



## 2023 TPI Aspen Forum AI - Separating Hype from Reality

Oren Etzioni:

As was mentioned earlier, it learns its algorithm, its approach, from data, from very extensive data sets. As such, it's much more nuanced but also much more surprising and unpredictable. A calculator will always do its job. It doesn't have hallucinations or hallucinations. Our new generative AI models really do.

So there's some important distinctions there, but the thing that I want to emphasize is that both are still tools. Some of the AI alarmism, especially the most extreme form about AI taking over the world and such, assumes that AI becomes a being with its own will, with its own initiative and so on. I don't believe that's the case.

I believe that it's a powerful tool. It may be difficult to control. Imagine a power saw or a modern car. We have issues with these cars. They pollute, they have accidents, et cetera. It's a very sophisticated and powerful tool. It's intelligent in some sense, AI, but it's not autonomous by any means.

Think of when eventually we'll have self-driving cars. They'll reduce accidents. They'll also make mistakes. But one thing they won't do is tell you where to go. You still tell the car, "Here's where I want to go." I could continue, but let me stop there.

Shane Greenstein:

In that sense, generative AI doesn't really change your views at all.

Oren Etzioni:

I think generative AI is remarkably powerful and creates a new set of threats. I actually disagreed with some of the comments earlier. It's always nice to harken back to historical examples and so on, but I actually think we are dealing with something quite different. That said, generative AI still is a tool with very, very limited autonomy, if any.

Shane Greenstein:

All right. Let's jump into policy because that's what this crowd loves. Most of us only get to fantasize what we would say if we were in front of the president and we got our three minutes to lobby. You had the privilege to do that and spend some time in roundtables organized by the White House about AI. What did you get to say in your three minutes? We all want to know.

Oren Etzioni:

It was an honor to sit with the president. He was about as far from me as you are. I was impressed that he was focused, engaged, sharp, even witty. Somebody told him that banning facial recognition is something that AOC and, I'm blocking in the name, an extreme conservative senator, both agreed on. He was like, "Help me, Father. Can there really be such a thing?" It was very funny.

I advocated for two things. One is the use of AI in assistive technology. I think there's an interesting fact that, and it was expressed by Danielle Li earlier, the weaker you are at

something, the more AI can help you. If you have trouble walking, if you have trouble seeing, say, AI systems can be a tremendous benefit. Whereas those of us who can see need to worry about that less.

He did not immediately cordon onto that and launch a moonshot project on AI and assistive technology. He indicated that he'd like to see a broader consensus than just me talking about it. Besides, other parties perhaps constrain his ability to fund.

But I was ready for that, so I immediately turned to something that doesn't cost money, but it does cost political will, which is immigration. We have a very difficult supply-side challenge with just not enough trained workers, not enough PhDs, students who come here and can't stay, et cetera, et cetera. If we even restricted it to highly skilled workers in AI, students in AI, this would be a huge boon to the industry. He told me that he'll get right to work on it, and he'll solve this in the next 15 minutes.

No. I'm obviously misrepresenting that, but those are the two things I highlighted in my three minutes with the president.

Shane Greenstein:

Did international rivalry also come up in this?

Oren Etzioni:

Yes. It did. The whole session was about 90 minutes. This was just my three minutes to speak directly about my ideas. There is a lot of concern about that. There were people there from National Security Council, chief of staff and so on. I would say that how we are competing with China on all fronts, but certainly in AI and its implications for the economy, but also for national security is top of mind.

Shane Greenstein:

There's this group of high profile scientists who wrote an open letter calling for a slowdown. That already also came up. What do you think about that? Please elaborate. Take all the 30 minutes.

Oren Etzioni:

Well, look. What would a six-month moratorium do other than, it was mentioned earlier, give a lead to our international competitors who won't abide by it, or maybe help Elon Musk, who's a signatory to that letter, catch up with his own AI efforts to the existing ones? Nothing. Nothing gets accomplished there.

The letter was vague in so many ways. The only motivation that I can see for writing something like that is that it's symbolic. It's not a specific policy proposal, but it's symbolic of the fact that we do feel like things are moving too fast. I just saw the statistic that 72% of Americans want AI to slow down, and only 8% want it to speed up. So we have this feeling.

