI. For example, to let them know what the power this AI has, and also along with the limitations as well. Really, I think this work really elevates explainable AI to a next level, not just for the data analysts, not for the people who are training, doing machine learning. This is for the masses. How can you explain to them, what is the magic of AI? What kind of magic your AI has, let's say it that way. And how should they use that magic in their application, in their solution. And in the meantime, to be aware of the AI limitations, AI imperfection. So that's one of one.

Second part, which is maybe the topic you would like to discuss as well, responsible AI. Currently our AI already can gain a deep understanding of a person during a conversation. This one can be used for malicious reasons as well. So if you knew that this person really likes to play a game. And it could be very easy to addict to games. And then you make a really bad AI to seduce, or maybe to allure, these people to just play a game every day, every hour. That's a part I'm worried about seriously. Because of this democratizing of the AI, what we have done, which literally means anybody can come in, create a very powerful AI that can understand people's strengths and weaknesses. Then, how do you prevent that from happening? Are there any principles, any measures can we apply? This AI is about responsibility and AI ethics. This is in our community. We have talked about now, especially now with the power of AI all the time.

Bruke Kifle: Certainly, yeah. I think we've seen without a doubt, many case studies or examples over the past decade where we've seen some of the negative consequences of AI. Even in the cognitive assistants space, we've seen cases where AI has continued to engage with users and result in model drift. I think it's really great to hear that at least some importance is being put around responsible AI principles as you're looking to democratize AI capabilities for the masses. So I appreciate you addressing that before I even got to the question. It's great to see that it's a top area of priority.

I want to focus the last segment of this talk around future directions. If you're familiar, one of the things that was recently announced by GitHub is Copilot, which is the AI tool pair programmer. Now, from my point of view, copilot and no-code are fundamentally different things. However, they do have some commonalities. They're focused on boosting development, they're focused on democratizing computing. What are your thoughts on the future of software,

now that we have the introduction of tools like Copilot, and how do you see this reconciling with the no-code movement?

Dr. Michelle Zh...: I haven't actually used Copilot myself. I did watch some of the demos. Actually, I saw it's a very cool tool. Because it's a tool just like our no-code reusable AI design studio is another tool. It really depends on the audience, depends on the users of the tools. I would say, tools like Copilot will advance as well, but their main users would be developers and programmers. And the no-code AI, like what we are developing, their main audience would be the people who don't know programming, who are, as I said, subject matter experts or domain experts. So I would say both of the tools would be needed. Actually, we're even right now contemplating the tool it is when we're talking about no-code AI. How about you use the conversation to actually design a conversational AI?

That's from our platform, that's even right now, can be totally supported. Which means that you don't even need the GUI anymore. You can just say, "Hey, what kind of AI assistant would you like to build? Let me walk you through, let me help you create one." And basically using the conversation to design a conversational assistant. So you can see it's all full of possibilities, but depending on who are the users of the tools. I think all these different types of tools are needed, because there'll always be programmers, there'll always be developers. They would need tools like Copilot. Other people would need the tools like ours.

Bruke Kifle: Yeah, I think that's a very great distinction to make. The end user or the target audience is different between those two tools. That's awesome. The next thing I want to actually pick your brain on is future directions for Juji, specifically. For me, I am quite familiar with scratch. I'm not sure if you're familiar with it. It's a free website or interactive coding platform that was developed by the Lifelong Kindergarten group at MIT Media Lab a couple years back, which has been a very powerful tool for introducing computing at an early age to kids all around the world. So I'm curious, with platforms now like Juji, do you see applications in early education? And more broadly, do you see potential for no-code AI to transform adoption and exposure to early education around AI?

Dr. Michelle Zh...: Thank you, Bruke, again, for this awesome question. Actually, there are high schools and universities that have already begun to use Juji as a platform to teach AI. What's a great aspect of that one is, for example, San José State University, their business school, to teach AI. They wanted to teach non STEM students, the business school students, about the core concept of AI. They use Juji as a platform. I also see high schools in Cambodia, and they use Juji to teach their students about the core concept of AI as well. We actually support them in this effort.

