MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

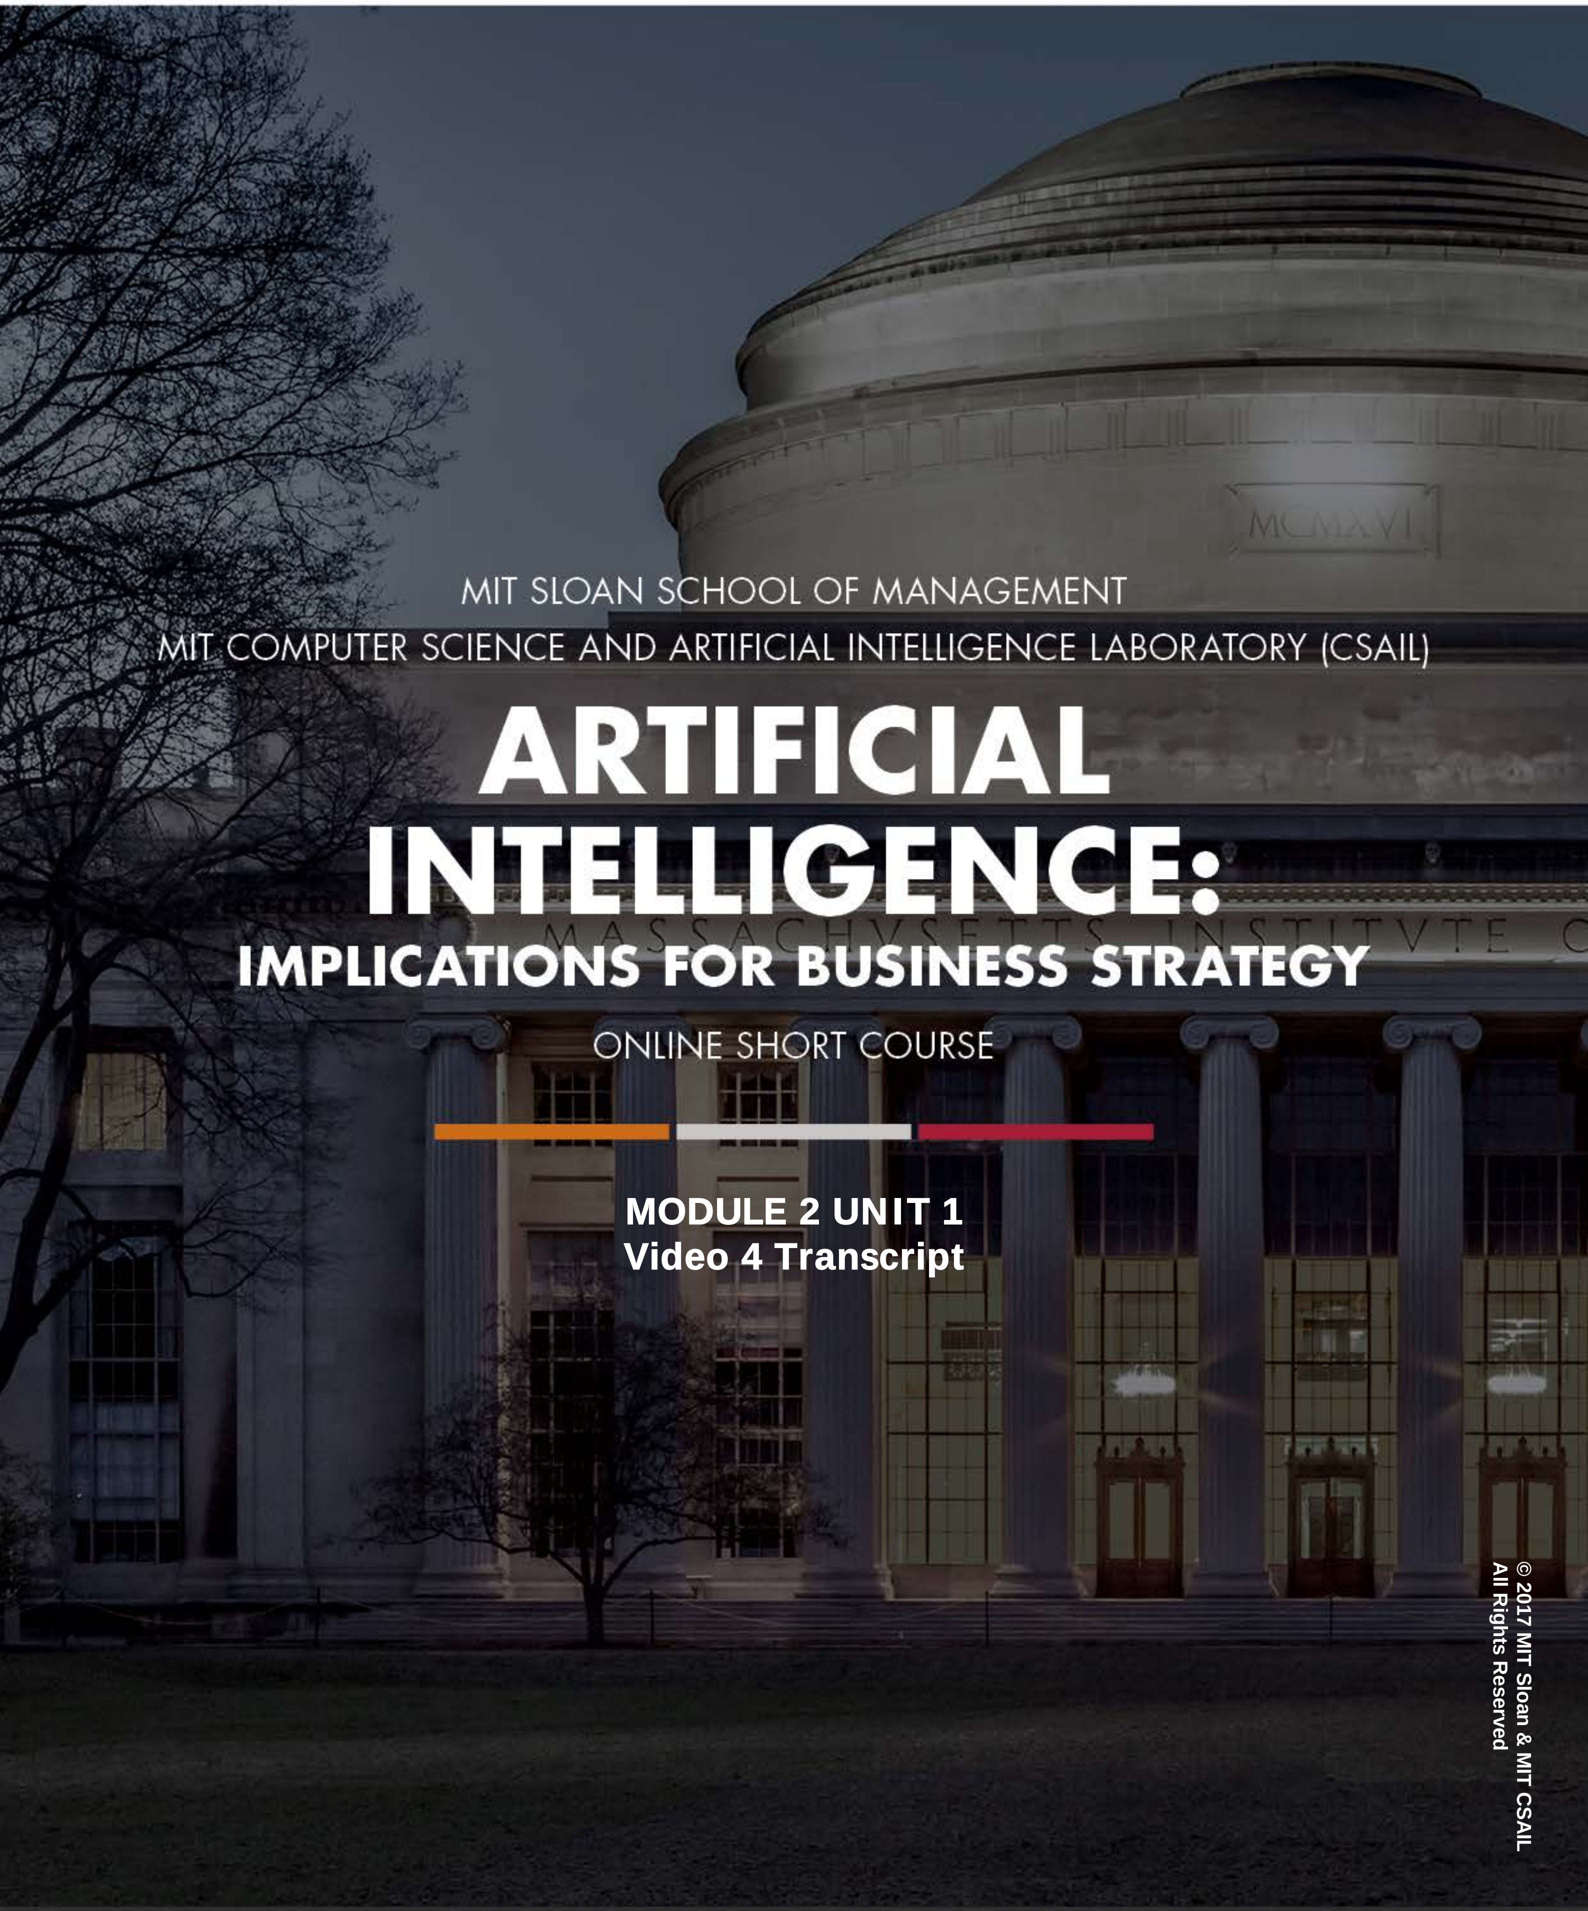

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 2 UNIT 1
Video 4 Transcript

DELIVERED IN COLLABORATION WITH getsmarter

Module 2 Unit 1 Video 4 Transcript

THOMAS MALONE (TM): You've given us a bunch of different ways that machine learning might go wrong at any of the three stages you talked about. Is there any other help, any other intuitions you can give our students to guess whether a particular problem would or would not work? For instance, suppose I say I want to use machine learning to figure out which are the most promising sales prospects for my salespeople to call on, which are the ones that are most likely to become customers if we spend time with them, versus which are the ones that are a long shot. How could I think about whether that's going to be easy or hard to do with machine learning?

