

AI and You

Transcript

Guest: Alexandra Mousavizadeh

Episode 77

First Aired: Monday, December 6, 2021

Hello, and welcome to episode 77. We're going to conclude the interview with Alexandra Mousavizadeh, who is a partner at Tortoise Media, which describes itself as a "slow-news" media company (so Tortoise is on brand there). She is the Director of the Tortoise Intelligence team specializing in global and industry indices and data analytics. She launched the Global AI Index in December 2019, which we talked about last week. She holds a degree in Economics and mathematics from the University of Copenhagen and was head of risk management at Morgan Stanley in London. Recently I was at the Cyber Summit she was chairing from London and watching a session where the panelists were Tom Hurd, the UK's former director general for the government office of security and counter-terrorism, and Nigel Inkster, the former head of MI6. Maybe it's a good thing I wasn't there in person because I couldn't have resisted the temptation to ask if people called him M. But that gives you an idea of the sort of thing Alexandra is up to.

She also started the Global Disinformation Index, which we'll be talking about this week. Their site says "Time to defund the coronavirus infodemic," and I have to tell you, I got emotional just reading that, because the disinformation battle has felt *so* one-sided for *so* long, and for years I was facing that without the terms 'disinformation' and 'misinformation' even being socialized to the point of familiarity, trying to figure out how to talk to people about something I saw as so important, that was affecting people around me, and now of course I'm not alone, and the odds are that you've been affected by disinformation – well, we all have, it's tipped key votes and elections – but I'll bet that it's also had a personal cost to you because just about everyone I know has someone in their circle of close people that has been infected by this, and I don't think there's a better word for it than that because the analogies between disinformation and infectious disease are quite numerous and revealing. And also point up how people are debilitated by disinformation; it's not just a... lifestyle choice, if you will, a foible like supporting Tottenham Hotspurs over Manchester United, say; it's costing our society in the way that people are harmed by these wrong beliefs – and they don't even know it. They're not just expressing a preference for a debatable option, they've abandoned science to embrace something else just because it appealed to them, as though that were all it took to defeat reality.

You know I talked about this a few episodes ago when I mentioned a few of the bizarre and ludicrous beliefs that are getting traction in an unreasonably large number of people, and you can experience this cost in a personal way when you're talking with someone whom you thought you knew and suddenly they come up with one of these lunatic theories and you have to push the reset button on how you think about that person. A direct impact is in how trustworthy you believe them to be, because if you can't trust them with reality, then what *can* you trust them with? Can you let them run your business, watch your kids, walk the dog? You get that I'm passionate about this topic. I can't and won't spend a large amount of podcast time on it, though, for the sake of balance, so I'd encourage you to listen to Sam Harris, because he's done a lot of shows about it and I agree with every word I've heard.

Our show is about AI - and AI has facilitated the spread of disinformation on a scale never before possible; largely through the recommendation algorithms on social media, which are driven by cluster analysis, and the goal of getting people to spend more time on the platforms, so that's both unsupervised learning and reinforcement learning in use there.

Stanford's 100-year AI project, the AI100 2020 update said: "AI systems are being used in service of disinformation on the internet, giving them the potential to become a threat to democracy and a tool for fascism." And that's why we're talking about it. I got into this conversation with Alexandra at a time that the Facebook hearings were peaking in Washington DC, so incredibly timely.

You'll hear us lead off in this part talking about the AI race between the US and China before we talk about the disinformation index and then a personal take from Alexandra on what's important to her about what she does. So let's get back to the interview with Alexandra Mousavizadeh.

Kai Fu Lee talks about the global AI race as being United States and China, and then a lot of much, much smaller players that, in his words, can't hope to compete. First of all, would that match your assessment? Or perhaps that's only considering a few of the parameters that you're looking at on the index? And secondly, how relevant is that to what the other countries should be doing with AI?

