AI and You

Transcript

2022 Predictions

Episode 80

First Aired: Monday, December 27, 2021

Hello, and welcome to episode 80! Today is an extra-special and extra-long episode, because it's our last episode of 2021, and so this show is all about predictions for AI for 2022. But instead of having to listen just to me, I've brought in a panel of talented experts. You will be hearing from **Katie King**, who is the author of the 2019 book *Using Artificial Intelligence in Marketing: How to Harness AI and Maintain the Competitive Edge*, and the 2022 book, *AI Strategy for Sales and Marketing: Connecting Marketing, Sales and Customer Experience.* She has over 30 years of consulting experience and has advised many of the world's leading brands and business leaders, as well as numerous schools, universities and colleges, and keynoted in many countries. **Prashant Natarajan** focuses on the intersection of business, technology, and human experiences. He is VP of Strategy & Products at H2O.ai and has served in leadership roles at Deloitte Consulting, Oracle, and Siemens. He is the author of the "Demystifying…" business book series on data, analytics, AI/ML, and digital transformation, and is a Co-Faculty Instructor at Stanford University. He was introduced to me by fellow panelist **Richard Foster-Fletcher**, who is the Founder and Executive Chair of MKAI, the inclusive Artificial Intelligence Community, and leads the MKAI Centre for Digital Trust. His first book will be published in 2023 about the economics of digital trust. Richard has advised the United Nations Environmental Programme (UNEP) and the United Nations Framework Convention on Climate Change (UNFCCC). And **Ben Goertzel** is a cross-disciplinary scientist, entrepreneur and author of the book *Ten Years To the Singularity If We Really Really Try*. As Chief Scientist of robotics firm Hanson Robotics, he led the software team behind the Sophia robot, and he leads the SingularityNET Foundation and the AGI Society, which runs the annual Artificial General Intelligence conference.

You'll be hearing first from Prashant, then Richard, then Katie. Prashant had to leave halfway through for a prior commitment, but shortly after that Ben joined us, coming from a prior commitment. Here we go.

Peter: Welcome, everyone to our AI 2022 panel and our distinguished guests here. And I just rated myself in the last episode on my predictions for 2021. I got five out of ten. So I said, "I got to do better, I need help." And here we have help. Why don't we start with retrospective? If you put yourself back in your body in 2020, a year ago, and you're looking forwards to the next year with that futurist hat on, and thinking about what would come in that year, can you recall what you were thinking about a year ago would be coming down the road? And rate yourself. How did the last year work out for your powers of prediction? Anyone?

Prashant: I'll start. One of my predictions for 2021 was that by end of the year, I'm going to be joining a panel of really smart people to discuss our predictions for 2022.

Peter: Thank you, Prashant.

Prashant: And lo and behold, I'm on your show here. So I think I'm 100% on that one so far.

Peter: Ten out of 10.

Richard: I'm definitely zero on three for the three that I had. So I remember thinking, "This is the year when Zoom is just going to transform completely from where we start the year to where we end the year." And to me, it looks the same. So I don't think that happened. And then I thought augmented reality was really going to start taking off this year. I thought we were going to get some good glasses, and we didn't. And then finally, I worked on a project with an academic institution where I said we were going to buy up some virtual real estate in Decentraland to get the kids playing around with it. Well, that's fine, but that would cost me about $3.5 million to do that now. So I called that one wrong as well.

Peter: Well, thank you. Thank you, Richard. What about the metaverse announcement, though? Even if it's not hardware, doesn't that count?

Richard: It's incredibly exciting. And I'd love to hear what my far esteemed panelists think about these kinds of things. Bill Gates says we'll be taking most of our meetings in the metaverse within 24 months. I know we're playing around with it. You can probably see my Oculus Quest behind me. And I think what's interesting for business is that if we have Zoom meetings like this, we can't position them in time and space. And if we had this meeting at your office, Katie, there would be visuals, there would be a scent that would remind us of certain chemicals and cleaning products or whatever that will place our meeting there. But we don't have that. We have every meeting in the same chair in the same room looking at the same screen. So I think that's really exciting for business, that we can go and have meetings in the forest and in the air and these kinds of places just to begin with.

So you're saying we need AR and VR that includes smells?

We need to place ourselves in space and time for sure, I think, and VR definitely lets you do that. The new Apple headset is about half the weight of the Oculus behind me, so that's interesting for usability. So dare I say it, maybe 2022 will be the year when we go more into AR and VR. But then I might have to apologize again in 12 months.

Peter: Okay, the first one is written down, carved in stone. Thank you, Richard. Katie?