TOMMI JAAKKOLA (TJ): Right. So, that would be an easy problem to formulate as a machine learning problem. Presumably, you have tracked your customers in the past, so you know at the starting point when you only were initially interacting with the customer what you did and what the consequence was. So, you have some historical data to see how a customer appears, and what the consequences are later on. And you can formulate that as now even a supervised learning problem. You know the outcome for those customers, you know you can featurize the customer in terms of their properties, and now you can learn to relate the two. And when a new customer comes, you can then apply the solution to that problem based on now featurizing the new customer in the same way, applying the prediction, and now you have a prediction for whether they are worthwhile to engage in terms of later.

TM: Okay, and by featurize you just mean represent the situation in terms of a number of features, each of which is a variable, basically.

TJ: Right. Correct.

TM: Okay. Now, suppose I haven't systematically tracked my sales operations in the past. Suppose I've only got, say, 50 examples of prospects that we started dealing with and, say, 10 of them eventually became customers. Is that enough to learn anything? Do I need 500 or 50,000? How do I know how much is enough?

TJ: Right. So, it all depends on what type of data you have. So, if you have now tons of data unannotated about customers, just have millions of customers and that's it, you don't know what the relationship is necessarily or what the outcome would be, but now you can learn about how customers vary. This isn't now an unsupervised learning problem. You learn to represent customers. You understand that some customers look like this, others like this, and you can find, sort of, a condensed representation of them. And it is that that you would then relate in the new task with just a few examples, and that would be an easier task to learn.

