

AI and You

Transcript

Guest: Alexandra Mousavizadeh, part 1

Episode 149

First Aired: Monday, April 24, 2023

Hello, and welcome to episode 149! We're going to continue commenting on the global phenomenon that is ChatGPT, after our interview segment in each episode for a while. Today's guest is returning to the show: Alexandra Mousavizadeh, who joined us last in November 2021 in episodes 76 and 77. Back then we learned how she founded some important indexes, which is to say rankings, measuring the performance of countries and companies on AI and disinformation.

Well now she has another one: The AI Adoption Index, curated by her new company, [Evident Insights](#), where they call it the Evident AI Index, the first independent benchmark of AI maturity in business. They started with doing it for the banking sector, for reasons we'll hear in the interview, but plan to expand it to other sectors.

Alexandra started at Moody's as a sovereign analyst covering Russia, Central Asia and the Middle East and then went to Morgan Stanley, later becoming CEO of ARC Ratings, a global emerging markets based ratings agency, and then Director of the Prosperity Index at the Legatum Institute. Most recently, she was a Founding Partner at Tortoise Media, where she ran Tortoise Intelligence, the Index and data business. Here, she was the architect of the groundbreaking Global AI Index which I mentioned earlier, the first benchmark to track the strength of national AI ecosystems. She holds a degree in economics, mathematics and game theory from the University of Copenhagen.

If you're thinking that ranking companies by their AI adoption rates sounds easy, maybe you've been watching too much *American Idol*. This is not like Simon Cowell hitting a buzzer. Listen in the interview for how much goes into those calculations and making them useful to customers of the index, which could be investors, which could be financial analysts, regulators, or reporters, or could be developers or executives looking to make a career move. So let's get right into it.

Alexandra, welcome back to the show.

Peter, thank you so much for having me. I'm delighted to be back.

So, what's hot in AI in London right now? What's it like being in the center of the society there that's talking about AI? Give me an idea of the hot topics and buzz.

Well, I mean, I don't think that AI has ever been spoken as much about in the last many years as it is today. And it's a global debate that we're having, not only in the UK and the US but across Europe, the Middle East, and elsewhere in the world. The UK is caught up in the same things as you are in the States with the release of ChatGPT and other large language models being released. And the big question is how businesses are going to be using them and for what, depending on what sector you're in. And of course, there's a big debate about what guardrails to put around these large language model platforms. And one of the big topics this week as, of

course, you have seen, has been the call for pausing further research on ChatGPT-type models for six months, which was a letter that has been signed by many who are worried about the long-term implications of AI in general, but large language models and deep learning in particular, so have asked for a pause or a moratorium on this research. So that is a big debate going on right now, both at a business level, but also at a policy and a regulatory level.

What did you think about that letter? The press made a big deal of Elon Musk because his name was near the top - and of course, it's Elon Musk - but it was signed by a lot of people that are very thoughtful players in this space that don't seek out the press and sensationalism. So I am interested in your take on the effects of that and perhaps some of the motivations.