I'm not big on making predictions. I don't even remember making many. I was too sucked into writing my books. But I know, for example, in some of the roads in the UK that we've had some trials of self-driving vehicles, so I didn't make many predictions. But I certainly thought that they would take off more in the past year than they actually did. But at the same time, I thought AI would be slower to take off. And I've been pleasantly surprised, particularly in the last three or four months at how much interest and applications I'm actually seeing, particularly in the marketing space. So something I thought would be faster and something I thought would be slower, have kind of surprised me pretty dramatically. So yeah, the acceleration that we've seen because of the digitization with COVID has brought everything forward. Luckily, vaccine

development is the biggie. But some of the other applications I think have just sped up, and I think that caught all of us unawares.

Peter: What are those sectors that you've been seeing this movement in? And what do you think they will do in the coming year?

Yeah, no, that's a really good question. I know I'm particularly focused in my work on the applications across sales and marketing and CX, but that is across all industry sectors. And we're certainly seeing it in retail and in banking, and professional services. So we're seeing subsets of AI, like machine learning being used, not at full scale, but pretty dramatically for fraud detection in the banks. So if I think of a banking organization, not even a FinTech, but a traditional bank, and I split it into sort of three areas, the front end is using AI for some of the interaction with the customers, that might be through a chatbot assistant. In the middle, various teams are using AI software, for example, for getting closer to the customer for personalization. And then in the back office, we've got fraud detection, you know, the machine learning, breaking it down and finding anomalies, and so on. So I think we've really seen more of that happening in some of the banks and financial services, insurance, some in law, and professional services. But quite a bit of it is operational, and quite a bit of it is from a marketing and sales point of view, as well. And lots more tools are being used too, which I certainly explore in my next book that's coming out in a couple of months.

Peter: We have plenty of opportunity for other plugs. So I want to explore with everyone, now, the question of what you were getting into there, which is I think, broadly, the commoditization of AI. Go back far enough in time, the only people that could do it were the R&D departments of major tech companies, and the only intelligence that could do it were in the departments of major universities. But now it's being commoditized. You see things like SageMaker from Amazon and so forth, that are trying to push this down to turnkey solutions where anyone can download it, push a button and run it. It's not obviously that far yet. But how do you see this movement happening? I picture it as a lowering of a bar. It's seeping down into more and more layers pursuing this adoption curve in society. Where do you see this now? And where is it going?

Prashant: So I think I'll take that one because I've been kind of playing in the space from both perspectives, both from the perspective of being a user of multiple types of these AI/ML technologies, not just applications, but also multiple platforms. And now I'm working for a company which again, I think has the best - obviously, I'm biased - soup to nuts AI/ML platform out there. So having said that, I think the key thing is not just the ability to democratize the access to technology to go beyond data scientists in order to empower operational users and business users. And, dare I say it, in the ways AI is being used to impact us, not just with the products that we know contain AI, but also the products and the decisions that contain AI, which we are being exposed to, as citizens, as community members, not just as creators, and designers and developers. When you look at that, I think the tradeoff has to be between, on one hand, how accessible you make it, and that still needs to continue, because the more

accessible something becomes in terms of who's able to use it, the better it is for all of us; a rising tide literally raises all ships. And it's not just the technical ships, but also the business ships and the operational ships. And like Katie, yes, I do have a book coming out too on a very similar topic. So I think, as a result of talking to several executives, globally, I think we are definitely seeing that shift. COVID has accelerated things, not to obviously restate the obvious. But the ability to go, shall we say, expand the usage, should also be supported by the increased ability to improve, for example, not just improve accuracy, but also improve the interpretability, the explainability, create conditions for trust, fairness, explainability, in ways that it comes together. So the simplest solution is to create a black box, which is where I think a lot of the cloud services are today. Without pointing the obvious, the API data science services, which just give us the data, "Trust us and we'll give you the output." No, that's not good enough. If you want to bring it to more people, part of the tradeoff also has to be, to be able to get into the depths of how that decision was made and not just "trust us," that's not going to connect.

Peter: Right. And so you're opening up these issues of transparency and explainability, and bias. And so I want to turn to Richard here because of the tremendous amount of work you've done on inclusivity in AI. And if we go back to the beginning of the development of AI, or any technology, it's localized, its de facto monochromatic, because there's only a few people doing it. Now, the adoption of AI by so many people and its effects on everyone forces these issues that you've been exploring of how do we make that happen fairly for people, and this is what Prashant was talking about there. How do you see that that has evolved and will evolve? A huge question, right? That's most of your business. But if you could condense it for us into five well-chosen minutes.

