AI Ignition

Ignite your AI curiosity with Oren Etzioni

AI for the common good

From Deloitte's AI Institute, this is AI Ignition—a monthly chat about the human side of artificial intelligence with your host, Beena Ammanath. We will take a deep dive into the past, present, and future of AI, machine learning, neural networks, and other cutting-edge technologies. Here is your host, Beena.

Beena Ammanath (Beena): Hello, my name is Beena Ammanath. I am the executive director of the Deloitte AI Institute, and today on AI Ignition we have Oren Etzioni, a professor, an entrepreneur, and the CEO of Allen Institute for AI. Oren has helped pioneer meta search, online comparison shopping, machine reading, and open information extraction. He was also named Seattle's Geek of the Year in 2013. Welcome to the show, Oren. I am so excited to have you on our AI Ignition show today. How are you doing?

Oren Etzioni (Oren): I am doing as well as anyone can be doing in these times of COVID and counting down the days till I get a vaccine.

Beena: Yeah, and how have you been keeping yourself busy during the pandemic?

Oren: Well, we are fortunate that AI is basically software with a little sideline into robotics, and so despite working from home, we have been able to stay productive and engaged. It's just a little bit less fun because one of the things I like to say about AI is that it's still 99% human intelligence, and so obviously the human interactions are stilted.

Beena: Yes, and I know Allen Institute for AI has been working in AI research and played a role in the fight against COVID. Can you talk a little bit about how AI research broadly has helped with the fight against COVID.

Oren: Sure. So, there are actually several ways in which we have signed up to help. We as the AI community and, of course, the Allen Institute for AI, or as we call ourselves, AI2, our mission is AI for the common good. So most broadly, we have looked at models for predicting, most broadly the AI

community has looked at models for predicting pandemic outbreaks, and there are reports that an AI company actually anticipated the COVID pandemic well before it was announced and confirmed by China. So that's number one, anticipation. Number two, the search for a vaccine, the search for therapeutics, that has been AI assisted in lots of different ways, at Moderna, by all accounts, for example, a Harvard Business School case study on the subject records that. In our own work, we have taken all the relevant publications, both preprint and journal, that are available and put them in a machine readable corpus that's a body of text that's machine readable, and a variety of people have built tools on top of that corpus to answer various questions, and at the height of the controversy, for example, about whether mass masking was valuable or not. We were able to provide a very data-based answer to that question and many other ones, like what temperature kills the virus and so on and so on.

Beena: One of the things that I believe what that research collaboration has enabled is how we share data for the common good, and it aligns with the mission of your company, AI2, which is about using AI for common good. Do you think the way we share data coming out of the pandemic has actually evolved beyond even the pandemic or solving for COVID-19, how do you think data sharing will look like in the future?

Oren: So data is a huge topic, as you know and, of course, modern AI methods are driven by data. I like to break data into different categories. So, there is personally identifiable data that is about an individual and could be used to target advertising, to surveil them, to castigate them, or even imprison them in nondemocratic regimes, and I think that that data is one kind and it needs to be very carefully regulated and guarded. A second set of data is more ambient data, maybe data about language or data that's commonly available on the World Wide Web, and that can really help us build AI systems that understand language better and so on. But probably the most interesting kind of data and the one that perhaps you had in mind is medical data, and there it's really tricky because on the one hand we really need that data in order to come up with treatments, come up with vaccines, generate insights in the medical sphere. On the other hand, that data is obviously very sensitive. Generally speaking, the solution to that is to anonymize it, to remove identifying information so that we can have access to the data en masse, and modern technology does allow us to do that. The question is regulation, HIPAA compliance and things like that, is it keeping up with modern methods?

Beena: Yes. That's the bane of AI's growth, can regulation and policy keep up with the innovation that's happening. And I know AI2 works beyond health care as well. Can you share some of the additional projects that AI2 is working on that you are most excited about?