This means that it's really opened up the space for computer science education. Traditionally, thinking about it, non STEM students who cannot program, they wouldn't know what AI is about. But now, the tools like Juji really change that. So people actually writing their resume, I have used the AI tools, I created this AI

and I accomplished this task. It's like they write down their skills like using Microsoft PowerPoint or Excel. Now, they say I can use the tool like Juji to create conversational AI agents, AI assistants. Actually, running startup is really hard, but those kind of use cases really make me smile at night.

Bruke Kifle: I'm sure. I'm sure. I think like you said, early education, early exposure, can really be transformational for opening up a whole new world of knowledge, of experience, of career opportunities. So I think hearing about some of the applications of Juji and transforming early education is certainly very exciting. One thing I want to touch on, I'm sure you've seen the recent headlines around the Google AI bot becoming sentient. As someone working in this cognitive AI assistant space, and having done research in the space for quite some time, what are your thoughts on if and when, if at all, we will ever achieve sentient AI.

Dr. Michelle Zh...: I guess I would take a step back to ask the question, "Why do we want to make a sentient AI?" So it really depends on purpose of the AI's use case and AI applications. How great would it be? You have an AI which has no emotion, seriously. Which means that this AI would not have any emotional burdens as we humans have. But in the meantime, to exhibit empathy, exhibit empathetic behavior, would that be great? Why would we require AI to have emotions, to feel anything at all, as long as they can exhibit empathetic behavior? So in this case, what I would say is, because we humans create AI, we wanted to really best leverage AI's strengths and avoid AI's weaknesses.

Similarly, we want to best leverage human strengths and avoid human weaknesses, which means that when we make AI, we don't want to teach AI something we don't want AI to have, because you know how hard it is about this. For example, I read somewhere about the call center workers, especially in the 911 call centers. They are under so much emotional stress because of the calls, because this, like EMT workers, the same thing. That's because they're humans, they have feelings. How great it is that we have AI that can do all of that and it can be empathetic still to show that behavior, but without having the emotional burdens.

Bruke Kifle: I think you posed a really good question. Why would we ever want AI to-

Dr. Michelle Zh...: Be sentient. How can it benefit us? I guess I'm maybe a human centric person. I always thought about how would AI help us humans? How AI could help us advance humanity. If we're standing from this point of view, then we'd know what we want to give to AI, what we don't want to give to AI, for the sake of the world. We are the makers. The makers should make the decisions on what you want make.

Bruke Kifle: Well said. To end off, I guess I would like to leave broadly with an open question. As someone who has been working in research for quite some time and now is pioneering such a great field with democratizing AI and cognitive AI assistance, what are some of the future directions, both opportunities and

challenges, that you see in the cognitive assistant or AI space more broadly?
What excites you?

Dr. Michelle Zh...: What excites me? Actually, what excites me also scares me. It's the same thing. It's really this power of AI. This democratization of the power of AI gives humans tremendous power, tremendous augmentation. To do things which we couldn't do easily in the past. But in the meantime, it's like a cliché, but it's real. With this great power, comes with great responsibility. And how should we handle that? And how should we actually better handle our responsibility of having this very powerful tool in our hands? That's why I say, I'm very excited to see, wow, this is amazing, we can do almost magical. But in the meantime, it also means that it could do great harm as well. And how can we prevent that from happening before it happens?

Bruke Kifle: Mm-hmm. With great power comes great responsibility. I think with folks like you pioneering the field, I think we are in good hands.

Dr. Michelle Zh...: Thank you.

Bruke Kifle: Thank you so much Dr. Michelle Zhou, for taking time to speak with us at ACM ByteCast. Very much enjoyed the conversation.

Dr. Michelle Zh...: Thank you, Bruke.

Bruke Kifle: ACM ByteCast is a production of the Association for Computing Machinery's Practitioner Board. To learn more about ACM and its activities, visit acm.org. For more information about this, and other episodes, please visit our website at learning.acm.org/bytecast. That's learning.acm.org/bytecast.