Richard: I'll do my best. It's a wonderful question. I think just to loop back around and catch you up very quickly, I feel that maybe just there's two bars and one is lowering, as you said, and I think that's the one, the business-as-usual machine learning bar is coming down. And that's incredible. And my colleagues here do amazing work to make sure that happens. But I worry about another bar that sort of drifted away. And obviously, AI is a funny old term, isn't it, that we're never quite sure what it means. But if we take it at face value, then the highest end of AI is probably beyond inaccessible now, and that there is a few companies that have monopolized GPUs and can do things. And yes, you can get a little bit of access to them, but I worry about the chasm that's opened up. You will never do AI like Google does, or Amazon does, or Tencent does, and so on and so forth. So to come back to your question, yeah, major, major problems with biases in data that have been well documented, and no easy answers on that. I think it's becoming quite regional and quite interesting. I saw some noises from the US, so they might be more interested in the EU's work on high-risk AI and what to do about it, which is exciting. China can get a bad rap but have a look at the ethical papers that China has put out over the last month, they're actually better than anything in the US and the EU, would you believe it. So have a look at those before you pass judgment. You still may wish to pass judgment but read that first. And then finally, from my perspective, the UK is a very interesting place to look, because they've very much put out a call to arms last week and said quite clearly, "We are

going to build an AI assurance industry in the UK." That's what they're calling it, AI ethics, whatever you like to call it. And they're saying it's going to be bigger than cybersecurity, and they want to create this with the various actors, parties. And I see Katie's nodding. And I just had a call with some of the representatives today and I said to them, "I think there's more that can be done on diverse inputs and wider perspectives." And they said, "Great." And we've booked another call for January to talk about how we might do that. So very excited.

Peter: Richard, which group is that?

Richard: The CDEI.

Peter: Which is?

Richard: The Center for Digital Ethics and Innovation. You caught me on the hop, but I think that's right.

Peter: Thank you. I want to go back to something Prashant said where you said COVID has accelerated all of this. And I want to get a handle on that because, on the one hand, I think it must have, on the other hand, I don't see headlines saying "Deep learning creates new vaccines." I don't see headlines saying "Autonomous vehicles have put the tracking part of the supply chain crisis to bed."

Prashant: I don't think we should. I think there's definitely been a relatively large amount of hype. I'm surprised that the AI toaster that not only toasts my bread, but before that, it slices it for me, and I heard that it even eats and digests it for me, so I actually don't need to exercise any more. I understand that that's the latest rage right now but that's the level to which I think AI washing and the hype has gotten, where it's far easier to sell AI by attaching those terms rather than prove AI. And what you're seeing, like any other technology curve, there is a reckoning and a reality that comes in. And with that reckoning and reality comes the question that I think all four of us are always trying to answer, which is how do we create value? How do we measure value? How do we communicate value? And how do we make it a repeatable process in order to make it a part of our culture? And I think that's the key thing over here. And I think the real test is going to be in measuring quantitatively and less so but also, in certain cases, qualitatively, the impact that's going to come out and then being able to create conversations where not just the users of the solutions, but the people who can be various stakeholders, even from an observational perspective, have a view into it. That equalization to your earlier point, or I think Richard made this point, you have certain giants, who are basically using AI in ways, and I think that equalization is a need of the art. And equalization is not merely standards, not merely legislation, but also improving access and ability for others to be able to play at that same level, if that makes sense, Peter.

Peter: Thank you. And I want to come to Katie here as we're talking about the direction that these things are going in. Is it fair to say that marketing is, like, a leading indicator, just because people have to market something before it goes on the market?

Yeah, I think so, Peter. I think we've had the digital world now, we've had social platforms, I think we can now go direct through those platforms to customers. I think over the last five to eight years, we've seen more of an emphasis on customer experience, and getting closer to your customer. And I think that's played into the hands of the venture capitalists who've funded the AI tools for marketing and sales. Because what people are looking for, to Prashant's point there, is value. And if you can give better customer insights, if you can have less trial and error, taking away that guesswork that you get in marketing, if you can be much smarter in your targeting, be more efficient, have more sort of human-like communications, effectively take away that term "artificial" and make it *augmented* intelligence, you give your salespeople or you give your marketing professionals, proper, real data insights. And when they've got those, they can deliver to the customer or the client what they need. And I think that's the difference. And I think the breakthroughs have been cool tools and they might be Concured or Phrasee that are used by Carphone Warehouse or by Virgin or even by some smaller law firms, for example. And if they can get those detailed insights that mean they can sell more, or do things more cost-effectively, I think the hype is gone, and more of these tools are now doing a good job. And I think that's where we've seen some breakthroughs.

Prashant: And if I may, just as parting thoughts, to Katie's point-- And again, it always makes me feel smart when I have to agree with the preceding smart person and reflect on their brilliance. It's like a satellite effect. All I'll say is Katie said something that is, I think, extremely impactful. It's what is artificial about artificial intelligence? The data that we use, it's a data-intensive thing. Data is created by humans or by systems that humans have created. It's really not too much artificial about it. Second thing is it's augmentative and it's also a system. We tend to focus too much on the automation aspect of it, which also creates this challenge around lost jobs, etc. But the point is, decision making becoming more of AI as complementary and supplementary to human needs, human wants, and like any other technology in the last few millennia that we have used, humans have always used technology to harness it to drive human needs. This AI is going to be no different. It's not going to be the toaster that's going to digest your bread. It's not going to be *The Terminator* that's going to do, pardon me, colonoscopies against my wish. It's going to be things that's going to help all of us. And I'll say that I couldn't agree more with my panelists. And my apologies that I have to drop off from this conversation just as it was getting going.

