AI and You

Transcript

Archive Interview: Michael Bowling

Episode 105

First Aired: Monday, June 20, 2022

Hi, and welcome to episode 105! Today we have something a little unusual. This is an interview I recorded in 2016 at the 29th annual conference of the Canadian Artificial Intelligence Association at the University of Victoria, British Columbia. Way back in our very first episode, I promised that I would bring you some of those, and here today is my interview with Professor Michael Bowling, of the University of Alberta. He is also a research scientists at DeepMind and a principal investigator in the Reinforcement Learning at the University of Alberta. He's been an associate editor for top publications such as the Journal of Artificial Intelligence Research, and he's been on television programs such as Scientific American Frontiers, National Geographic Today, and Discovery Channel Canada, as well appearing in the New York Times, Wired, on CBC and BBC radio, and twice in exhibits at the Smithsonian Museums in Washington, DC.

He was at the conference to give a keynote, titled "von Neumann's Dream," on the subject of AI playing games. When you think about computers and games, what game comes to mind first? How about chess? The von Neumann in the title was the legendary computing pioneer John von Neumann, and Michael quoted him: "Real life is not like [chess]. Real life consists of bluffing, of little tactics of deception, of asking yourself what is the other man going to think I mean to do. And that is what games are about in my theory."

And Michael was there to talk about AI and poker, because he is the leader of the Computer Poker Research Group, which has built some of the strongest poker playing programs in the world. In 2008, one of these programs, Polaris, defeated a team of top professional poker players in two-player, limit Texas Hold'em, becoming the first program to defeat poker pros in a meaningful competition. And what Michael had done by 2016 was create a program that was provably superior at playing heads-up limit Hold'em, as in, the best any player could do against it was draw, if that player was as good as this one, and that was mathematically provable. In the same way that you can force any game of Tic-Tac-Toe to a draw, right? And you could prove that, by drawing out all the possible move trees if necessary. Of course you can't prove a strategy for poker that way, but Michael had a smarter approach.

If you play poker, like I do, this kind of result comes as a shock. There is a huge gulf between a board game like chess and poker. Chess, checkers, and Go are what's called Perfect Information games in that each player can see everything the other players can see. Not so in poker, because players have hidden cards, or "hole" cards. And winning at poker is all about knowing how to bluff, because on average, you won't get any better hands than anyone else, so the only way you'll win is by betting in such a pattern that the other players think you have better or worse hands than you really do. And that's all psychology, right? You poker players know what I mean; you're following how the other people are playing, you're looking to see if they're consistently conservative, or reckless, whether there's a pattern to their bluffing, and anything you can figure out from how they behave at the table.

A computer should be useless at all of that, right? How are you going to code the ability to read players' state of mind? Although the computer is going to be very good at computing probabilities and not making emotional decisions, which are the other traits you need to be a good poker player. But I didn't think that a computer would be able to learn the bluffing strategies to win at poker. And yet, it has. Don't play against one if money is at stake.

Now, I've got to tell you, everything Michael has to say is as useful and interesting as it was when we recorded this. On the other hand, I was just starting out and I was pretty green behind the ears about how AI worked and where to pay the most attention to it. So I cringe when I listen to my questions. Forgive me, okay? I did rerecord them, though, because in the original my microphone wasn't working well and you should get the best audio I can give you. But I left the words the way they were. So let's get to the interview with Michael Bowling.

And your specialty is artificial intelligence.

Specifically, I do a lot with machine learning and game theory.

And you're at a conference to present your poker playing bot.

That's right. So, a presentation I just gave today was on some work we had done a year ago, where we had shown for the first time that we could solve the game of Heads-Up Limit Texas Holder. So, it's a variant of poker, it happens to be the smallest variant of poker that humans play. But we have a perfect strategy. So, a strategy that cannot lose, no matter what it reveals to somebody else, it doesn't matter how much they play against it, no matter how much they figure it out, it's guaranteed that it can't lose money.