I mean, first of all, I think it's, it's really worth mentioning that we know, that the US and China are far ahead, right. And Europe is quite far behind third place. And in Europe, the UK is leading on AI to ecosystem and development, but they are really far behind. If you look at AI index, for example, it gives a very sort of clear picture of exactly sort of the gap between the US and China and the rest. But there's also actually quite a significant gap between the US and China. Now China is, is because of its quite focused policy are our leaders on many, many things when it comes to AI, probably also quantum compute. Very far ahead in facial recognition, very far advanced in AI adoption, for many reasons that we can get into. But the strongest ecosystem, by far, is the US ecosystem of measured on all of those, you know, parameters that go in, it's not very narrow, it's very broad and very deep on every aspect.

China has said that they intend to overtake the United States in AI by 2025. By whatever criteria they're using, do you think they will?

Well, it's interesting. And you as you were mentioning, Kai Fu Lee just before, and their position is that they will on the things that are, I think they deem important than AI, which is quantum computing, facial recognition, machine learning from sort of the tools perspective, they have a very good chance of leading (maybe already are) by 2025. But the question is, there's much more to AI than that. And there is this sort of the ecosystem that spins out lots and lots of companies that use AI across many, many sectors. And if it's on those parameters, and looking at leading research, not just on maybe a few things, but on many things that are related, then the US definitely is light-years ahead when you're looking at the talent pool. China has a very strong talent pool, but the US, you know, it is in the US where the strongest talent pool lies for sure. But and when it comes to breakthroughs, there will be I'm sure will be aspects where China will lead there's no doubt. But it is also very clear that having a deep bench and having sort of a wider scope is also really important for AI innovation and AI adoption. So, I think he's right; I think

they may lead on certain things, but not as an ecosystem and a strong AI ecosystem in general, I think they're quite far off.

So, on your index, the US has benchmarked at 100 to put some numbers on the US and China is 58.2. And then the next players, United Kingdom 39.7, Canada 37.6, the rest of them are grouped together quite closely. Do you see on your index China catching up with the US?

No there's a very, very big, I mean, we measure the strength of the entirety of the AI ecosystem. So, we, so in terms of the ecosystem, China's very far behind. And I think that I mean, we've now run this index for three years, and you can see the gap closing every year. It's definitely catching up on the ecosystem. For sure. If you extrapolate that out, its which I actually haven't done as in, I haven't built the model that would say, actually, on this trajectory, it will take X number of years, but it is quite far behind. So, I would say, at least, to getting close to the US will take 10, 15, maybe even 20 years. If at all; because it's a very different policy environment, it's a very different ecosystem, it's a very different business mindset, there's a much more top-down approach, which might prevent the creation of that kind of ecosystem that the US has, which is basically the opposite is a bottom up. It is very diverse, it's not solely focused on certain aspects, it sort of has it all there. So, I think it will take a long time before China catches up to the US on our index, when it comes to specific breakthroughs. You know, I think probably, China's already there, but hasn't already surpassed the US and certain breakthroughs for sure.

I'd like to talk about the global disinformation index, which you started, and find out more about that, because that has a lot of intersections, as we already mentioned, with AI. It's also incredibly volatile. I mean, people are dying as a result of this daily. And the political and military stakes are enormous. We talked earlier on the show about Taiwan and what they face from China and, and some of what the United States faces from Russia. And the consequences of this are measured in decisions like Brexit, and the US presidential election, and Coronavirus, misinformation causing under-vaccination, and many, many excess deaths, and we're already talking about things that cause people to riot. And you stepped into the middle of this maelstrom and said, "At the GDI, we look at disinformation through the lens of adversarial narratives that undermine trust in our social, political, economic and scientific institutions." What was your goal in doing this? And what effect has it had?

It's a Herculean task to try and measure disinformation. And many people are attempting to do this, I think that we set out to try and use AI and machine learning to the best of our ability to capture the markers of disinformation. This is not a new technique. This is something that Facebook and others use internally. But we felt it was very important to have a third-party independent index that set standards and tried to capture this in an independent body that could then start to map where we saw disinformation incidences crop up. It was also to start the conversation about, What is it? How do we measure it? Who should be measuring it? Could this lead to global standards on disinformation, could it then lead to thinking about sort of the regulation behind it? You know, that's obviously very much being talked about today with Facebook and what are the responsibilities of the tech platforms. But trying to create a set of metrics and thinking about this, that could be adopted across not just Facebook and Twitter, but