Peter: No need to apologize, Prashant. Thank you for joining us. That was a great summary. We've still got half the episode to go. So we'll try and step into your shoes there. Thank you for joining us.

Prashant: Thank you so much. Bye.

Peter: So I want to continue with Katie then. From a marketing perspective then, what pressure has the pandemic put on AI?

Good question. I don't think it's put pressure; I think the digitization has made people bring forward their decision making, for many people by a few months, for others by a year or two. So if anything, it's alleviated the pressure. I think it's played into the hands of the-- whether it's AWS, or whether it's an independent, VC-backed provider of these tools, in a marketing context, we have the compute power now, and we have the data. So actually, the last couple of years, it's been simmering, and then I think the pandemic has helped it explode. And again, not pulling the wool over people's eyes, that doesn't mean that everybody's using this, it doesn't mean that all industries are using it at scale. But I think it's made it acceptable. And Richard talked a little while ago about the regulations starting to come together and things starting to be taken a bit more seriously and getting a bit more global consent around the need for some of this. Again, I think that's helped with the credibility and giving people a bit of a way forward to say, "Yes, we've had GDPR, now we're starting to see tighter regulation." We need that at a trade body level for it to really be pushed hard. But I think if anything, the pandemic has sort of twisted that knob on that pressure cooker and made things a little bit easier and starting to sort of brim over a little bit.

Peter: And since you brought up regulation, Richard, thinking about issues like the environment. AI gets a lot of press for consuming huge amounts of electricity in machine learning training. What sort of conversations have you been involved in? What do you see happening with AI and the environment as we go forward? It seems like this is something that it's about time that something happened in this space. What do you think?

Richard: Yeah, thank you for the opportunity to talk about this. There's a couple of different angles, and I'll try and go through them quite quickly so that I don't take up too much time on this one topic. There is a lot of bad press about the consumption of large, particularly the language models and GPT-3, and so on, the large carbon footprint that they have. I think Katie might know numbers better than me, but people talk about AI that uses the same consumption as Holland and things like this. So those are kind of scary when you hear those headlines. And of course, we're not the only ones. The cryptocurrency market also gets its fair share of bashing. But when I speak to experts in this space, they're pretty clear that we're making great progress in this area. I mean, if you look at crypto as an easy example, you know, Ethereum 1 to Ethereum 2 uses 98% less power. Now, if you widen that lens out, and I know I've spoken about this before on various platforms, if we did chance upon civilizations out there in the galaxies, which we won't, but if we did, then we could pretty much work out quickly how advanced they were with one metric, how much power they were using. So I think it's safe to say that in 100 years, 1000 years from now, I think we know we're going to be using a lot more power. So I don't think the answer is to say, "Don't work on that AI model because we've got to reduce our power." If you're talking about CO_2, well, I'd say I'd start with concrete and cement, if that's what you're worried about. That's 11% of all CO_2 right there when you tear down old buildings and put new ones up in their place and call them more efficient. So before we even start picking on football teams flying all over the world and F1 and whatever. Now, these are all minor carbon footprints, but my point is, you can argue any way you want on these things.

Peter: Well, I think I wanted to look at how AI might repair or rehabilitate that reputation. For instance, Google used AI to reduce power consumption in their data centers by 40% just by analyzing the usages with deep learning. And now I'm reading about neuromorphic architectures that consume orders of magnitude less power. So there's perhaps the potential for a movement, maybe, where AI declares itself to be carbon neutral; I don't know, like there are carbon credits for AI that you can run AI engines and if they save as much CO2 as they consume, then you can put a little green check mark on it and advertise that. I don't know, I just made that up. Any thoughts?

Richard: I think if you want to see a real change in this, of course, you must have carbon tax. Otherwise, you're just asking companies to care, and they don't. But of course, things do, of course, get more efficient over time. What will quantum computing do to this? But certainly, in the longer-term future, when you look at how we might have sort of living cells involved in computation, DNA, perhaps, or bacteria, way out of my wheelhouse, but certainly, you could imagine that things could be very, very different in terms of efficiencies. There are some great examples right now, Peter, nVidia are of course working on how they recalibrate the resolution of calls like this so that you don't have to send all of the data packet down, you can actually calculate on the other side what it would have been and then recreate it, rather than sending all of the data. So that's interesting. One thing I just wanted to put in, and you can throw it straight back out again if you want, is I spoke to the UNFCCC after they returned back from COP26. One of the questions I asked was, "When you hear these headlines, 'China wants to be net zero by 2060, and this country wants to be whatever by 2030,'" I said, "How much do these leaders put this number in and just go, 'Guys, we don't need to worry about it because some piece of tech will come along between then and now and save us? So if we just pick a long enough timeline and cross our fingers, then we'll achieve it.'" And there was a bit of quiet like, "Oh, whoops, possibly some are thinking that." I'll hand back to the floor.