Yes. We've been thinking a lot about what was the motivations for Elon Musk, as we know he supported and funded OpenAI in its early days, to then step back and be critical of the release of what essentially he initiated some years ago. That said, he has voiced concern about AI and sort of long-term implications of AI and questions around existential risk around AI and AGI in particular. So in that regard, not so surprised that he would be leading in the drafting of such a letter. There are a lot of people who have signed the letter, the great thinker Stuart Russell from Berkeley, Gary Marcus, among many others who I deeply respect their opinions of and know that they have for many years been voicing concern about AI in the long term, AGI in particular, specifically with the - well, from Stuart Russell's perspective, the lack of understanding of what controls need to be in place in order to not create really big issues for humanity long term. And in Gary Marcus's view, there are some really significant issues relating in the short term to misinformation and disinformation with the models that have been released. So I think the concerns are both for the short-term implications of the release of these models and the long-term implications of the further development of these, alongside deep learning and large language models and increased compute power with the anticipation that innovations will lead to some form of AGI long term and what that might mean. So I think that it's not a clear yes-no answer to whether the letter is good or not because I completely understand the concerns that everyone has. I think the pause was also meant to generate some room to think about what the regulatory response would be and what could be put in place. And there has been ideas floating around on whether we should be setting up a global independent organization like the IAEA for nuclear development and disarmament, which I personally think is quite a good idea. And it's not to say that there are not a lot of efforts out there on the responsible AI side in terms of thinking through the principles that businesses and governments need to be signing up for in terms of explainability and fairness and bias and so on. There are lots of bodies that are trying to find solutions, but it's quite a fragmented space, a little bit like ESG was maybe 10 years ago. But it's still quite fragmented, so to have one body that pulls it together and sets in place some founding principles and guidelines that people should adhere to is not a bad idea. The question I have though is whether six months is enough. The other question that arises is whether, listen, if we in the Western world put a pause on this research and development, China is definitely not going to pause their research and development. So the question for us, I guess, is whether we do want to put that pause in and let the rest of the world keep developing it. Would we run the risk of falling behind and should we not, instead perhaps, I think the expression is "knit the parachute as we're jumping out the plane," which means let's be a bit faster in thinking about what those

guardrails could be. And maybe having these large language models being developed and released is really putting some pressure on us to think about what the response should be in terms of a responsible framework.

And it's a fascinating topic and not even the reason that we are having this interview, but it's also an indication of how interesting a time it is when you have a document that's signed by both Elon Musk and Gary Marcus. And I note the International Atomic Energy Agency was formed in response to proliferation of devices that had a known effect of blowing up cities and yet the effects of what we're talking about here are not understood. This is not even a singularity-type response. Anyway, so much that we could talk about here, but it's a pleasure to be talking to you today about how you're a serial indexer. You have started so many indexes and now you have formed Evident Insights to start the AI Adoption Index. And so tell us, first of all, what does adoption mean, the dimensions of that. Let's edge into this gradually, and so you tell us how this started, why this is important to you.

Yes. Well, thank you, Peter. I think, if you don't mind, I will just take a quick step back to how this all started. As you know, we were the architects of the Global AI Index that looked at how the strength of national AI ecosystems was developing. And that was back in 2018 where there was a real surge in national AI strategies to start to think about how to support the performance of AI development and deployment. And that time leading from 2018 and before, maybe a few years before, up until now, it's been more exploratory, and it's been looking at what are the components that go into this. And we've definitely reached a point now where it is private sector that is leading it. And we've reached a point where of course, many sectors and nations are still in an exploratory phase, but there's certainly come a point, a tipping point, if you will, where we've moved from a phase of exploratory AI development and deployment into a phase where the use cases and the use of AI is very clear. And the businesses that are adopting AI are now looking at not so much, "let's try this out," but it has been tried out, and there's a real question around return on investment. What does it mean to be AI-ready? How do we build businesses that can harness the breakthroughs in AI and we can harness them as quickly as possible? So given the demand there was at a business level and looking at what are the components that go into determining AI readiness, we formed Evident to build this new index on companies assessing these components that go into AI readiness and AI maturity, if you will. The components that go into AI adoption, to go back to your question, the ecosystem that forms that maturity is really four areas that we look at. One is the leadership component. Is the leadership very clear on what they want to achieve with AI? And what are they investing in and how are they thinking about their organization? In many cases, it actually does need a complete retooling and restructuring of the organization to be one in which that is more AI first or AI native if you're going from a traditional business and making that transformation. So the leadership component is very important. Then we look at the whole talent stack, all of the types of talent that go into AI development and deployment. And that we measure at a very granular level. We also look at innovation within a company, which means some will be focusing on developing in-house research capabilities. They might be doing that through research that they keep inside the company or produce and submit at conferences. We look at patents. We look at the citations of

the research and the patents. We look at what companies are investing in and acquiring in terms of AI capabilities, AI companies that are absorbing in order to make their AI capabilities stronger. And then we also have a fourth pillar looking at responsible AI. So while we take a deep look at the engine room of AI maturity, we also look at whether it's done safely and responsibly. And that is looking at adapting existing oversight to being able to also assess specific AI risk. So we look at responsible AI, we look at AI ethics, whether that is incorporated in the businesses, and so on. So those are the components that go into our assessment of companies on their AI maturity.

