

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 1 UNIT 2
Casebook Video 3 Transcript

MIT AI M1 U2 Casebook Video 3 Transcript

TOM MALONE: So, Patrick, you just gave us a great overview of the history of artificial intelligence. Let's talk now about the future of artificial intelligence. People talk about human-level AI or general AI. How far away do you think those things might be?

PATRICK WINSTON: I can't answer the question, Tom, because, you know, it's hard to make predictions in the absence of a plan. And we don't have a plan.

We don't have a plan because it's several breakthroughs away. You can make a plan if it's building a bridge or some other kind of engineering project, but you can't make a plan in science. You have to wait for those breakthroughs to happen. Sometimes they happen sooner than you think; sometimes they happen later than you think. But it's all intuition. It's extremely difficult to talk about it.

TOM MALONE: So, it sounds like you'd say that the people who, kind of, believe that, kind of, human-level AI is going to be here in a few years or ten or twenty years are probably being overly optimistic.

PATRICK WINSTON: I'm not sure if they're overly optimistic or not. I can give you both ends of that spectrum. Sometimes I wonder. Sometimes I despair and think, "Well, will it ever happen?"

You know, I once found myself taking care of an orphaned raccoon. And this was an amazingly intelligent animal. Within a day or two it learnt how to open the refrigerator. Now, we found ourselves playing tricks on each other. For a long time, Miss Raccoon would only eat chicken wings. And I wanted her to eat hotdogs because they were a little easier to work with.

Finally, I hit upon the idea of sticking a chicken bone through a hotdog and putting it in the garbage can. And she loved it. But, you know, I looked at that raccoon and said, "Wow, she's really smart, but she's not smart enough to make a machine that's as smart as a raccoon." So, I think we humans have a lot of hubris when we think that we're smart enough to make a machine that's as smart as we are.

On the other side, you know, a lot of things are happening. A lot of people are in the field. There are surely some geniuses. And maybe they will make those kinds of breakthroughs that we don't expect. Maybe where we are is where biology was in 1950. My biology friends tell me that in those days we had antibiotics. We had much understanding of how to take care of sick people, but we didn't really understand how things were inherited.

And some people despaired of the idea that it would ever happen, but then along came Watson and Crick. We figured out the structure of DNA. And everything changed. That could happen tomorrow in artificial intelligence too. So, I see both ends of the spectrum. And it's hard for me to sign up to any kind of prediction in terms of how long it will take.

TOM MALONE: So, I find it interesting that when people have done studies of the surveys people do about how fast people think, general artificial intelligence or human-level AI will occur. People, typically – and this is both experts and non-experts – they typically predict that it will be about fifteen to twenty-five years in the future. And that's what they've been

saying since the 1950s. So, in other words, people have been predicting general AI is about twenty years in the future for the last sixty years.

PATRICK WINSTON: Yeah, I might be one of those people actually. But again, it's so hard to say how long something is going to take. It's easy to say how long it's, you know, what's going to happen, but it's so hard to predict when it's going to happen. We really don't understand a lot about ourselves. We don't know why we are the symbolic species. We have an inner language. We use an outer language to communicate.

We can think in ways that we just don't understand. And so, it's very hard to make a prediction about how soon we're going to understand it when we don't have, you know, when currently we don't have any idea.

TOM MALONE: So, it sounds like you wouldn't absolutely rule out the possibility that would happen quickly, but you would certainly not be convinced that it definitely will.

PATRICK WINSTON: I can't rule it out, but – and I am an optimistic person – I think, we'll figure it out in my lifetime. But I'd like to add that my mother lived to a hundred-and-two years of age. So, I have a lot of time left on that prediction.

TOM MALONE: So, Patrick, another thing I liked about your presentation was how you talked about we need many different kinds of AI, no single one is the answer. As I was listening to you say that I thought, "Well, maybe we need not only different kinds of artificial intelligence, but we also need different kinds of human intelligence."

PATRICK WINSTON: Absolutely, yes. We need partnerships between people and different kinds of people. That's certainly true. You know, there are so many things that we humans do so flexibly. And it's not just our cognitive abilities, it's also our visual intelligence. We see things; we see things differently from dogs, I think, because we can direct our visual apparatus to do stuff at the command of our cognitive and reasoning parts.

