

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 3 UNIT 2
Casebook Video 2 Transcript

Module 3 Unit 2 Casebook Video 2 Transcript

THOMAS MALONE (TM): So, I'm here today with Frank Levy, an Emeritus Professor of Urban Economics here at MIT. Nice to have you with us, Frank.

FRANK LEVY (FL): Nice to be here, Tom.

TM: So, I want to talk today with you about what AI is good for and what it's not so good for in natural language processing. So, I know you've studied at least two areas, one is call centers and the other is the legal field. Let's talk about what AI is good for and what it's not good for in those two areas.

FL: Well let's begin with call centers. If you think about what a call center is trying to do, if a customer calls for information they want to leave the customer satisfied, they want to give them an answer that will leave them satisfied and coming back to the company. If they're calling to make an order, then you'd like to encourage them to do as big an order as possible. On the other hand, you want to do that in a way that minimizes cost or at least keeps cost down. You're balancing that tradeoff back and forth. And what artificial intelligence can do is help you make that operation more efficient, which is to say you're able to get a given level of performance at lower cost.

TM: Good. And so how can AI do that? It sounds like a good thing to do. How does AI help you do that?

FL: AI can help you do that by in essence doing triage on callers' questions. It determines which questions can be answered satisfactorily with an automated response and if a question can't be answered with an automated response, it determines which operator ought to get the question and what other aides the operator ought to have on their screen that can help them get to a better answer faster.

TM: Okay, that's great. So, it's not just trying to do it all itself; a big part of what the AI may be doing is just helping to find the right person to do the task or to answer the question.

FL: That's right.

TM: And can you give us, let me just dig a little more deeply here into how the AI does that, how does the AI recognize that it has an automated answer for this question and for that question that should go to Susan who really knows that topic? How does it figure that out?

FL: One way to think about this is to think about the fact that there are really three components operating here. The first one is automatic speech recognition, which is the process of understanding the words the customer is saying. The second piece is natural language processing, which is to extract the meaning from those words. What is it that the customer really wants or wants to know? And then the third part is information retrieval, which begins to get at the heart of your question. So, to think about information retrieval think about your own experience with Google Search. If you're asking about a date, you

know, what was the last day of World War II? Or if you're thinking about how big something, how many storeys in the Empire State Building or things like that, if you're thinking about factoids, Google Search is pure genius. It can get that answer for you right away. On the other hand, if you're thinking about how do I get rust off the bumper of a car, Google Search is going to give you 3,800 links about that, and a caller doesn't want 3,800 links. And so, the first dividing line is whether this is a kind of factoid question or whether this is a more complicated question that you just can't get from a database. The one caveat for that is that if we're talking, say, about an electric utility where a lot of people really have the same question, which may have a kind of longish answer, but it's always the same answer, then you can take that answer and turn it into a factoid. And so there what the natural language processing is doing is, is this person really asking that question for which I already have a pre-canned answer and I'll just give that to them?

TM: And for that, the things it can do well are the things that it's seen many times before.

FL: That's exactly right. That's exactly right. The whole idea with any kind of artificial intelligence is that it's responding to situations, to contingencies for which it's already been trained. And if you throw a totally new contingency at it, it's really hit and miss. I mean I'll give you an example. Several years ago, for a paper I was writing, I asked my iPhone, I asked Siri, "Can a dog jump over a house?" Now, I have a four-year-old grandson who could give you a good answer to that question. What Siri did was to say, "Here are a list of kennels that are located near your home. One of these ought to be pretty good." Well, it hadn't been really trained on that question, so it couldn't make the link, even though it understood the words I was saying, it couldn't make the link between what I was asking and something in its database, and that only happens with training and it hadn't been trained on that.

TM: So, it knew it was a question about dogs and it knew to give you something about dogs, but it didn't know, it didn't understand your question.

FL: Right, it was dogs and houses and that's what made the link to kennels, and so that's right, that's what it did.

TM: So those are some of the things that AI can do well in call center operations. What are some of the things that are hard?

FL: Harder questions that require complicated answers that it hadn't seen before. If you step back for a second, what's going on under the hood is that the software is producing a number of answers, each one with a probability, its best estimate of we think this is the correct answer to the question with probability 0.9, 0.83, 0.72. And it's the operator of the call center, the corporation running the call center who sets a cut-off and says okay, if the probability of this answer being correct is above 0.9 then give it to the machine. If it's below 0.9 then we've got to send it to an operator.

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