Now, the question is what do we do about it? How do we translate it, not to policy proposals, but again, as one of the speakers earlier said, to policy proposals that actually have a positive impact? I'll leave that to you. I'm not a policy expert, but I think you would agree with me that a six-month moratorium is not the most effective way to go.

Shane Greenstein:

How about the principles? That was also brought up, but let's go into them. The White House brought these CEOs in from seven companies. It was a big press conference. They all said they'd voluntarily abide by these safe AI principles. What did you think of those principles?

Oren Etzioni:

I think they're lovely, lovely principles. Can we get to work now?

Shane Greenstein:

Any of them that you liked in particular or you didn't like in particular?

Oren Etzioni:

Here are some ... Yeah. I did like some of the work on transparency, some of the work to test something before you launch it. At the same time, I would say that we need to make some important distinctions that were not made in that statement of principles that I would add.

The first one is the distinction between AI research and AI deployment. I think research, there are many benefits to allowing research to proceed relatively unfettered. When we've tried to regulate research, like with stem cell research, it's had a very negative effect. So I would first make that distinction very, very clear in any regulations or principles.

The second one is I would make a difference between regulating broad technology and regulating applications. There are different applications that we should think about very seriously, say AI weapons that kill people without human in the loop. There are other ones that are less dramatic but just as important.

We should also figure out who's responsible in various ways. So if ChatGPT persuades me to do something wrong, or distributes instructions on how to create bioweapons, who's liable? Who's responsible? I think all these are extremely important.

Some of the other vaguer principles, I just think somebody ought to connect the dots more explicitly to outcomes. Let's say we abide by this principle, then what?

Shane Greenstein:

All right. We'll come back to the threats in a moment. Your remark about deployment makes you think about a lot of algorithms have already been deployed, and AI is being deployed now, just now as we speak, in a lot of parts of daily life. To those who know it well, you can see the AI already in a lot of the things we use.

So let's be practical. For this audience, what does everybody need to know? What does every college graduate need to know? How about every lawyer? How about every lobbyist? I'm not sure what's most common in the room today. What does everybody need to know about the things that are being deployed today?

Oren Etzioni:

I think that everybody needs to know what the technology can and cannot do, and what is a reasonable trajectory for it. Sometimes, when you point out to people that the technology does not do nearly as much as it is hyped to do, people say, "Oh yeah, but in six months or 12 months," or, "Look how rapidly it's progressed in the last 12 months." Beware of extrapolating, particularly along an exponential. We've had an exponential, but that doesn't mean that it's going to continue that way. In fact, I don't believe that it will. Certainly not in the very short term. So I think people just need to understand what it can and cannot do.

In addition to that, let me go out on a limb. When I was in college, we had a course that I helped to teach. It was a one-week course, and it was required for graduation. That was a course to teach you how to write a computer program. Curious. How many people in this room know how to write a simple computer program? Well, I'm delighted. I can assure those of you who don't, in a few hours with me or with code.org, which has wonderful tools for that, and other places as well, you can learn it. It's just not that mysterious. I would go further that you could easily learn, in another week or so, some simple uses of AI, simple abilities to program it.

I think it's really important because it informs the opinions. It's not something mysterious, ineffable that only techies could do. Everybody has the ability to do this. Even if you don't take me up on this, if you have kids, nephews, nieces, I strongly encourage you to have them learn so that they don't come to this new world with this kind of shadow-boxing feeling. There's some AI thing out there and it's going to take over. Certainly spending some time yourself, and I'm sure people have done that with ChatGPT or the equivalent, is also extremely important.

So I would say that being empirical about it, as opposed to conceptual, because conceptually, I can sell you all kinds of things. Be empirical about it, and understand what's going on.

The last thing I would add, to bring to your set of tools, is never trust an AI demo. What I mean by that is not don't trust it like check its sources. I actually means something else. I mean that sometimes it'll do seemingly amazing things, particularly if it's being controlled by somebody else whose job it is to make it seem amazing. But if you kick the tires systematically, you'll also discover some of the weaknesses. So don't just go with cherry-picked examples with the best case scenario or the worst case scenario. Understand what the boundaries are and what the playing field is.

Shane Greenstein:

I'm employed at a university that requires every undergraduate to pass a swimming test before they can graduate. So let's add to that every undergraduate should be able to write a simple program before they graduate, even the English majors.