So, now poker goes the way of checkers and chess.

Yeah, it's a long time coming in the sense that this is actually solving the game. Whereas some of the other milestones we often think about, particularly in games, that there's different milestones, there's a point where we can play competitively with humans or surpass human play, or we're able to beat them and so in checkers that happened in 1994, and poker for again, the variant of Heads-Up Limit Texas Holder that happened in 2008. But then you could also solve the game and so checkers was solved in 2007. This is what we did for poker in 2015 and so go, for example, we've just now crossed the threshold where we're able to beat top human players in the game of Go, but we're never going to, that's out of the realm of mathematical possibility that we would be able to solve a game, the size of Go.

Although if you can win all the time anyway, it doesn't really matter?

Well, it kind of comes down to you can't guarantee it right, maybe humans can continue to get better, and maybe the program don't and so that's one of the, when you really solve the game, it's kind of special, because what you're actually saying is, it doesn't matter. Like this isn't about how good your competition is, this is just establishing that for all time now forevermore, humans, at best can just tie the game, because they can only reach a perfect solution themselves.

So, those are relatively narrow applications of AI, and there's a lot of attention lately being given to general artificial intelligence. We've got statements by Elon Musk, Stephen Hawking,

Bill Gates, saying we need to be very concerned about where this will end up. Contentious, perhaps. But in what sense does that impact you?

Well, in some sense, the poker work that I presented on today isn't all of the projects that I work on, and some of them are actually really trying to get much closer to that sort of, I don't want to say we're getting closer to general intelligence, but we're certainly trying to build an AI that has a more I like to use the phrase general competency. We've seen over the years that if you pick a particular problem, we can get computers to achieve a very high-level performance on it. But what happens if we want to have computers be able to show a more broad range of competency and how do we even begin to evaluate that I think that's where things get really interesting and so one of the projects that I started about four years ago now was the arcade learning environment of using video games, being that there are so many of them made. So, we started with the Atari 2600. There're over 500 games made for the 2600 and let's see if we could have an agent that could play any one of those but just sitting down, never seen the game before and learning to play the game. Now, I think this is an interesting problem and it's very different than the problem of how do we get to like super high level of Go play? How do we get to super high level poker play? Because to be able to do this, you really have to cover the variety of situations that people have made. There's sports games, there's puzzle games, there's chess built into the Atari 2600. So, if you really could do that really well, then you've certainly done something categorically different than what we've done before.

Does that mean that one day Halo 2 or World of Warcraft fall under the scythe of computers?

Well, I think the day will come where that happens, even before we see a more general competency system get there. Right now, I think if an AI group wanted to tackle that problem that became a focal point of the community, I think that you would see an individual game fall rather quickly. But how do we move to the point where we're really able to have a system that we don't have to engineer that particular game and learn the intricacies of that particular game, but the system itself can divine that and mostly, we're going to get that from learning and so this is why machine learning is a big part of what I do, is I want to see our machines be able to figure much of this out for itself. Now, I bring that up, because you asked about sort of these fearful statements, and I have trouble with them in the sense that I don't fully know what perspective they're coming at it from because on the level of in the future, we're going to want to be careful about this, I'm totally on board and it's useful to have the conversation now because it's much harder to have the conversation like after when we're arriving at the point the conversation needs to be had, it would be much better off if we've already been used to having those conversations. And so I think on that level, I think most of the AI community is in agreement. On the level that oh crap, what you're building in Atari is going to somehow escape Atari and run amok in the real world or the virtual world or any world at all is just seemingly utterly nonsense, right? Like my agents make chickens cross roads and they're not going to suddenly think of anything else and it's hard to imagine how there's a danger in that.

Nick Bostrom speculates on a scenario where you don't have to; that an intelligence whose mandate is to maximize paperclip production, can if it becomes sufficiently powerful and unscrupulous consume the earth and everything on it to make paperclips.