Peter: Katie?

I don't have much to add there, really, apart from Richard's last point, which is actually training the AI to come up with the solutions to exactly those kinds of complex problems. I think that's what we have to do is to be thinking strategically and to be using the technology and feeding in the right kind of data inputs to try and come up with the answers to some of those problems, be they operational, be they strategic, and so on. So I think that's where it's making sure that the AI is being used in a way that benefits mankind, benefits the planet, as well as benefiting UK PLC or China PLC.

Peter: And two statistics that, I think, play into this. One is the often-cited Moore's law, that computational power for constant dollars doubles every 18 to 24 months. And the other one that I've heard, and I don't know how current it is, but it was current about a year ago, was that AI engines are doubling in size every three months. It's not hard to see that that means the power consumption must be going up because it would have to be every 18 months on the AI engines to stay constant. So that's something that's heading for a collision. And if you want to

go out on a limb here and speculate about what sort of innovation revolution might handle that, why don't you do that? But first, we have our next guest, and Ben is here. We'll pause. Hey, Ben.

Ben: Hey, welcome, welcome.

Katie: Hi, Ben.

Peter: Fantastic. You join us *in medias res*. If you wanted to speculate about what sort of technology might come, *deus ex machina*, to save that, be it neuromorphic computing or something else, do you have any thoughts about that? If you're ready, Ben, do you want to take a crack at that?

Ben: Yeah, I think that's really a complex and multifaceted issue. Because I think that the power consumption and general just resource utilization of the deep neural net models that are sort of the trendy thing in the AI field now, and among big tech companies, in particular, this, I think, is largely unnecessary, and is a result of the limitations of this particular sort of AI technique. And indeed, shifting to neuro-symbolic AI, or just more sophisticated forms of recurrent neural networks, I think there are many variations on AI technology that could be emphasized more, and that would decrease the amount of resources used for the same amount or more intelligence. I would say the situation now is that big tech companies have a negative motivation to pursue such ideas, because having AI techniques that require all the world's data and require a huge amount of processing power, means they have a differential advantage. Whereas if you have AI techniques that require less data and less processing, then anyone can do it, and the Googles and Microsofts of the world have less differential advantage. So I think a big part of the situation we have now is the incentive structure that Big Tech has, actually biases them to the most data-intensive and resource-intensive AIs so that they're the only ones who can deliver them. But what I would say, though, is even once that is overcome, and we move to smarter AI techniques that can do more with less resources, even after that's done, then AI will still grow and take up more and more resources because that's just the trajectory that we're on. I mean, just as the human brain uses quite a lot of energy per unit mass. Cognitive processing is just an energy-intensive pursuit. And if by giving more energy to an AI, that will figure out a better way to solve other problems-- I mean, if that's an option, on the whole, that's probably more efficient than doing less thinking and solving other problems in a less efficient way. So there's a number of different aspects here.

Peter: Right. And I'm reminded that I can run my artificial general intelligence on a burrito for a day, which is quite a gap from our hardware. We were talking earlier a lot about AI democratization and you raise a very important point, which is, I think that if effective or leading state of the art natural language processes can't run until they've ingested a terabyte of the internet, and are running on very expensive hardware, then is that counterdemocratizing? Is that the way that the big companies capture this space and crowd everyone else out? Have

we reached the point in AI, in other words, where the little players will not be able to compete? Richard or Katie?

Ben: Yeah, that's certainly a threat. And I think, however, that there are also pathways to a technological workaround to the phenomenon you note. And just as Bitcoin is decentralized money, we can decentralize AI processing, which is what we're doing with SingularityNET. You can decentralize provision of processing power, which is what our SingularityNET spin-off, NuNet does, NuNet.io. You can decentralize management of big data, which is what Ocean Protocol, another blockchain project does. So I think that there's a number of small projects out there. I mean, small relative to Google and Amazon; these can have dozens or hundreds of employees. But there's a number of relatively small projects out there that have, I would say, already demonstrated early stages of solutions that can decentralize everything that big tech companies now do. But I would say that, at the current stage, these technologies are more awkward to use than the centralized technologies, and they don't yet have the same economies of scale that the centralized technologies do. So there still needs to be some maturation of the decentralized infrastructure. But on a computer science level, this stuff doesn't have to be centralized. Again, it comes down to economics and the industry's choices as to where to allocate resources.

Peter: And so then from an economics perspective, and maybe, Richard, you have a thought about this, is there concern among smaller players, ones that are from some disadvantaged population, perhaps, that they are not going to get to play in this sandbox because--?