Oren: I would love to. So, I eat and breathe AI, but AI2 is really my professional home. Again quickly, we were founded by the late Paul Allen in 2014 as a nonprofit research institute. We have grown to more than 120 people, mostly based in Seattle, but also Irvine, California, and Tel Aviv, Israel. We also have an incubator that has now spawned off quite a few start-ups and has maybe 100 people in the different companies working there. We work in two broad areas, one is natural language processing. So, any kind of text signal, audio signal, we love to analyze and draw conclusions from, try to understand what's in text because, of course, it's quite opaque to the machine, and often it's not just the explicit text, it's what's implied and so on. The second area that we work on is computer vision. So, there are again a lot of questions about the visual signal in images and pictures, in diagrams and videos, and we have a team that is at the cutting edge of analyzing that. We have published hundreds of papers, garnered best paper awards, released open source software. So, we are trying to be a force for progress in these arenas. The

last thing I want to point out is that this is a kind of data, text or visual data, but it applies across the board. So, if you give me an arena—health care, construction, you name it—and I will give you five applications of AI in those arenas. So, it's really across the board.

Beena: Yes, besides what you are doing with AI2, I know you are very plugged it into the start-up ecosystem as well, having founded a few start-ups of your own. Broadly, what are some of the challenges that AI is addressing in the real world that is most exciting for you besides the work you do at AI2?

Oren: I think Andrew Yang, formerly of Stanford, said it best. He said, "AI is the new electricity." Actually, I take a little bit of exception to the electricity metaphor because it suggests that AI is plug and play, you just plug something into the wall and AI flows out, that's not quite right. There is a lot of blood, sweat, and tears. But the part that Andrew Yang just absolutely nails is the fact that AI is transformative in the same way that electricity was. So if we take any arena—loan processing, criminal justice, health, cars, you name it—AI is there. I think that there is a tremendous potential to save lives. There are plenty of issues to think about, there is privacy, there is job displacement, and so on, and much has been written about that, but I think not enough has been written and talked about the issue of how AI can save lives in medicine, on our roads. There are 40,000 highway deaths each year in the US alone, and studies suggest that we could cut that by 80 to 90% using AI methods.

Beena: Yes, and I have heard your TEDx Talk where you are talking about AI. We focus so much on the headlines of AI being evil and destroying a lot of things, whereas we should be focusing more on the value that it is creating and some of the challenges that could come along with the value creation. Are there any industries that you think AI will not transform?

Oren: I think that's almost an oxymoron because what AI is about is making decisions better, faster, and at greater scale. So, any place where you have decisions being made, and that's literally everywhere, AI is at the table. We have a hard job to do. We have to worry about bias. We have to worry about privacy, as I mentioned, but we also have the potential to make better decisions, to actually remove bias. So again, a lot has been written about how AI because it's databased can capture old biases and then in the models reflect those, even amplify those. But let's also remember that people are highly biased and we can actually use AI to counteract bias, both cognitive bias that leads people to make mistakes, for example, it has been documented that doctors that are exhausted and overworked aren't necessarily good at the probabilistic calculations that might drive a key diagnostic or a treatment decision, and also people of course have their races, sexes, and other kinds of biases and AI can serve to counteract that.

Beena: Yes, I remember reading in one of your interviews late last year about the prediction for 2021. You said, "2021 will be the start of a renaissance in tech as people prioritize jobs that have real and direct impact, direct benefit for humanity and the planet." So, I absolutely love this quote. I am a technology optimist too, and I think there is a lot of good that can come from using AI and technology as long as we think about the ethical impacts. And I have spent a lot of time thinking about the impact of technology beyond just the value creation. So one of the things that I have been focusing on is the ethical impacts and how do you solve for it. How do you get beyond the philosophical discussions to actually solving for it in the real world, in large enterprises, even beyond big tech, in enterprises where AI is being used, whether it's to make a loan processing decision or to predict when a machine might fail, which is on the field. So my question to you is, what are some of the tangible signs that you see that

makes you so optimistic that we will see significant improvement in ethical considerations in companies that are using AI today? What is some of the positive progress that you have seen in this space?