Oren Etzioni:

It's remarkable. When I was in college a good, gosh, I was going to say 20 years ago, but it's more like 40, time flies, it was a requirement. That's why I was one of the computer people teaching the course to English majors who delayed it until right before graduation, but they needed to do it to graduate. Guess what. They all learned, every single one. Nobody didn't graduate because of this.

But they removed the requirement. As far as I know, very few, if any, colleges today have the requirement that to graduate you have to be able to write a computer program. But this lack of computer literacy is a huge problem.

Shane Greenstein:

What about a mid-career executive? Not to say that there are any of those in the room, but what would you suggest for them? Many can't take a week out of their lives to learn how to program.

Oren Etzioni:

Well, I do disagree a little bit with that premise. Imagine that you're a mid-career executive, but you were illiterate literally. You couldn't read. You've somehow made your way through there. Would you take a week off to learn to read? I would encourage that.

And it doesn't take a week. Take a day and see how far you can get with the modern tools. No excuses. Tell them it's part of your job. There's on-the-job training. They teach you these workshops with dubious merit. Yeah. You've been to some of those. Yeah. During those workshops, put it on mute and learn to program.

Shane Greenstein:

All right. All right. Let's lighten up. Let's lighten up. We have a couple questions coming up. Let's just go to it, the AI apocalypse. Go straight to the lightest topic on our list.

AI has found its way into a lot of society, military weapons, and there are malicious actors. What's happening? How do you think about that? That also has come up. You've worked in this area for a long time. How do you think about the benefits and the dangers?

Oren Etzioni:

Very related to that, I see in front of me a question, anonymous, three minutes ago. "It makes sense to worry about AI-driven weapons that can kill, but people have an awful record killing accidentally or maliciously. Would AI be worse?"

I want to make a few points here. The first one is that I do think that the right standard for judging AI isn't is AI going to be perfect or ideal, anything like that, but is how does AI compare to the status quo. Self-driving cars is a great example. We have 40,000 people dying each year in American highways using these killing machines. We can reduce that with self-driven cars. They'll still kill some people, but far less. That is the point of comparison. Are things going to be better or worse?

Now, directly to the question, I do think it makes sense to worry about AI in this context for two big reasons. The first one is that it's a power tool. You can kill people a lot more efficiently with AI-powered weapons. So I do think we have to worry about that.

Secondly, the compliance game really changes. When people build nuclear weapons, uranium enrichment, it's a lot easier to track what they're doing and so on. When people are using AI in various nefarious ways, it's in a computer somewhere, much harder to track.

AI is changing the equation in terms of cyber attacks, in terms of weapons. What that means is that right now there is an AI arms race going on. It is worrying. Again, from my point of view, the only thing more worrying than the Pentagon and our military building AI weapons is us not building them but having our adversaries build them more effectively.

Light topic.

Shane Greenstein:

Yeah, light topic. Yeah. Wow. I'm not sure where to go with that. That is very worrying. All right. Let's end what we've planned to have, a conversation. We'll go to questions. How about that? So let's end with my last question that was planned here.

You get the privilege, probably the joy of playing with a lot of frontier tools. Some of them are toys and some of them are tools. For this audience, what's your favorite new tool, new toy, and what would you recommend to everyone these days?

Oren Etzioni:

Hands down, spending time with these GPT tools, whether it's Google's or OpenAI's or other ones, is, again, completely essential, and surprisingly entertaining, and actually helpful in a number of ways.

One of the things that OpenAI released very recently is something called Code Interpreter, which helps you if you're doing coding, but also helps with data analysis. You can upload data and ask it to analyze and manipulate it in various ways. If you work with data, and I'm sure everyone here does in one way, shape or form, take Code Interpreter for a spin. You have to make sure that you click the right options to unveil it at OpenAI. The good news is you can do a Google search, and it'll tell you exactly how to do that.

Shane Greenstein:

Let's open for questions. See what's coming up. You can call.

Oren Etzioni:

Do you need a mic?

Speaker 4:

I'm coming.

Oren Etzioni:

Bravely sits in the front row.