And so as you say, it's entering a new phase, and I want to get at what the use cases are for an index like this. If we were still in the initial startup innovative phase of using AI as a research effort in a company, then that would be an idea generated out of their R&D department and it wouldn't be important to a company who else was doing that. It would be, "Let's look at this thing." But now as you say, we've moved on and so there must be some case for people being interested in what the other companies are doing in a space. Who uses an index like this and what is it useful for?

Well, yes. Again, to step back, we have created our initial index on banks. And the reason why we chose the banking sector to establish the first index on is because, as I said before, they've been investing in AI for a long time, they also have an enormous wealth of customer data, and they have lots of different types of use cases that they've been able to test AI solutions on for a while. And while we are anticipating to go in and cover lots of different sectors, and going from the 25 banks that we have in our initial index to maybe 2000 companies that we'll have mapped in the next four to five years, we decided to start with banks for those reasons. They have invested for longer, they have a wealth of data, they have a number of use cases that they have tested and tried. So really a great sector to do the initial index on. And this is the first public benchmarking of banks on their AI maturity. It has previously been done by consultancies that would go in and use a survey approach and have the results anonymized and so really looking at more sort of a general benchmark. But what we've gone out and done for the first time is actually rank the banks on their AI maturity in a public ranking. And going back to your question on the use cases, which was also part of the reason why we chose the banking sector first, is that there are a number of use cases. And just to mention a few where AI has been tested and tried for a number of years is in fraud detection - one that AI solutions has been used for many years. Also, customer service. Banks use them for personalized product recommendations. Also for risk assessment, things like credit underwriting, portfolio management, cyber security, they look at trade finance, market analysis, and more often now also for recruitment. So those are the use cases that we've been taking a closer look at in the banking sector. But the research that some banks have created a specific hub for is more general research to develop things that are more about the cutting edge of the AI innovation space. Some banks have AI research hubs that would look at synthetic data, neural networks, large language models that they do deeper research in with the view that it will become useful across the banks, across all of these use cases eventually. And having a research hub within a bank, you'll be much better prepared for when the problems get raised by the different divisions and use cases bubble up to the surface in terms of where they're saying, "We need to find a solution for this. Is this something that our research hub might

have looked into already?” So that’s why some of the banks we’ve seen established these research hubs as far back as five, six years ago in order to be able to accommodate the questions that arose and the problems that arose from those use cases very quickly. Now, some banks go the other route, which is more of a bottom-up applied research route, and that means that it is the divisions and the use cases where they have research within and embedded in the divisions. So each bank takes a slightly different approach to where they place their research. And some opt for both pure research and applied research and some banks just simply use the applied research that grows within the divisions of the banks. So it’s interesting to see these two models, if you like, on how to position the research across the banks.

And I want to update my understanding of the maturity of the banking industry with respect to IT at this point because it seemed to me that there was a dichotomy - and from what you’re saying, maybe that’s been superseded - that you had, in the finance industry, on the one hand, very cutting-edge investment bankers using quants and AI to do ultra-fast trading and prediction, but also the banking industry was famous for maintaining the largest installed base of COBOL in the world. And at least the last time I looked, that was still the case. But you have been describing parts of the banking business where that might have been the case and you’re saying that those are now the subject of some research. So is my understanding of the amount of legacy code dominating the banking industry, is that now out of date? Should I revise that? Just scratch out COBOL from my lexicon?

On the trading side, as far back as the late ‘80s and early ‘90s, using advanced mathematics and models to have very sophisticated quant teams within the banks has been the case. And really, a lot of the early adoption of AI have definitely been in those divisions. The thing is that today really, every single division is touched by AI, and that’s the difference. It is everything, as I mentioned before, from customer service to loan models to efficiency gains from organizing your data better, then being able to use that for clients across many different divisions. So the change has been that now it’s universal for the bank. Now it’s not just in specific divisions where AI is being used. And essentially what we’re seeing is that all of those divisions are going to receive an uplift from an understanding of how to use AI in these divisions. And I guess that is the change. Because if you didn’t have a quant division in your bank, you might not really think about AI for the future existence of the bank. But today, it really is a do-or-die question for the banks no matter what they focus on, whether they’re pure retail banks or they’re more on the investment banking side. So that’s the shift. No matter what type of bank you are, you cannot avoid this question. So that’s what we’re seeing.

