MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

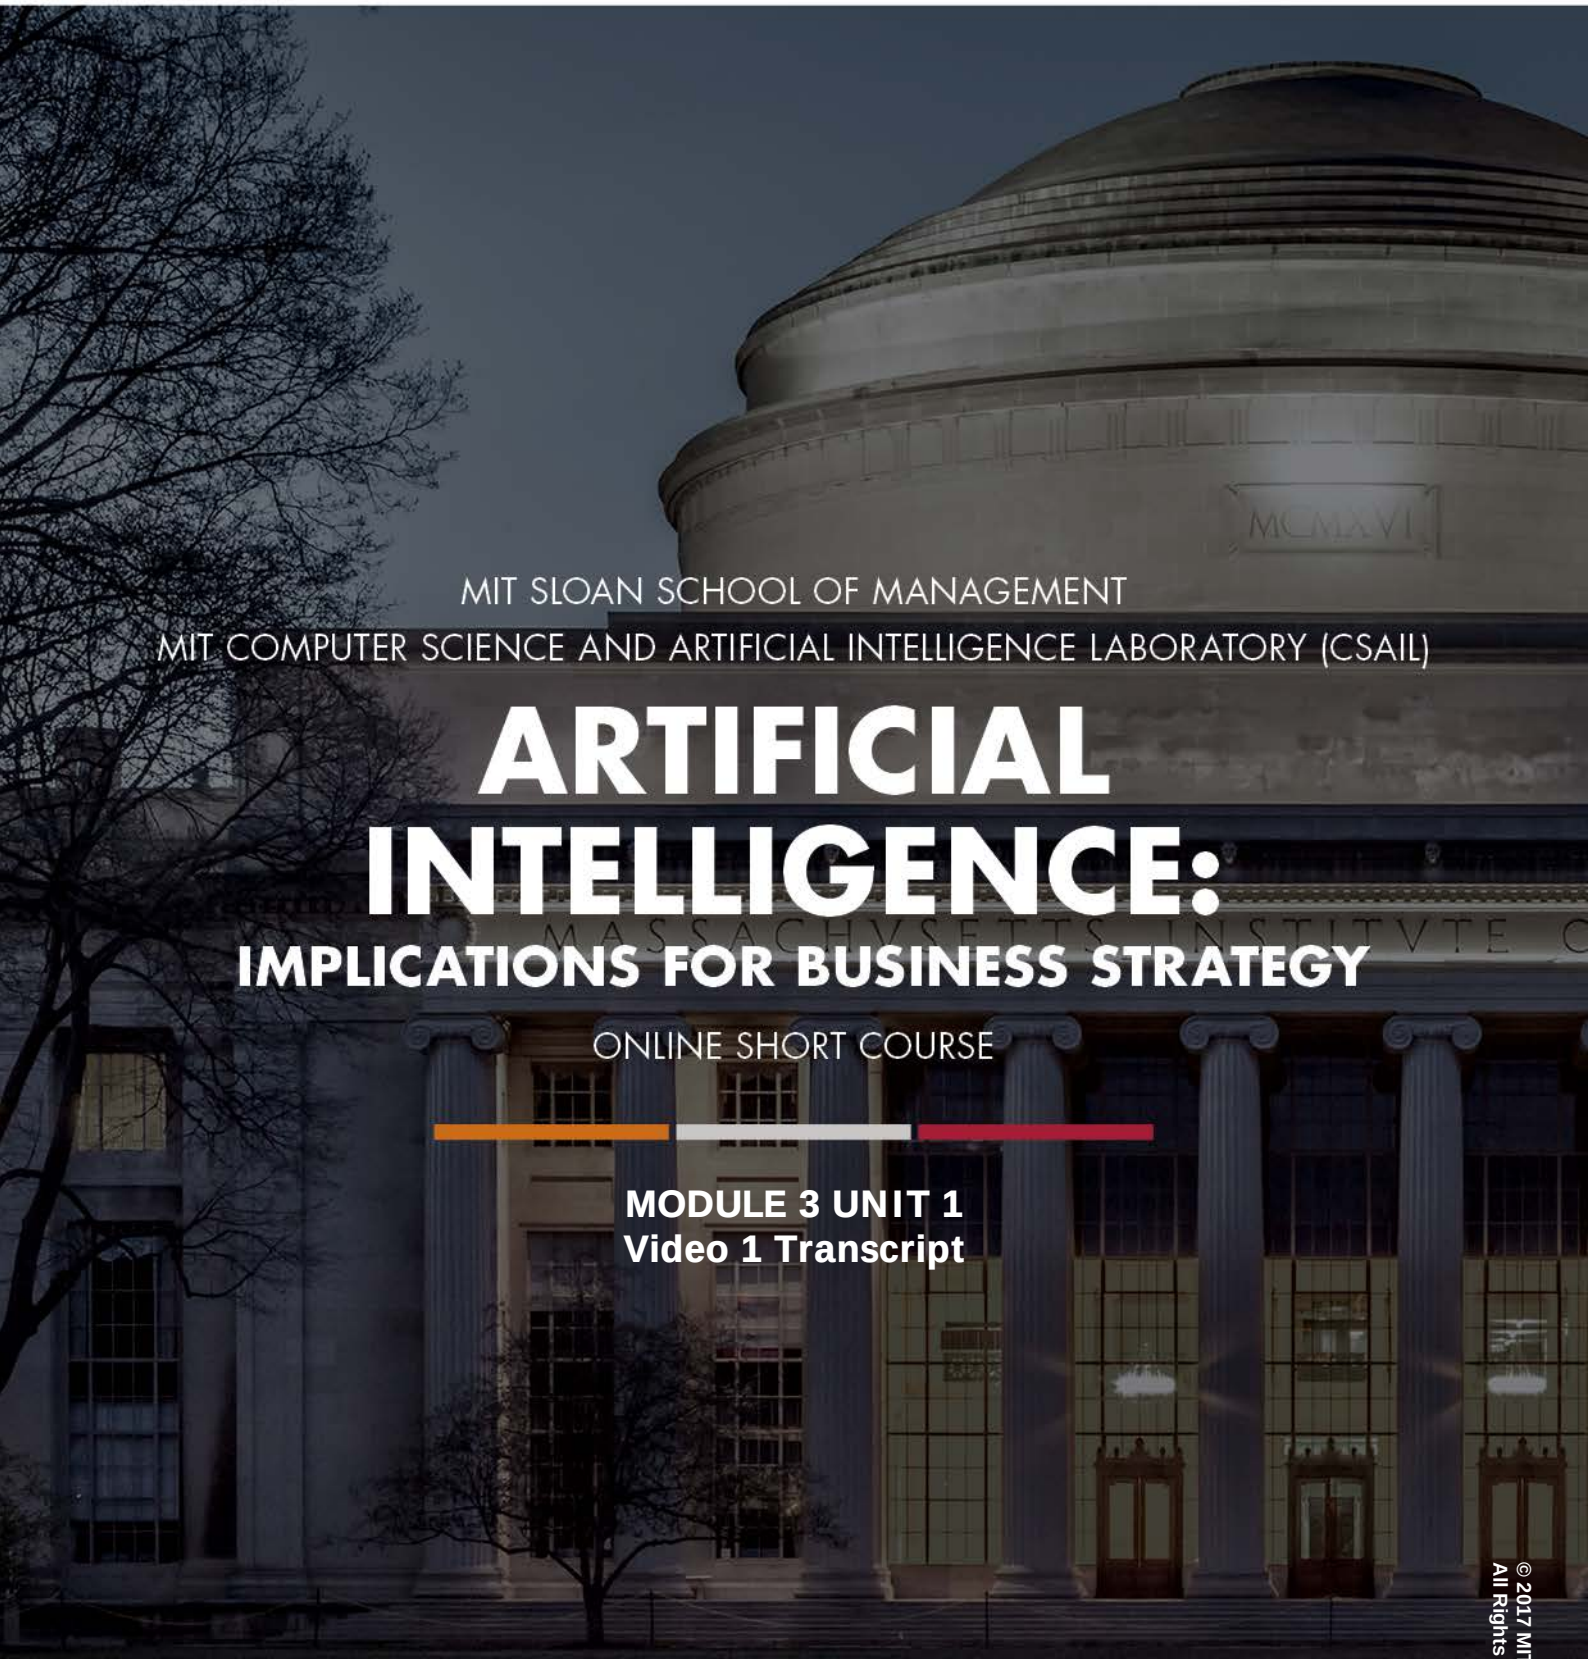

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 3 UNIT 1
Video 1 Transcript

DELIVERED IN COLLABORATION WITH getsmarter

Module 3 Unit 1 Video 1 Transcript

THOMAS MALONE: Since very early in their development, computers have been able to understand highly structured languages, like computer programming languages, such as C, Python, and so forth, and the commands used in spreadsheets. But huge amounts of human activity don't use these highly structured languages. They use natural languages, like English, Spanish, and Chinese, and it's surprisingly hard to get computers to understand and generate these unrestricted natural languages with anything like the flexibility that any normal human five-year-old can. The part of AI that's focused on getting computers to do these things better is called natural language processing, and Professor Regina Barzilay is here to give you an introduction to that field.

REGINA BARZILAY: Today I will introduce you to natural language processing, and the goal of my presentation is to teach you what machines can do today and what they cannot do in terms of understanding of natural language processing, and to provide you some intuition in how machines today can understand natural language. When we are thinking about natural language understanding, there is always a question: what does it mean to understand? Do we mean to understand like the dog on the picture where you can just pick up some strings, or do you mean understanding the same way as humans can understand? When we