Richard: All I would add to this conversation is I think Ben is right and therefore, these companies will live or die by this centralized strategy, which is interesting. The world will change, of course, a lot of sideways forces. Yeah, a lot of work yet to be done in decentralized technologies. They are exciting, as Ben describes. A lot of work to do. But there's sideways forces such as climate. And we've got a decade of some really, really intense fighting going on. We haven't seen climate terrorism yet; we will. And I guess when I look at somebody like a big fast-food manufacturer and retailer, I would say, "Don't worry about whether your profits are going up or down by 2% next year, worry about the day when nobody comes to work." The tipping points, I think are things that they can't see - these corporations - one, because they're just over the horizon, and two, they don't want to see them.

Peter: And Katie, so to the question of is there a crowding-out happening?

Katie: Yeah, I think there potentially is, and there's a fine balance, isn't there, with innovation and regulation. And I think we need trade bodies, we need governments, we need the VCs, we need the scientists, the businesspeople to kind of have equal voice, and to have that, for example, codes of conduct, treaties, all kinds of things in place to ensure that there is a bit more of a level playing field. And that might be the beginning of a new form of economics, it might be, in the shorter term, some simple codes of conduct, like I say, and elements that force certain areas of taxation or mandates or stipulations of some sort. So I think that's been the bit

that's missing. The tech companies have, by blinding people with science, been able to be above the law, and have been able to be above regulation. And I think, as happened with social media, and as is now happening with this area of technology, the law, and the regulators and others are gradually catching up. And I'm involved with the All-Party Parliamentary Group and involved with the AI and Ethics Journal; that piece of it, and Richard talked to it five, 10 minutes ago about big changes. It's all been happening in recent months. We might eventually see an equivalent of Geneva treaties and all kinds of others. But at the moment, it's burgeoning, and I think it will lead to a shift in the way that the smaller players can operate and more pressure on some of the big technology players. I'm optimistic that that will happen.

Peter: You know, we've been talking about the developments within the AI industry, we've been looking at regulation oversight, and I think a third leg of this is public perception. And just as an example of how that has changed, now, you can go into a lot of mainstream conversations and find that they understand the value alignment problem, that they've heard of it, or the control problem, yes, they know about that, where it wasn't on their radar, had no place in their minds 10 years ago. And so I'm wondering how you think that public attitudes towards AI, in general, perhaps, even this existential question or something more immediate, are shifting and may continue to shift in the year to come? Who wants to take that?

Katie: Yes, good question. Well, I think we've been ingesting AI, haven't we, in small doses in our everyday lives for quite a few years. But I think, post-COVID, the public is beginning to understand a lot more and to accept a lot more. I look back to when I wrote, you know, my first book came out in 2019 but I was talking to people about it in '17 and '18. Nobody had a clue, and people were very scared. And I think now they are visiting their GP virtually, their technology is more in the lives of everyday people, even my mum in her 70s. And so I think people's perceptions are greater, their understanding, they've seen the impact in the medical world that AI is had on COVID, speeding up vaccine development. So I just think we're seeing the beginnings of an adoption journey that probably would have been a lot longer had COVID not been around.

Peter: Ben, what are you seeing with audiences, people you talk to in terms of shifts in awareness of, and attitudes towards, AI?

Ben: I think that there certainly have been shifts in attitudes. I think there's still very large differences, demographically, based on age and based on geographic regions. So I lived in Hong Kong for 10 years and moved back to the US last year. I'd say in Asia, attitude toward AI and robotics is overwhelmingly positive and the basic assumption is that they're there to help people and will make everyone's life easier. In the US where I am now, I find that young people, say Generation Z and younger tend to have that default attitude, by and large, they're like, "I don't mind if AIs take my job, I didn't want to have to get a job in the first place. Great. Bring on the universal basic income." I think there's a segment of older Westerners, which is the category I'm in, but I'm a strange guy, who are more worried about existential risks of AI and worried about AI taking everybody's jobs and so forth. And things are being polarized there as

with many other things where I think some people are learning to stop worrying and love the AI, and others are getting more and more worried, like, “Can’t someone do something to stop the AI from taking my job and leaving me bankrupt in the street?” And I think that all these aspects will get more and more intensive as we progress from narrow AI to really having AGI systems out there. Once you have the first AI that palpably comes across as having human-level intelligence, even if it’s not fully as smart as a human, I think then all these attitudes are going to amp up to the extreme, and you’ll see quite polarized attitudes on whether that’s a good or a bad thing, which we’re probably not well prepared to deal with.

Peter: I think we’re going to creep up on that gradually, as opposed to the Hollywood vision of someone invents AGI in their lab overnight. I think it’s going to come a tiny little step at a time.