Oren: Well, first of all, the focus on the issues of ethics is unprecedented. So if you look at academic AI and just search for ethics in the conference paper titles, there is an enormous amount of that. There is a whole conference that's been founded that focuses on fairness, accountability, and transparency. Secondly, in my conversations with members of Congress and their staffs, this is an issue that often comes up. So the awareness that we have ethical responsibility here has always been part of technology and software. There is an almost ancient organization called Computer Professionals for Social Responsibility. But with AI, because of the prominent role it plays and because of autonomy that it can potentially embody, I think there is a very strong awareness, that's number one. Number two, there are a set of technical problems that come up. It's not just philosophy and vague talk, and we are actually making progress on these problems. So, for example, the first set of techniques that we built to support decision making did make gross errors, did amplify bias. Now we have learned how to counterbalance that and impose more constraints on the behavior. I still think though, actually particularly if we think about AI in the military, and there is definitely an AI arms race going on right now, the notion of autonomous weapons. So again, there is a difference, an important distinction I would draw between intelligent systems and autonomous systems. Intelligent systems are ones that use advanced algorithms to come up with hopefully better decisions. Autonomous systems are ones that make that on their own, make those decisions on their own. Those aspects, intelligent and autonomous, are sometimes conflated because we are both, we are intelligent and autonomous, but they really need to be separated. So, intelligent weapons should result in less loss of life and more targeted behavior. On the other hand, autonomous weapons are very scary because you have a machine deciding on its own whether to take human lives, and I don't think we want that, particularly given some of the brittleness that we have seen in AI.

Beena: Yes, I remember seeing that graph that you were showing in your TEDx Talk with autonomous and intelligent and that was a great explanation.

Oren: I feel like you remember my TEDx Talk better than I do. I really appreciate you delving into it.

Beena: Because I have a follow-up question on that because in 2016, I think it was in 2016 that you said that AI can do a lot of good and we are nowhere close to autonomous intelligent systems. Now about five years later, do you still think we have not moved toward AGI or that superintelligence that you hear hype about or have we made progress in that direction as well?

Oren: It's great to have the opportunity to reflect on that question roughly five years later, and in fact what we have seen is continued progress on what I would like to call AI savants. These narrow, targeted systems, whether they play chess or go or analyze radiological images or predict pandemics, they are very narrow and targeted. There has been no progress whatsoever on autonomy. Show me an autonomous AI system. Show me an AI system that can cross the street. Show me an AI system that can formulate its own problems the way that we do. Show me an AI system that can prepare for an interview the way you clearly have. It remains to be the case that we are very good at these thousand points of light, at these very narrow targeted tasks, but we really struggle to build general intelligence and autonomy in a broad sense is very far afield.

Beena: Oren, I have listened to your TED Talk and I always like to look back and kind of see have your thoughts evolved, has how you thought about it at that time changed, because we hear so much about the progress. I also give the example of we are very much in the artificial narrow intelligence phase, where you can solve very well-defined specific, narrow problem, but it is not as catchy a headline. So, most of what you see in the media is on the artificial superintelligence, which is where AI is taking over the world and replacing all of us, which is so true. You have seen this whole space evolve. You have studied it for a long time, and AI has replaced some jobs, but it has also created new jobs and new roles. Have any of the new jobs that were created been a surprise to you compared to 10 years or 20 years ago? Did you anticipate some of these new jobs coming up in the future?

Oren: I do think that job displacement is a complex issue. Let's first of all start with the broad context, which is computer technology and digitization does cause increased efficiency, but a job loss. I think that that's actually completely independent of AI. I like to say that email has cost us more jobs than AI has. But if we look to the future, particularly in industries like retail where Amazon is opening automated stores and transportation where autonomous vehicles are still far off in the future, but they are coming, and the trucking industry, for example, Uber and Lyft driving, these are very vulnerable jobs for millions of people. I don't think it's realistic that all these jobs are going to be replaced by AI-created jobs. The AI-created jobs can actually be better, they can be better paying if you understand AI, if you can use it as a partner. But we do really need to think about the millions of people and they are some of the most vulnerable who will experience job displacement. I saw a headline a while back that talked about coal miners to data miners. The idea that people will be retrained. And we should invest more in retraining, but we can't expect all of our coal miners and all of our Uber and Lyft drivers to become data miners and so on.

I want to touch on the point of what's new most recently in AI, and that also implies a whole new class of jobs and a whole new class of capabilities. So, one of the things that's really emerged in the last couple of years is generative AI techniques, and generative is a technical term, but kind of like it sounds, it's not just AI saying yes or no, good loan, bad loan, this email message is spam, not spam, categorization decisions, but also generating information. I am sure you have heard of GPT3, and what we are seeing there is the ability to generate documents, generate emails, even generate pictures and artworks and music and even videos. So, this automatic generation capability of AI is remarkably humanlike, which both opens tremendous challenges—we have to worry about AI forgery—but also tremendous opportunities. AI can help us author things for the first time, can help us scale the creation process, can change customer service to be much more responsive and immediate. So there is a ton of potential around generative AI techniques.