Is there a tension in banks in terms of risk tolerance? I mean, a bank has to be risk-averse where it comes to the security of its funds and its transactions, its data. But to be competitive, it has to take risks in some ways. So where do these conflict when it comes to AI adoption within banks?

Well, banks are highly regulated, and already have well-established risk monitoring systems within the banks that are monitored very, very closely, and even more so after the ‘08 financial crisis. There are very well-established risk management systems, and it’s not that easy to

implement AI systems in a highly regulated sector, especially when it has something to do that touches customer data. So that means that you've got - I wouldn't say a hesitation, but any use of AI systems on customer data in particular - is carefully, carefully vetted before it's implemented.

But I think people listening might be still thinking, "Well, that regulation didn't stop Silicon Valley Bank from failing."

Right. Silicon Valley Bank is, in my view, a slightly separate issue. They were well aware of the mismatch in their portfolio, maybe even a year before it collapsed, maybe even before that. They were aware of the mismatch. That was not an AI issue. You could argue that actually if they had better AI-driven risk management on the portfolio of the bank and had better alert systems for what might go wrong, and created deeper scenarios of what would happen with a rise in the interest rates, maybe that would've been something that they could have avoided the collapse of the banks if it was alerted to them with the scenarios that had shown how catastrophic this portfolio mismatch had been. However, I do not think that the Silicon Valley Bank collapse had much to do with AI and not installing AI management. It was more to do with I guess it was sort of a willingness to hope for the best and a lack of local regulation or regulators to really oversee what was going on in the accounts and the portfolio of the Silicon Valley Bank.

So looking at the pillars of the Index here and the results, what was interesting or surprising to you about the conclusions?

Well, again, a quick step back. We decided that we were going to measure the biggest North American banks and European banks in our first situation. That means that we didn't cover banks in Asia. So with that in mind, looking at how the rankings and the results of the Index fell, it was really interesting to see that the top 10 were really dominated by North America. And in the top 10, I guess my biggest surprise was to see two Canadian banks in the top 10, and that was the Royal Bank of Canada and TD. And I couldn't help but think whether it's driven in those particular cases by very visionary CEOs, a very deep understanding of the opportunity of AI. But also, I'm wondering whether the government's focus over the last five years of supporting the private sector in commercializing AI and the AI adoption maturity of the businesses might have also been a factor. So it was interesting to see two Canadian banks in the top 10. It was interesting to see that the different approaches that the US or North American banks had taken compared to the European banks. The North American banks are much stronger on the innovation side. So investing in, partnering with companies on AI, sort of really trying to look at ways in which that they can bring in the ecosystem and bring that into the banks. Whereas the Europeans who have been much more hesitant to build big partnerships with companies like NVIDIA and IBM and so on are investing somewhat less in AI companies and making acquisitions in AI companies to sort of bring that into the banks to build the AI capabilities that way. European banks have taken an approach of looking to build in-house rather than looking at where can we build partnerships that's going to enhance our banks and maybe we can get to strong AI maturity quicker that way. So that was interesting to see the different profiles of the banks and the different pathways to AI adoption that the North American banks had taken versus the European banks. So those were sort of my initial take takeaways. Now, with the expansion of the Index, we're going to be including banks in Singapore, in Japan, in China. And I'm very

curious to see what that profile looks like standing up against the North American approach and the European approach, and whether there might be some really interesting things to learn from the Asian banks, which I'm sure there are going to be some interesting takeaways from that.

It strikes me some of those conclusions are reflective also of the Global AI Index in terms of the rankings of the US versus Europe. And I wonder what will consumers of this report, people looking at this, what will they do with the conclusions that you've been describing here? Will they look at that and perhaps a researcher will say, "Well, I think I'll approach Royal Bank of Canada because that looks like a place where they would value my work"? Or what else might happen as a result of people learning the insights that you've uncovered?