Ben: I don’t know that I see it that way. I think if you look at the recent evolution of AI, there’s maybe a three-year period where computer vision transitioned from being a laboratory thing to being commercially omnipresent, say 2014-17 Starting with that AlexNet. And then NLP, starting with Google’s launch of the BERT model, it was pretty much bang, bang, bang, one NLP result after another. You have GPT-3, and a whole bunch of neuro models for NLP which are now rolling out all over the place. And so I think the pattern of the advent of AGI may be similar. So yeah, maybe not overnight just the AGI crawls out of someone’s basement and takes over the world. But it could be a two or three-year period in which you see new results in AGI coming out every couple of weeks like we saw with vision and NLP during certain periods of time. And on the one hand, it’s slow enough that you can meaningfully adapt; it’s years, not minutes. On the other hand, it’s fast enough to unsettle and confuse a lot of people and not give them time to retrain themselves and reorganize their lives.

Peter: Richard, your perspective on how people’s attitudes in general, people outside of our bubble, our echo chamber, are perceiving and especially feeling about AI. How’s that evolving?

Richard: Thank you. Yeah. And the opportunity just to catch up will be great as well. I wanted to echo what Katie said, I think we are getting much progress and I think we’re raising awareness for sure. We’re dealing with a very confused public, as been mentioned from news headlines and sci-fi, which have not helped. And even just calling it AI didn’t help, did it? So we’re having to push back and all that kind of thing. I think that I really struggle with this AGI concept. Maybe because it sounds like a line we cross; but Ben, you’re saying it’s very much not a line. But it’s certainly divided as such, isn’t it, into narrow and then maybe broader and then general, but of course, it’s sophistication is going up over time. But this line problem suggests that once it crosses it, it’s out of the bag and it’s all thinking and planning and--

Ben: Well, I mean, was the computer revolution or the Industrial Revolution, were those lines that we crossed? I mean, if you zoom out enough, there are lines that we crossed, if you zoom in month by month, they were gradual, right?

Richard: Yeah, but I think people are worried about AI having its own intentions, which is just utter nonsense.

Peter: Right.

Ben: I don't think that's utter nonsense, I think to say that's utter nonsense is utter nonsense. I think at the moment, indeed, our AIs don't have their own intentions. And similarly, before manned flight was a thing, people didn't fly. You could say that was other nonsense. So the advent of manned flight was also gradual, right? I mean, you had balloons, you had the Wright brothers, you had jets, gradually, you have spaceflight. And so getting to new things that haven't been done before, it does happen by steps, right? And you're likely to have some sort of AGI toddler before you have an AGI Einstein or engineer. So I think it will be step by step, not all of a sudden, this superhuman AI pops out of the robot lab and takes over the world. On the other hand, I guess I'm enough of a physicalist that I don't believe that human beings have intention and free will given in a unique way by God that we can never replicate in a machine. So I think that what we call "will" and "awareness" are associated with the physical processes in our brains. And I think the same attributes can be associated with the physical processes in machines. And I think we're not there yet but we're in that direction.

Peter: I think we're just talking about a difference of timescale here. And we're looking in this panel, especially [at] what can we say about 2022? And Ben, you're going at a longer scale.

Ben: Yeah, if you go back to Ray Kurzweil, who I don't agree with on everything, but I think he's been a decent futurist. I mean, he's predicting 2029 as the year when human-level AGI will achieve, which is not 2022, but also not 2200, right?

Peter: Exactly. But now, to bring it back to something that's immediate, lethal autonomous weapons, for all of the *Terminator* hype, which is obviously nowhere on the horizon, is an immediate concern. You've got the sequel to *Slaughterbots*, the video, that's just been released; there are organizations actively campaigning for some kind of treaty, some kind of regulation. And where might this go in the immediate future? Katie, do you want to chime in?

Katie: That's interesting. I think any of that level of scaremongering and exaggeration isn't going to help. We've already had if we didn't get the Gartner Hype Cycle, you know, we already had that peek of paranoia. And I think these kinds of approaches are really detrimental. However, of course, it highlights some areas that are credible. And the use of AI and weaponry in defense is real and is serious and is out there for governments and attackers, and for people all over the world to take advantage of, so it would be naive to think that that isn't possible and isn't potentially likely in a future defense perspective. But I think the usage of it in this kind of environment doesn't help what we're trying to achieve. Back to what we were discussing in the previous conversation there, I don't think it's all about demographics. I'm going around schools, and colleges, also working with senior execs in their 60s and 70s, trying to find jobs, and working with people of my children's age in their 20s and people in their 30s. I think it's about mindset. And I think it's about levels of education. And I think it's about what you're exposed to, and what you're prepared to do for yourself to upskill yourself, and to think openly, and to move forward because you see that things are inevitable. And so I think these kinds of videos

really paint that bad picture and unnecessarily villainize the AI, which doesn't really help anybody.