Beena: Has there been any job or role that you saw come up recently that caught you by surprise and you were like, huh, I never thought there would be a job like this. I will give you one which caught me, which was like where I had to take that pause, it was a job posting for sensor cleaners and it is for the vehicles which have a lot of sensors. It was actually a job posting because those need to be cleaned in a specific way. You think about the technology and the kind of data that you can get by plugging in IoT and getting that data and being able to drive insights. But there is basic work that needs to be done to keep those sensors clean and functioning. So that was one where it's not a cool, sexy job, but it is something somebody has to do it right, and I am sure we will come up with some way of robotically being able to solve for it, but this one was one of those where I paused and I was like, I didn't see this one coming.

Oren: Well, on the data side, there are data annotators. There is still no substitute for human ability to label data, whether it's Amazon Mechanical Turk or many of their competitors, and not just labeling simple data. It can actually be quite complex. There are also data curators. So, it's not just a matter of labeling or cleaning the data, but even deciding what bodies of data to show the machine. And then there are jobs of auditing AI. So, it's very important, if we have an AI system, we need to look at it and decide is it doing what's intended, auditing and analysis, and then last but not least, it's partnership. So, we have for a long time, decades, talked about human and computer interaction, but now a new field is emerging, which is called human AI interaction. So, how humans, let's say a doctor or a knowledge worker, a secretary, how do these folks interact with AI systems so that together they can do a better job than either the AI or the person alone. So it's new jobs, but it's also new tasks and new capabilities and new activities within existing jobs.

Beena: Yes, absolutely. It's new jobs, job descriptions that is about driving more of how do humans and machines or humans and AI work most effectively to make progress together. I love it.

Oren: Recently, I had a medical issue and they told me I needed surgery. I am very thankful that I didn't end up having to have it, but they explained to me that the surgery was going to be performed by a doctor and some robots. So robot-assisted surgery is very common today and it results in better outcomes and particularly the robot isn't making the decision should you have the surgery or what kind of surgical procedure to have, but when it comes for stitching and various other low-level activities, the robot sure has a steady hand.

Beena: Let me ask you this. How did you feel about it as a patient that it will be a human but there will also be a robot involved in the surgery? How did you feel as a patient?

Oren: Actually, I felt great because as I understood it, there was a perfect division of labor. I would not want a computer to just make a decision about what they need to do, but when it comes down to something that requires a steady hand, complete accuracy, I want a system that is infinitely patient, that doesn't get exhausted, that frankly doesn't make mistakes when it comes to rote tasks and stitching, moving something inside my body one millimeter, that is exactly a job for a robot.

Beena: Yeah, so true. So we work with different clients across different industries and sectors, and we have seen companies that are very early in their AI journey and companies that are more advanced in their AI journey. And I am sure you see the same with AI too. What suggestions or what advice/guidance would you have for CEOs and board members of companies that are still very early in their AI journey where they have just started doing their first few POCs and put out their first AI project into production. What is your advice to companies who are just beginning their AI journey? What should they focus on? What should they thinking about?

Oren: The first thing to realize is that this is a long-term major investment. So you can't just wave a magic wand, flip a light switch, plug something into the wall, and boom you have AI flowing. I would say that it is an iterative process and one needs to have patience and investment. The second thing is the notion that we are going to use AI is, of course, no substitute for clear business thinking. So questions like what is the return on investment and when will this be ready, what exactly is the problem being solved here, where are we going to get our data from, all those sorts of basic conceptual questions and framing questions have to be answered and have to be answered by humans. The AI is not going to do it

for you. So I think that there is a lot of potential here, but at the same time, AI is not a panacea nor is it an easy journey.

Beena: Yes. What about the companies who are further ahead in their AI journey. They have AI being used in almost all their functions and their co-products. They are well ahead in their journey. What would your advice be for those companies?