Well, the hope that we have with the Index is that not only the companies on the Index can take a very granular look at the profile of their banks when it comes to their AI approach and see whether it aligns with a conscious strategy and vision and be able to refine and reflect and maybe update the AI strategies for the banks in the Index because there's a lot of detail and data that underpins the ultimate rankings of the banks, but it really is also a tool for those about to start an AI transition or build an AI strategy for the first time. Because it really lays out all of the components of the areas that you need to be considering when looking at building a strong AI ecosystem for your bank or for your company, more generally speaking. So that is the hope that we can provide that insight into what does it mean to be AI-ready? What are the pathways to a successful AI adoption? And yes, there are also definitely those banks that are doing well in the Index would also be able to take those results as a talent attraction tool because there certainly is a war on talent right now between the tech startups and the Big Tech and other sectors other than banks. But banks themselves are also fighting for the same talent with the capabilities that are attractive to many sectors and there are lots of interesting problems that AI talent can work on these days. So in order to display that your bank is a place where you do interesting research, as you're doing, a lot of innovative thinking, there are really interesting use cases to sink your teeth in and actually that it has practical real-world application. Now that can be really attractive for fresh talent that are coming out of university or maybe the tech industry and want to poke their head up and see what else is around. This is an opportunity for the banks to show that they're doing really, really interesting work in and across the banks. So yes, it could be for a researcher looking at which bank or which company is doing interesting work that I might be interested in doing, and is it a place that will harness my work, also retrain me, and build my skills as I go, and a place where I'm able to grow.

That's the end of the first half of the interview. I thought it was very enlightening as to how an industry I'd thought of as being very conservative in its technology choices is evolving with cutting edge AI technology like ChatGPT, which is a good segue. On the subject of ChatGPT, well the first comment is that I'll use that as the most convenient label for the large language systems that are upending the world at the moment, since it's the most widely used term in the public sphere, even though ironically it's been succeeded by GPT-4 by now. I mean, I saw an episode of *South Park* all about ChatGPT, and I don't think you could ask for a more obvious indicator that something has gone mainstream. As you know, we're not big on sound bites here, our style is to give you a useful story that goes behind the sound bites. And there are a lot of sound bites about ChatGPT out there right now, so we'll probably do

another episode dedicated to it before much longer. We've been talking in the interview about banking, and there's a prime example of a sector that's going to have a love-hate relationship with large language models. On the one hand, banking transactions need to be completely airtight, there is no room for creativity, let alone hallucination, in making retail, commercial, or investment banking transactions in any way. On the other hand, when it comes to customer service, marketing, and other content-generating functions, banking has as much potential as any other sector to leverage large language models. One of their own, Goldman Sachs, is extremely interested in the impact of large language models, and – here's our news ripped from the headlines for this week - [estimated](#) that AI automation could affect up to 300 million jobs worldwide. Now, that doesn't mean 300 million people out of work, it means that Goldman decided that generative AI could substitute up to one-fourth of current work, which in aggregate multiplies out to 300 million full time jobs but in practice is going to mean a lot more people having some parts of their jobs automated, which may not impact their employability at all but instead accelerate business. We had a similar discussion about the Frey and Osborne 2013 study on the future of employment which concluded that 47% of job tasks were at risk of automation. If this sounds like a uniformly negative analysis, consider the possibility that this automation could increase productivity to the point where it reverses a worldwide recession that we've been flirting with lately. Of course that doesn't help you if you're worried that you'll be laid off. But unless you're writing headlines for BuzzFeed, that's not a sizeable threat at the moment. Look instead for the opportunities to use large language models to improve your personal effectiveness and efficiency, and you can become more employable, not less.

Next week, we'll conclude the interview with Alexandra Mousavizadeh, when we'll talk about the reaction of the banking industry to ChatGPT and how they're using it, what the Index shows about the movement of talent within the industry, and how the AI Adoption Index will evolve into other sectors. That's next week, on *AI and You*.

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

<http://aiandyou.net>