are going back to the history of artificial intelligence, from the early days, the thought was the machines should be understanding language the same way as humans do, for instance, Alan Turing, one of the fathers of current computer science thought that, you know, in 1950, that in about 50 years we are going to have computers that can understand you, and that communicate with you in such a way that another human wouldn't know whether it is a computer that communicates with them, and people were trying to build these kind of conversational devices, and I would like to mention one of these devices called Eliza that was developed here at MIT, where they developed the first machine where people can kind of talk to their machine.

Here you can see an example of a dialogue, where the young woman says men are all alike, and then Eliza, which is a machine, say in what way? And then they continue this kind of conversation, and the story goes that when this program was introduced at MIT, it became addictive to the secretaries, and they had to stop using it. But clearly in this case, the machine didn't really understand the language, it was just using some expression to continue the dialogue. For humans who wanted to share it was enough to continue moving forward. So, let me just step up and look at the whole landscape of natural language processing technology, what actually works and what doesn't. So, there are some tasks related to natural language processing that today most people consider solved. For instance, spam detection that you are using in your mail server, or programs which can identify named entities, the names of companies and people and locations, those are very simple tasks.

So, let me tell you about a task where we are making good progress, and can actually deliver a robust performance. For instance, a program that can do sentiment analysis, which can take a review and say whether it has positive sentiment or negative sentiment. Or, for instance, machine translation. Millions of people today are using Google machine translation, and if you don't know another language, this is the best you can do. So, the tasks which are still very hard – it's, for instance, the dialogue task, like the one that I showed you, or question answering, when you want to ask questions to the machine, and

if these questions are nontrivial it is hard for the machine to answer them. Now, there is a big confusion in the general audience in what exactly natural language processing can do.

I would like to show you some cases, which would enable you to see how far this technology is actually developed. So the first case is actually we will start with a task that many of you have done in high school, where you try giving it a sentence to decide, you know, what is a verb and subject and object, the reason we care about this task is because it's like a basic step of understanding the structure of the sentence, so machine translation system and other big application need to understand this information, and as you can see here on the slide, for English, we can do this task with very high accuracy, which actually approaches human performance, and even in a complicated task such as Arabic you already get quite close, it's above 80%.