Speaker 5:

Thanks very much. Really enjoyed this. Something that I've always been curious about is you opened this with this point about people have this fear of it because we don't know exactly what it's going to produce. It's different than the older school stuff where there's an algorithm, we know what it's going to do.

The thing I've been curious about though is I think when people hear that, they have this tendency to then animate the technology and view it as then it's some sentient being, because we don't know what it's going to do, just like I don't know what another person's going to do at any given point in time. But I don't believe that's exactly how it works.

I was wondering if you could enlighten us on what are the sources of the uncertainty. Is it because you don't know what the data will be that's training it? Is it doing some probabilistic draws from distributions when it's creating its response? What is the source of the uncertainty in what the tool is producing?

Oren Etzioni:

It's a wonderful question. I want to go back to something that was written more than 40 years ago by Herb Simon, who bridges the fields of computer science also Nobel laureate in economics. He had a metaphor or parable that was called Simon's Ant. How many people here have heard of Simon's Ant? Okay. All of you except this gentleman will learn something new today.

Herb Simon said, "What if you take an ant walking along the beach, and you drew its trajectory, its path, on graph paper," or in three dimensions if this was a more modern version of this. You would see an extremely complex and unpredictable trajectory from what is ultimately a very, very simple reactive mechanism.

That's exactly what's happening with AI. It's entirely a function of its reaction to its input, which in this case, it's the training data and the questions that it's being asked, what's called the prompts.

So you're right. It is somewhat uncertain, certainly unpredictable, but it's entirely a function of what it's fed, just like Simon's ant.

Shane Greenstein:

Following up on the question, also we've been asked anonymously, does that require a different kind of computer literacy to understand that?

Oren Etzioni:

Sophistication with the concepts of computers has always been much more important than knowing how to code in a particular language. I taught computer science for 20 years. The students would always come in and it's like, "I got to learn the particulars of the latest technology that's going to help me get a job at this place and so on. Let's get down to business." I said, "This is not a vocational school. The languages will change. Even the concepts will change. I'm going to teach you the fundamental ideas and techniques to withstand the test of time that you can then apply in different contexts."

The question was, is pseudo coding or more being able to structure a program more important than coding in a particular language? It always has been.

Shane:

I'm the other Shane at the conference. Actually, on the point of vocation, how do we get to start to democratize AI and lower it down into the community colleges and high school in the education point? I think that's part of it. All this is very interesting, and I think we need to go faster rather than slower on a lot of this, but any thoughts you have on how we really get it out there beyond the higher institutions?

Oren Etzioni:

I have a very concrete suggestion. What I described is literacy, which is different. You're saying we need to go beyond that. You're absolutely right. Particularly younger people need to go way beyond that.

I think it's a combination of improving our online resources, which is happening very rapidly, and it's the easiest to do. code.org just launched an AI sequence. Khan Academy, all these places have increasingly powerful tools.

The second thing will take longer but is essential. Improve the teacher training. Most of the teachers that I encounter don't know the basics of what they really ought to be teaching. So they gravitate and agitate and do various things to teach what they want to be teaching because otherwise they'll be out of a job. That's a real shame.

Shane Greenstein:

All right. One more. You get to pick.

Oren Etzioni:

Please. I was overruled.

Speaker 7:

This is, I believe, more than a comment. It got me thinking. Okay. Learn to code. Do you really need to learn to code when we have ChatGPT now, and you can ask ChatGPT to code

something for you? Or in the future, would you need to learn how to drive if the car is going to drive? I think, at some point, it's going to switch to how to use these tools effectively rather than fight against them.

Oren Etzioni:

That's a wonderful point. Let me just clarify what I said, because I actually agree with you 100%. Learn to code is always using existing tools. We used to code, and I'm old enough to have done this, code using what's called machine code and assembly. Very, very primitive. Now, we code using high level languages.

Learn to code. You can certainly use GPT to help you learn to code, but it's not as simple as take the problem, give it to ChatGPT, cut and paste the answer and we're done. Not at all the case. You have to test what came back. You have to do a few other things. You have to often clarify the question. So I would modify it, completely agree with you, learn to code using the most cutting edge tools available.

Shane Greenstein:

Well, I can see more questions, but you can take them during the break and ask Dr. Etzioni on your own.

That was wonderful. Thank you very much for taking the time to be with us. Look forward to seeing more.