Peter: I have talked with people like Peter Asaro that are in these campaigns to limit the development of these systems. And so I think they're faced with a difficult target of trying to gain awareness of this without it looking like sensationalizing, without it looking like *The Terminator*, and that's a very narrow needle to thread. But their concerns are legitimate in terms of how the proliferation of that technology could destabilize certain situations, Richard, responses?

Richard: If you look at India's strategies for AI for education, AI for business, and then you look at their AI for defense, it's nothing like the others, it is a completely different set of writing and standards and ethics. I obviously don't have an agreement with the vast majority of the population because I'm a vegetarian, and I don't believe in war or violence, so I think any discussion around the ethics of autonomous weapons wouldn't be something I could contribute to. Only that, if you're talking about the ethics of how to do this, I think you've lost the plot, personally. An AI autonomous drone, to me, is the same as a heat-seeking missile, you're just programming it with a certain end in mind. Obviously, there's huge risks and failures in its ability to do that. And it's just a testament to how far away we are from where we need to be as a population if these are the kinds of conversations we're going to have.

Peter: Speaking of where we need to be, we need to be wrapping up. And so everyone seems to have a book in the works, including me, I can get to announce that later on. But if you could talk about whatever is coming up next to you, but first precede it with here's your-- I'm going to nail you down, give one prediction for AI in 2022, that we can revisit in a year from now and test to see whether it came out. Who wants to go first?

Katie: I think the use of AI in the business function or business discipline of sales and marketing will be the biggest development that we'll see in the business world and the enterprise and the business adoption of AI in 2022. I think they'll be the ones that stand out more than any, and that will be visible to us in many different ways that companies market themselves. And I think if we trust and if we give our consent to that, then there will be a whole bunch of benefits, but we need that trust to be able to do that. So I predict that that will really be one of the big areas that we'll see. But we as individuals will only trust that if the companies can give us the assurances, whether it's GDPR, whether it's other regulations, but I think the latter part of 2022, we will see more and more usage of that. And the book that I've written is very much with that in mind, but it talks case studies of successes, failures, it gives examples as to how you can strategize, who the key players are, and so on.

Peter: Thank you, and we can put a link to that in the show notes. Who's next?

Ben: Well, I don't have a book to sell at this moment, I'm working on one in my spare time. But trying to create democratic decentralized AGI is taking enough of my time that writing books keeps getting postponed into the future.

Peter: Where would people find you doing that if you want to give something for people to locate you with? SingularityNET?

Ben: Singularitynet.io, and my own website is goertzel.org, plenty of information there. I think in terms of predictions, AI is going to be rolling out all over the place in this next year. One thing that excites me particularly is the use of metaverses, which are trendy and buzzy now, but as training grounds for AIs and early stage AGI systems. So I'm hoping that the metaverse hype and explosion, among other things, leads to the maturation of virtual worlds as early stage AGI training grounds. So let's keep an eye out for AI in the metaverse to become a thing.

Peter: Terrific. Richard?

Richard: And as an example of that - and I love Ben's comment there - would be if we are all going to have an AI buddy by 2030, 2040, whatever, next year, why don't we test it? Why don't we see what it would be like simulated in some of these worlds, and we can start to see what the risks and challenges might be now? So I'd encourage a Unity or a Fortnite or somebody, step in now and help us do that, and let's have some fun with that. And then to Katie's point as well, I mean, the one that we've been waiting for in business, I guess, is for language processing, language understanding, to get really, really good, really good. And I'm really excited about what that will do for all of us working when we have that transition to when we can reliably rely on this and when the context starts to be good. And maybe these guys will tell me that's still a bit a way away, but I'm really excited about that one coming through.

Peter: Thank you. Anything you want to plug, MKAI, or some of your other work?

Richard: mkai.org, yeah. Come and see as if you want, we're busy away. We did finally agree to write a book with our friends Bloomsbury, out August 2023, where we're exploring the new dynamics of digital trust and some of the things in this conversation, really.

Peter: I can't wait. Thank you, Prashant, who has left, Katie, Richard, Ben, it is all too short a time to do something like this. On the other hand, having that pressure of compressing our thoughts into a brief period can sometimes make diamonds out of coal. Who knows? All the best for 2022. This show is going out on December 27th. So this is our final show of the year. And thank you for contributing to our predictions and forward-looking for AI in 2022.

That's the end of the panel. What a pleasure it is to be a ringmaster for such amazing minds. Well, I'm either the ringmaster or the clown act that comes in between the trapeze artists and acrobats that dazzle you; not sure.

That's it for this show and that's it for 2021, may it stay in the trashcan of history. Next week I'll be talking with Tannya Jajal, UAE Chapter Lead for the Global Women in Tech Movement, technology contributor at Forbes Middle East and author of the new book, Thinking Machines: AI and the Intelligence Explosion. That's next week on AI and You.

Until 2022, then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

http://aiandyou.net