So, I would say that syntactic parsing is very close to a solved problem. So, let's look at real application. For instance, machine translation. I selected intentionally a difficult case of machine translation where you are translating from Hebrew, which has a very different word order and structure, and comes from a very different family. And here, you can see that, let's say if I show you just the Hebrew text, if you don't speak Hebrew there is no way you will understand it. If you check the translation, you definitely can see there are some ungrammatical element in the translation, like the appeal to the Mayor of Netanya remain in custody, but you definitely can understand the meaning of this whole text, and translation systems become better and better.

Let me show you the case where machine translation as produced by Google really didn't work. So, what you see here on the left side is a recipe that I got from a friend in Finnish. I don't read Finnish, so I took this particular recipe into Google Translation, and from the very beginning you can see how badly it does on this translation, so it translates the first sentence: bake the flat rod and cut it into 20 pieces. Pretty much the translation of this recipe is totally unusable, and you can be asking yourself how come it did a reasonable job in translating the news and it does such a poor job in translating the recipes. And the recipe seems to be even simpler than the news, the sentences are short and use a simple vocabulary.

The reason is that machine translation system trained on news, so they know how to interpret news, they didn't see a lot of, you know, recipes written by normal people, so as a result it may produce really funny outputs. Let me go to another example. It's IBM Watson, when they demonstrated they can win the trivia game, and the question which was extremely hard for human, and I'm sure it's a hard question for you if you haven't seen that particular competition, is William Wilkinson's *An Account of the Principalities of Wallachia and Moldavia* inspired this author's most famous novel. I mean, who knows? You know, nobody does.

And when machine can give you the answer – and machine did – it sounds, wow, the machine's really incredibly smart, and they really outperform humans in their level of intelligence. However, if you just do one query to Google, you just paste this question into Google and one of the first links that you will find is this piece of text which actually very clearly states that it was Bram Stoker, who is the writer, and then they describe about his most famous character, Count Dracula, so you can clearly see that in this case, the machine doesn't possess extra intelligence that human does, it just has an ability to go and

search huge amounts of materials available online, and find the place which uses a phrasing very close to the original sentence.

Finally, like, when we are so impressed by the ability of a machine to solve these very complex questions that humans cannot solve, let me show you a question that even, you know, seven year old can answer is the small text. It's called reading comprehension task, so says Sally liked going outside, she put on her shoes, she went outside to walk. Why did Sally put on her shoes? So, you can say this is trivial, because she wanted to go outside. So, on the set of these questions which are very basic understanding questions, machine can do today with 71% accuracy. This difference, this huge discrepancy between the ability of a machine solving, like, a very complex question, where they can utilize all the data available online, and inability to solve this kind of reading understanding question may be puzzling, but it again very logical with how machines do it, and we will see a bit later how actually a machine can understand the text, and it will make more sense to us.

Finally, let me show you another example. This is example of text summarization. We all have a lot of material available online, and we would really wish somebody take this very long text and crunch it for us, so we don't need to read the whole thing, and here you can see two summaries: one written by machine and one written by humans. You maybe guess, the sentences which are longer were the ones selected by machine, but it's a very, very reasonable summary. And in this case, again, the machines which are trained to generate highlights of the stories can do a pretty reasonable job. But let me show you an example where we cannot use machines to do summarization where there is actually a huge need to have machines.

So, Cochrane is a community of doctors which is span, you know, across the whole world, which try to take clinical studies and trials and summarize them in such a way that the physicians can decide whether overall body of work in medicine support a particular decision or not. So, they may say here is an example of one of the studies where they will take all the studies that supported describe for which group they supported and so on. This kind of very deep analysis and understanding the connection at this point is beyond the capacity of our machines.

THOMAS 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 on the relevant button.