Sure… I don't know, I feel like I mean, the scenarios are nice for science fiction storytelling, but the step between the technology we have, and the fear of these science fiction storytelling elements, it's not even just that it's not here; it's totally unclear what would be the steps between here and there. That seems like a chasm that is so far-fetched that our systems right now don't look like that. They look like toasters. And even as sophisticated as they've gotten, even the ability for AlphaGo to learn to play Go, just through millions and millions of games against itself, our poker program learning to do that in poker, it still just looks like –

The only thing it could ever output as a slice of toast.

That's right and it's not going to suddenly decide to have some other impact on the world until we set it up to do that. So, I think there is a conversation that I don't want to say it's a fearful conversation, I think there is an important conversation to be had sooner than even the maybe the Nick Bostrom one, which is what are we using our tools to do and I think that conversation is the one human should be having and have been having for thousands of years, there's nothing, should we be using the ability to make explosive powder into the bullets that we propel very very quickly and kill people? Like that's an important question to ask the actual building of guns and I think that there's no reason to believe that AI is in a somehow a different category than that one. If we choose to hook our AI systems up to our factories, then we could have the problem that when we program to make paperclips that they might go and have much more effect than we want.

Might happen if one factory does it, gets a tenfold increase in productivity, than the rest will do it.

So, we do; I think we do have to think about how we use our tools. That is a conversation that we're used to having we happen a lot and we can should continue.

And perhaps in the context of artificial intelligence. Is there a conversation happening about ethics in the sense that we need to have more of that conversation?

Yeah, I think I mean, in some sense, the doom and gloom speakers, I'm not really sure how doom and gloom they are. I mean, like Elon Musk is one of them and yet, he's a big investor in not just AI and ethics, but he's a big investor in just straight up AI, and pushing the AI agenda as much as anybody else. So, he can't be that much of an AI Luddite about what's actually happening. And so I don't want to put too much on him, but I think making those statements has certainly driven AI researchers to say, okay, you're right, there is value in starting that conversation. I think the problem is when it's put out there in the public in that way, I think it does, I mean, the media in particular likes to, the cool story is when it sounds more doom and gloom than it even is and I do fear that this can change funding situations and I don't just mean selfishly, doesn't entirely matter to me. But I think that that can affect our progress and I think there is a lot of good that can come out of this and right now, I think what everyone sees is far more good than fear. But yet, if we change the public conversation, we lose the potentially the opportunity for all the good. So, I think there was value in saying, let's have that conversation as researchers, because there's no reason not to. I think I mean; they continue to push the doom and gloom agenda of what the dangers of this could be. I think that to me, that's dangerous. I think

there's a number of opportunities for which if AI research were to slow, we would be losing out on chances to improve people's lives on every level, from convenience, to life giving ability.

The upside is as big as the downside.

That's right.

Do you perceive that artificial intelligence is getting not just the doom and gloom, but in general, a lot more attention has it taken off?

On some level. I mean, obviously, the circles Iran in it feels like lots of stuff is happening. There's a number of like milestones being passed that we thought were decades away, and they're falling quicker than we expected them to. On the other hand, when I've talked to people who were in the general public, and I asked them, have you heard of AlphaGo? Have you ever been following the Go matches? The answers are, "What are you talking about?"

And in some circles, though it is. Gartner came out with a report about what the impact will be like two years from now and with some large numbers about the impact on jobs and other capabilities.

So, it feels like in the circles that I typically run in that, that AI is both hot, and it's exciting, and there's a lot of stuff happening. And yet at the same time, I think if you actually were to pull the mass public, I think it's a smaller blip in the whole perspective.

They're typically behind the curve. We've had this thing more than once in the past have an AI wait into following overblown hype. Are we now at the beginning of a summer? Or are we at the crest of a summer, waiting for the bubble to burst?