sort of really across or similar type platforms around the world. So, trying to convene some thinking and converge on some metrics that said, "Look, these are the ways that we should think about it, and these are the ways that that we can measure it." So it's really about building the framework, because collecting the data is - there's millions of fields that that we trawl through, through sort of a fishing net, if you will, that look at it, and sort of just from a technical metric perspective, like what are the what are the characteristics of disinformation, what are the markers of them? And given what we know about it today, what are the kind of filter and the metrics that we can use to capture it. And the issue is that it's constantly changing, so to start one place, and it's go with the change, and keep refining the metrics, and keep going, even if we're on the back foot, because the people who are trying to infiltrate political elections and change mindsets and do all the things that we know that - I wouldn't say our enemies, but people who want to impact societies in a negative way, are doing that they're far ahead, right, they're miles ahead. And it would be naive to think that we could actually prevent that from happening. But I think we felt very strongly that we at least had to take a stab as an independent third-party body to try and set some metrics and some thinking around what this could look like. And as we get better, and faster at reading the amounts of information that comes through these filters, at least, our system will develop as the disinformation also grows. But if we didn't try, we wouldn't have a starting point. And we need to have a starting point.

Well, and I really have to take my hat off to you for doing this because I have been despondent about disinformation for some time, in that it transforms the reaction, we have the standards that we have about what's right, to those people that it affects. So it really has the potential of having no bottom and you have done something here. And you've created these measurements, which you mentioned. And I wonder if you could briefly describe what do you measure? How do you think about measuring disinformation?

The way that - and by the way, it's it is constantly changing. But if you sort of just as a baseline, you look at what has been found through quite meticulous analysis of the disinformation that we've that we know of that have come in. And we look at those markers, and what are the markers on that metadata. And those markers are what we use as indicators to try and sort of catch things that are coming in. And what I mean by that is that in disinformation campaigns that we've seen in the past, [20]16 elections and things that have happened subsequently is that campaigns are often from, you look at the location - is it from somewhere that might be slightly obscure, you look at whether it's a very new platform that is safe, it's only two weeks old, but it has 200,000 posts, there's something off, right? You look at ownership, you look at the structure of the platform that is that is that usually are markers for disinformation. There's no bylines, it's established a minute ago, it has very, very high volume. There are certain things that are very hard to find on it. So, the sort of those there's about 30 of those markers that that that have been established from previous disinformation campaigns. And so that's what we sort of set out you know, they fall into five pillars and so on. And that's what we set out as a sort of the first filter. And as those changes happen, and we identify, maybe as we go, we identify other markers that are prevalent when it comes to disinformation campaigns, we would then add that in. So, we're at least at a really good starting point. And we can then add on to that. And those are the filters that we create sort of with our tools and using the filters and the algorithms that we use. And we

set that up to track that and see what comes in. Then the question is, well, what do we measure on? We can measure on the pipeline of information that we're measuring on. But it's a big ocean out there. So, we're just also you've got to, you know, how do you filter all the information coming out? And you've got to make some decisions on what publications are you looking at? What tech platforms are you using this filter on and so on?

Well, that's really interesting. And I hope we have a chance to dive into that more at another time. I know our time is limited now. And you have a story on the Tortoise media site that I have to read here because it's compelling. And so here it is. "Once upon a time, there was a race to understand the world. Everyone hurried off, all over everywhere to find out about everything, all but one who set off deliberately and thoughtfully. We have too much information and too little time, she observed choosing to explore a few things that really mattered more deeply. The others moved faster and faster, produced more and more, but it added up to less and less. They saw it all, it was a blur, they were dizzy and had to have a lie down. Meanwhile, she took her time, chose some subjects, sought out the people who knew and listened, she paid them attention. She saw less and saw it clearly. She won. "There's a moral to this story," she said - she was something of a moralizer - "Slow down, wise up." So is that the voice of the tortoise?