Oren: First of all, I would say congratulations, you are ahead of the curve, and then I would say how do you continue to build that sustainable advantage with of course companies like Google, Amazon, Microsoft, and others moving so quickly? So I would say in a fast-moving field, if you have that AI capability, focus on what is a sustainable edge and what allows you to remain ahead of the competition. And of course, two huge components of that are (a) data, so proprietary data can be a tremendous advantage, and (b) talent, and that is talent both within the company but even more broadly. I would love to see CEO support for improving AI education in this country, improving AI immigration.

Beena: Yes. So, Oren, there is one way I use to describe this, and I would love to get your thoughts on where we are with AI in terms of the industry and applications. There is a core AI research that is still happening because the technology is not fully mature. We are still discovering new ways that AI, the next wave of deep learning or quantum, there is continuously that core technology is evolving. Then there is a second stream, which is the application of the technology, which is what we just spoke about, taking a technology that is still developing and applying it in finance, in HR, and in industries like agriculture and education. The applications of technology that is a second stream that is also still evolving, and then the third stream, the consequences of using that technology beyond the value creation. The second stream is very focused on creating value with the technology. And the third stream is what are the ethical implications, what kind of regulations need to be put in place, are there health impacts? The consequences of using the technology. So there are three parallel streams. And if you think about it, the first one is around academia and research groups. The second one is out there within industries and companies. And the third one is you need regulators, but you also need these two other cohorts to come together to solve for the third one. And you straddle very well between academia and research, and also with the industry and also with regulators. What do you think will help us move forward in an ethical, responsible way because of all these three streams that are still growing, [there is a sound here like a utensil hitting against a dish. It shows as 31:03 on my audio] what is the best way for us to move forward? How do we bring these cohorts together to have these conversations about not only value creation but the other implications of using technology in the real world.

Oren: I think that the taxonomy you described with core research, application, and analysis of consequence is really excellent. The thing that I would add is it is important to have the feedback loops, so that the people doing the applications talk to the people doing the research and bring up problems. For example, brittleness, corner cases, how do we deal with that, how do we generate explanations of the program's behavior in the real world. Likewise, the people concerned with consequences like privacy and bias and so on need to be talking to people in these earlier cohorts, earlier streams so that we can do a better job designing the algorithms, designing the applications. I would say that is number one. Number two, I think it is really important, while we want to think about this globally, and across all these different arenas, to realize that what happens in medicine, what happens with toys, what happens with cars, these are all different places where AI is being applied, is very different.

Beena: You are absolutely right. That is why I am so hopeful for 2021 on discussions of ethics. Like we have moved beyond just those headline and broad blanket statements about we need to solve for bias, to what does bias look like for this specific industry versus what does it look like for another industry. For example, bias is super important if you are in the retail world or health care, but bias is not as important when you are thinking about predicting a machine failure in a manufacturing plant. The ethical implications for those are much more different, so getting to that nuanced discussion is the next step before we do blanket level regulations. We need to be able to look even within an industry, it also depends on the use case where you are using the AI solution. The implications are going to be very different and the regulations, policies, the guidelines should be different based on where and how you are using AI, as opposed to having these broad challenges. But I do think there is progress happening. And I am hopeful, just like you, that 2021 is where we will actually make progress on solving for ethics as opposed to, we have talked a lot, so getting to that level. Oren, this has been a great discussion. I have really enjoyed our conversation today, and before I let you go, a quick question. How can people stay connected with you, where can they follow you, how do they stay on top all of the amazing work that you are doing?

Oren: Well, first of all, I want to thank you for the wonderful dialogue and great questions and the discussion. We have a newsletter that reports on the latest and greatest from the Allen Institute, AllenAI.org, you can sign up for that, and we regularly release data sets and articles, both more popular and more technical. For example, a colleague and I just published an article today in TechCrunch, literally today on GPT3 and its implications for the business world and start-ups. So, welcome engagement from your listeners and viewers, and also from you and your team.

Beena: Oren, thanks again for being with us on the show. And I want to thank our audience for tuning into AI ignition. Be sure to stay connected with the Deloitte AI Institute for more AI research and insights.

Thanks for tuning into another episode. Check out our AI Ignition page on the Deloitte AI institute website for full video and podcast episodes. And tune in next time for more thought-provoking conversations with AI leaders around the world.