I was once, when I was doing some work on computer vision, I decided I would take a course in painting. And so, I signed up. And there I was taking a painting course. And one day the instructor said, "Let's go outside and paint a tree." And I looked at the tree like I'd never seen a tree before. You know, I thought I'd pull out the brown crayon or something like I did when I was a kid, or something analogous to that, and that would be the tree trunk.

But I suddenly discovered that having been directed to look at a tree it was full of grey and purple and white and black. It was very different from what I expected. And I was doing that because I was commanding my visual system to do something that it wouldn't ordinarily have done. So, there's a lot of magic in the interaction of our cognitive parts, our more symbolic parts with the parts that are not so symbolic, not so perceptual and not so unique to human intelligence. And that interaction is extraordinarily important and at the same time very difficult to understand.

TOM MALONE: I guess many people would argue that inside a human brain there are actually many different kinds of intelligence. As you said in your presentation, inside future computer programs there may be many kinds of artificial intelligence. So, perhaps the real future of intelligence on our planet will be combinations of human, many kinds of human

intelligence, many kinds of computer intelligence all creating a new kind of collective intelligence.

PATRICK WINSTON: As far as anyone can see into the future, I think, that's absolutely true. There are so many kinds of intelligence that it will take a long time to understand them in anything like a complete way. So, in the meantime, the most important applications, I think, will be partnerships between computers doing what they do best and people doing what computers can't be made to do at the given state of the art.

You know, there'll be the kind of things that we do when we solve problems, when we do visual problem solving. When things are not steady state. When, you know, the statistics are, as they, to use the technical term, when the statistics aren't stationary, that's when people have to come in and use the kind of intelligence that's unique to our species. Thinking about precedents, thinking about similarities between this situation and that situation, those are all things that remain mysterious from an AI point of view.

TOM MALONE: I think one other interesting thing that happens when you start thinking in this way is that you realize that human intelligence is just one of many possible intelligences and that one of the most interesting things that computers may do is help us understand some of the many other possible kinds of intelligences that there can be in the universe.

PATRICK WINSTON: Yeah, it takes me back to the discussion I had in my segment about suitcase words. Intelligence is a suitcase word because it's a label for many different kinds of information processing capability. And we know about some things that are unique to our human species. We know about some things that computers have been made to do better than us, fast multiplication, for example.

And no doubt there will be other kinds of things that we humans can't do so well and machines don't do very well today that may be, eventually, part of this collection of things that we label as intelligent.

TOM MALONE: Okay. So, it sounds like there are a lot of things for us to look forward to, a lot of interesting possibilities for new kinds of combined human and computer intelligence. And I think we also are concluding that there's no reason to believe that it's likely that computers will have human-level intelligence, whatever that means in the next few years.

PATRICK WINSTON: Not in the next few years. And, you know, just as collections of people are smarter than any individual, I think, collections of people and computers will be smarter than the computers and smarter than the people. And it's those collections and partnerships that, I think, are where the most interesting applications lie because it's out there in those partnerships that, I think, we'll be able to make systems that can do things that can't be done today. And those kinds of things will be, I think, very important.

And just to give you one example, you know, in my own world I think a great deal about story understanding and story processing. And I dream of a day when we can make systems for lawyers, and diplomats, and doctors, and business people, all kinds of people, that would help them to understand situations, help them to understand where they might be making blunders because of their greater ability to find precedents and apply them in new situations.

That would be almost like a spreadsheet for a financial analyst. You know, you say, “Here’s some parameters. I’ll adjust them and see what happens.” You could do the same thing, I think, in the, you know, in the medium term with systems that understand our story, understanding processing, and can help us to make ourselves smarter by showing us the kinds of things we might want to be thinking about.

TOM MALONE: It sounds great. And we’ll hear more about some of your story understanding work in a future module of this course.

PATRICK WINSTON: That’s right.

TOM MALONE: Thank you very much, Patrick.

PATRICK WINSTON: Okay, you’re welcome.

TOM MALONE: Did you understand all the concepts covered in this video? If you’d like to go over any of the sections again, please click the relevant button.