Well, I think when we talk about the AI Winter, we don't feel like the summer was real, right? It's sort of like suggesting that there was some smoke and mirrors that were causing us to believe more was happening. I think this is mostly a summer; I don't think this is a bubble that's going to burst. But at the same time, I think, inevitably, hype can be both real and hype at the same time. Right? The reason that people are super excited is because it's real, but it doesn't mean that we fully understand what's actually going to happen and I think there will be some realization that, I mean, this is in particular to say, deep learning, that this isn't going to answer the intelligence problems, that fundamentally, there are still many questions, many of the exact same questions we were asking, we've already been asking, are going to manifest themselves, we just now have another hammer in our toolbox to address them and so I think we're going to realize that it's real, and simultaneously start to understand that it doesn't solve everything, and there's still work to be done.

Do you foresee the day when you have a conversation with an artificial intelligence that's as interesting as this one?

I'm guessing, not in my lifetime, but I actually kind of feel like I'm an AI pessimist, I'm actually working on the problems I think can't be done. So, it's not surprising that I think maybe we're not going to get to that point, I think how we start to think about our machines is going to be very

interesting. I mean, this feels like a topic for a sociologist and not a computer scientist. But I think it's going to be very interesting how we conceive of our machines, as we begin to interact with them in ways that we haven't previously.

Speaking of not computer science, one last question, outside of computer science, what do you think offers the most opportunity for cross fertilization that would improve or accelerate AI research in interesting ways?

That is a very interesting question. I don't know the answer to that question. I mean, historically, people have drawn a lot of inspiration from, say, biologically plausible suggestions for how AI systems might work and I've never found that useful way to think about things. I mean, if we all need a place to get ideas from and so in that regard, I don't think those ideas are inherently bad. But I don't think that's a benchmark by which we're really going to be able to push forward on I think, we are going to discover general principles of intelligence, but I don't think the way we're going to implement them is going to match the way biologically we've arrived at those same skills. So, that's never been a useful place for mining ideas for myself. I'm not sure I know the answer to that I often get most of where I get ideas from are trying to, when you make a problem concrete, you start to realize what are the pieces that are missing and so often, I need good concrete problems that aren't so far-fetched that the chasm is too big to make the next steps and so I feel there are these moments when things crystallize around the problem of GO or poker, that's what they do. I think they crystallize sort of this is exactly what we're missing and we can begin to think about what would a system look like that could do that. So, my go to is almost always games. I think that there's some obscure games that most people don't know of that I think are really fascinating, that I've been thinking a lot about how could an AI system ever sort of do that, I'll give an example that most people know better. So, the game I'm thinking of is a game called Hanabi. But most people haven't heard of it. But I'll take bridge for a sec. So, we can actually build a bridge program that can play the cards of bridge very effectively. So, here's a different question. Suppose you were to pitch the game of bridge to a computer and say, devise me a bidding strategy? How are you going to communicate? Because bidding isn't about trying to gain utility or get into the best position, it's about me communicating with you what's in my hand and so building up a communication protocol is effectively what a bidding strategy is, the people who looked at bridge previously, they mostly, it's partially how bridge works, you don't get to really invent a bidding strategy, you have to write it down and show your opponent's there's these really strict rules and so when people have looked at bridge doing computers for bridge, they almost are always restricted by the existence of bidding strategy. So, they have to play within that space, and they end up building rules on top of it and what they're doing isn't the problem, I think is interesting, which is the problem is the communication problem. How can you actually build communication protocol and the kind of problem when you distil it down to its very core isn't one that looks like any of the search base problems that we've seen before it looks really challenging. So, this isn't really pulling an idea out from another field. But I give that as an example of where games can help us think about what might be missing in our intelligence capabilities.

And going in the other direction. You mentioned dating, anyone on Wall Street would start listening in another context and with game theory and bidding. You're speaking their language, are you not?