That is our voice. That is that is what we that's how we that's what we're guided by every single day, when we wake up, we try to cut through the noise. I run the Intelligence and Research side at Tortoise so through the indices, we try and do that with these longitudinal studies, and taking our time and getting really deep. But on the journalistic side and on our events and our audio, also there is the guiding principle that we slow down, and we wise up. And I hope that that is something that you really feel when you read our pieces or listen to our podcasts and stories that these are very carefully selected topics that we track. And we track a lot of topics, but we they're very carefully selected. And we take we take our time. And we really go deep into the sort of not only the investigation, but also what does this mean? So we do pick some things that we really run with as campaigns alongside our investigative journalism, we'll be going to all kinds of topics such as, we just did one on the diamond trade, but we've done on hidden homicides, on the unreported violence against women we've done or something that we've worked on a lot, because we actually happen to have a journalist whose mother was murdered in Malta, for being an investigative journalist herself, and the government had her assassinated. So, we go into the injustice of that and what that means and also what that means for journalists that are exposed elsewhere and what risks that they run. And we go into, obviously, deep into the vaccination campaign. We look at inequality, we look at climate, also from a journalistic perspective, but also from a data perspective. So we're really trying to cut through, but slowly.

And you're clearly driven by this enormous passion to make the world a better place. You've attracted people in that respect. If people are listening to this that are equally passionate and want to get on board and think that they have the resources and skills to help you, what do you want to tell them?

Well, come on board. I mean, become a member, I would say. So, send me an email Alexandra@tortoisemedia.com; please do get in touch. We'd love to hear from you. One of the things I think worth mentioning is that we host what we call think-ins. Sometimes we have three a day, but we definitely have our evening slot. Those are unique in the sense that this is where we open up our newsroom. So, if you want to get involved come and join us, we open the platform almost right away to listeners, online. They're all hybrid, or in the newsroom. And actually, our stories are very much driven by what we hear in the newsroom. And a lot of the things that we run with are things that get raised there. And we find that actually, that is really where the stories are. It's people who come in with their opinions to topics that we raise. And so, get involved by participating, come and speak up, speak your mind. We want to hear your opinion and get involved in our journalism that way. Or just read us.

Alexandria Mousavizadeh, thank you for all that you do. And thank you for coming on the show.

Thank you, Peter, for having me. It's been such a pleasure. What a lovely conversation. Thank you.

That's the end of the interview. It was useful to get Alexandra's take on the competitiveness of China as a counterpoint to a lot of popular rhetoric about their being on the verge of overtaking the US. I'll quote from the disinformation index site about their methodology, for some detail here. The index is based on four pillars: Structure, Content, Operational, and Context. **Structure** looks at the metadata of news domains and uses AI trained on a sample of 20,000 known disinformation domains. **Content** is assessed manually to look at the credibility, sensationalism, hate speech, and impartiality of ten of the top-shared articles which are anonymized so the researcher evaluating them doesn't know where they're from. **Operational** looks at the policies, standards, and rules of the domains to focus on conflict of interest, integrity, and accountability. And **Context** assesses the reputational practices, reliability, and trustworthiness of a news domain, using an independent expert survey of respondents from across the political spectrum. You can find a link in the transcript to the [GDI](#) or just google for [Global Disinformation Index](#).

In today's news ripped from the headlines about AI, maybe you've heard about ImageNet. It's a collection of labelled images that was pivotal in the early training of deep learning algorithms and was largely the result of groundbreaking work by Fei-Fei Li out of Stanford. ImageNet is used to train image recognition AI. Well, in a big step, the curators of ImageNet have blurred out all the faces in it. They wrote in a blog post, "The dataset was created to benchmark object recognition — at a time when it barely worked. Today, computer vision is in real-world systems impacting people's Internet experience and daily lives. An emerging problem now is how to make sure computer vision is fair and preserves people's privacy. Experiments show that one can use the face-blurred version for benchmarking object recognition and for transfer learning with only marginal loss of accuracy."

I don't have any editorial commentary on that, I just think it's a notable milestone.

Next week, I'll be talking with John Zerilli, a philosopher who is a Leverhulme Fellow at the University of Oxford, a Research Associate in the Oxford Institute for Ethics in AI. We'll be talking about his book, "A Citizen's Guide to Artificial Intelligence." 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>