TM: But suppose I was trying to do a supervised learning example. Suppose I say, oh, we want to do this, let's start keeping track of things now. So, we got all the data, including whether they did or did not become customers. How many cases, how many examples do I need to have any hope?

TJ: Right. So, that's a harder thing to give you a specific number. The reason is that easy tasks are quickly learnable. Hard tasks take lots of examples to them. So, if the task is easy in the way that you formulate it, it could be that only a few hundred examples is sufficient to get to the accuracy level that is actually useful to you. If the relationship is very hard, ambiguous or things like that, it may take a large number of examples to pull out the right thing.

TM: I know you can't or don't want to be pinned down to any exact numbers because you don't really know, but it sounds like at the minimum you need a few hundred or hundreds of examples to have much hope of learning anything non-trivial with machine learning.

TJ: No, I would not say that. One single example is enough, as we discussed.

TM: If you've already got the features.

TJ: If you already have the feature representation. But that took a lot of data, as it takes a human many years to learn to understand the world, and once they do, you can instruct them with very simple feedback. The same is true for machine learning. But if you pose it as a tabula rasa learning setup, meaning that you give it no prior information whatsoever, completely empty brain from the point of view of the machine learning algorithm, then you will need lots of examples to guide it to the right solution.

TM: So, I guess maybe the lesson to take away is it's really hard to say how much data you need. You may have to try some things out or get some expert opinions and so forth. Is that a fair answer?

TJ: Yes and no. So, if you formulated the task and I could do some initial checking of what it looks like, I could tell you how many examples you should have, but since some tasks are easy, some tasks are hard, unless you tell me which one it is, I can't really give you a number.

TM: So, basically, you need a machine learning expert to give you some guidance based on initial exploration of the problem.

TJ: At this point, when, sort of, generic solutions that are not trained from lots of data prior to addressing this particular task are not yet available, you need an expert to guide the solution. But that will change.

TM: Okay, good. So, another question people often have about machine learning systems is that an issue they have is that machine learning systems can't really explain how they got their answer, they just tell you the answer. To what degree is that true today, and how do you think it might change in the future?

TJ: Right. It is currently an active area of research. Interpretability is another name that it goes by. People are trying to force, say, deep learning methods to explain themselves so that they can communicate how and why they make the prediction in terms that we can understand. And that gives us some confidence that at least it paid attention to the right aspects and its reasoning seems correct, in an easily verifiable way, even though it actually made a much more sophisticated prediction. But roughly speaking, it did the right thing.

TM: Okay, so that's an area of research now. In terms of what's commercially available right now, can people expect any kind of interpretation or explanation for an answer?

TJ: Commercially available, I'm not sure that that's available today, but it might be available next year, for example. So, it is this close that you will start getting initial systems that try to provide some explanation for the method.

TM: So, today's machine learning systems don't give you much of any explanation, but soon, within the next year or two, that might be more available.

TJ: Right. That is based on the need. As they become more prevalent in, say, medical context. It is important to make the physician comfortable that the medical decisions are based on accepted medical knowledge, even though the method can scour vast databases and image databases to make that prediction, but at least its reasoning seems consistent with physician understanding.

TM: Great.

TJ: It's much easier to verify a solution than to come up with a solution yourself.

TM: Yes. Okay, good. Do you have any other kind of advice for people who are thinking about machine learning and business applications?

TJ: Formulation is the key. In many cases, the formulation is easy to do, so learn about machine learning to the extent that you start recognizing machine learning problems all over the place. Almost any problem can be formulated as such, but it takes a little understanding to start seeing where they are.

TM: And what you mean by formulating then would be what are the inputs and what are the outputs that you're looking for?

TJ: Right. Starting to think about problems not in terms of solutions directly, but in terms of formulating it as, in this case, this is what I want, and here are illustrations of what I want, and creating such data sets and now they can be passed on to a machine learning algorithm.

TM: Great. So that's one piece of advice. Any other advice?

TJ: The world is on an accelerating pace, and it will be very interesting to see, and there are lots of applications and business opportunities that are available in that space.

TM: Okay, thank you very much Tommi. 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.