Yeah, I think moving out of strictly competitive games. So, even with the poker work, which I often contrast with chess and checkers and go because poker is a game of imperfect information. There's another way in which all those games are the same, which is that I am trying to be you. This is a zero-sum game and my actions are like, if I benefit me, it's directly hindered you. I think there's a whole slew of games for which that isn't the particular setup and many of the games that we sit around we don't operate that way. That isn't how we generally approach someone else is not to assume that if I take an action, that's good for me, it's going to be bad for you and so I think there are opponents and bridge but the part of it that I'm thinking about is the part where no we're on the same side here and I think that is historically we don't have a lot of experience in those kinds of games and so teamwork based games, like video games that involve lots of people working together, whether it be like Counter Strike, or World of Warcraft, or any of these games, where you really work together to achieve some sort of goal. How do we begin to build systems that can play in those spaces, I think is also I think it'd be an important step forward and you can also imagine, what if one of the other people you're playing with isn't another computer agent? What if it's another human age? How do we actually be able to do that, and there has been some work that I started doing some of this a while ago, and more people have been pushing this idea with something that's called impromptu teams is what I call it. It's also been called pickup games and there's another name for this, it's eluding me but this idea that could you build an agent that can jump on any team arbitrarily, and can figure out how to contribute onto that team and this obviously, feels like this would be great from a societal perspective that we would love to have our agents to think that way and those problems turn out to be much harder than you would initially think you would think. But we're all on the same team. This should be easy. But I guess maybe we do know this. We do know, when we get a bunch of people together with a common goal. Sometimes we're less efficient than we would be if we just did it ourselves.

Yes, with a committee. I did want to mention what you brought up in your talk about the application of your work to medicine, and how you can apply the same strategy to diagnosis.

Yeah, so the point is that doctors when they're making a medical treatment decision, one of the things they have to deal with is uncertainty. They have uncertainty and knowing exactly what the patient is suffering from exactly how this particular patient with this particular DNA plus medical history is going to respond to some treatment, how the person themselves might choose to take the treatment, they might not even follow the advice that was given to them. All of those are uncertainties and once you sort of described those uncertainties in sort of the form of a game as being I have a series of decisions, and there are outcomes that I can't predict, and there are other agents in the world taking action. That's a game and the kind of game it is turns out to be an imperfect information game and so we had this paper where we showed that in fact, making robust decisions making robust policies, exactly corresponds to an imperfect information game and the form of it was one that our techniques can exactly handle. So, I feel like the idea of poker

was to say, here's a game that is all about uncertainty and if we can advance AI in that game, then we're going to see applications in lots of places and this is just one of many examples of that.

That's great. Thank you very much.

That's the end of the interview. In terms of what might be different now, I think the only thing is that people are having more interesting conversations with GPT-3 than we thought was imminent at the time, although even GPT-3 isn't capable of coming anywhere near the depth and interconnectedness of the conversation you just heard. As far as today's state of the art in poker-playing AI, the Pluribus bot by Carnegie Mellon and Facebook plays six-player no-limit Hold'em. Expert poker player Jason Les said he felt "very hopeless. You don't feel like there's anything you can do to win." And it will do things like bet two or three times the pot, which is gutsier than most human players. Even while they didn't release the source code, you can hardly expect other people not to replicate that eventually. So I wouldn't want to play poker online for money.

In today's news ripped from the headlines about AI, researchers at UC Berkeley, UC San Diego, and Google have shown how they can turn a few 2-D images of a scene into a 3-D image, in order words, turn pixels into voxels. They call this Neural Radiance Fields, which may be just so they could abbreviate it to NeRF. The videos are just, wow. They need just a few dozen still photos of a scene from known angles to be able to render a 3-D scene within milliseconds from which they can generate a video that looks flawless. They say it's a small neural network that reconstructs the scene by predicting the color of light radiating in any direction, and even works around objects blocking the view. Search for Nerf AI and you'll find it and the videos.

Next week, I'll be talking with Ben Goertzel, AI researcher and founder of SingularityNET. Ben was on our 2022 prediction panel, but it was clear that we needed to talk with him solo about artificial general intelligence, and that is a mind-blowing experience. 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