MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

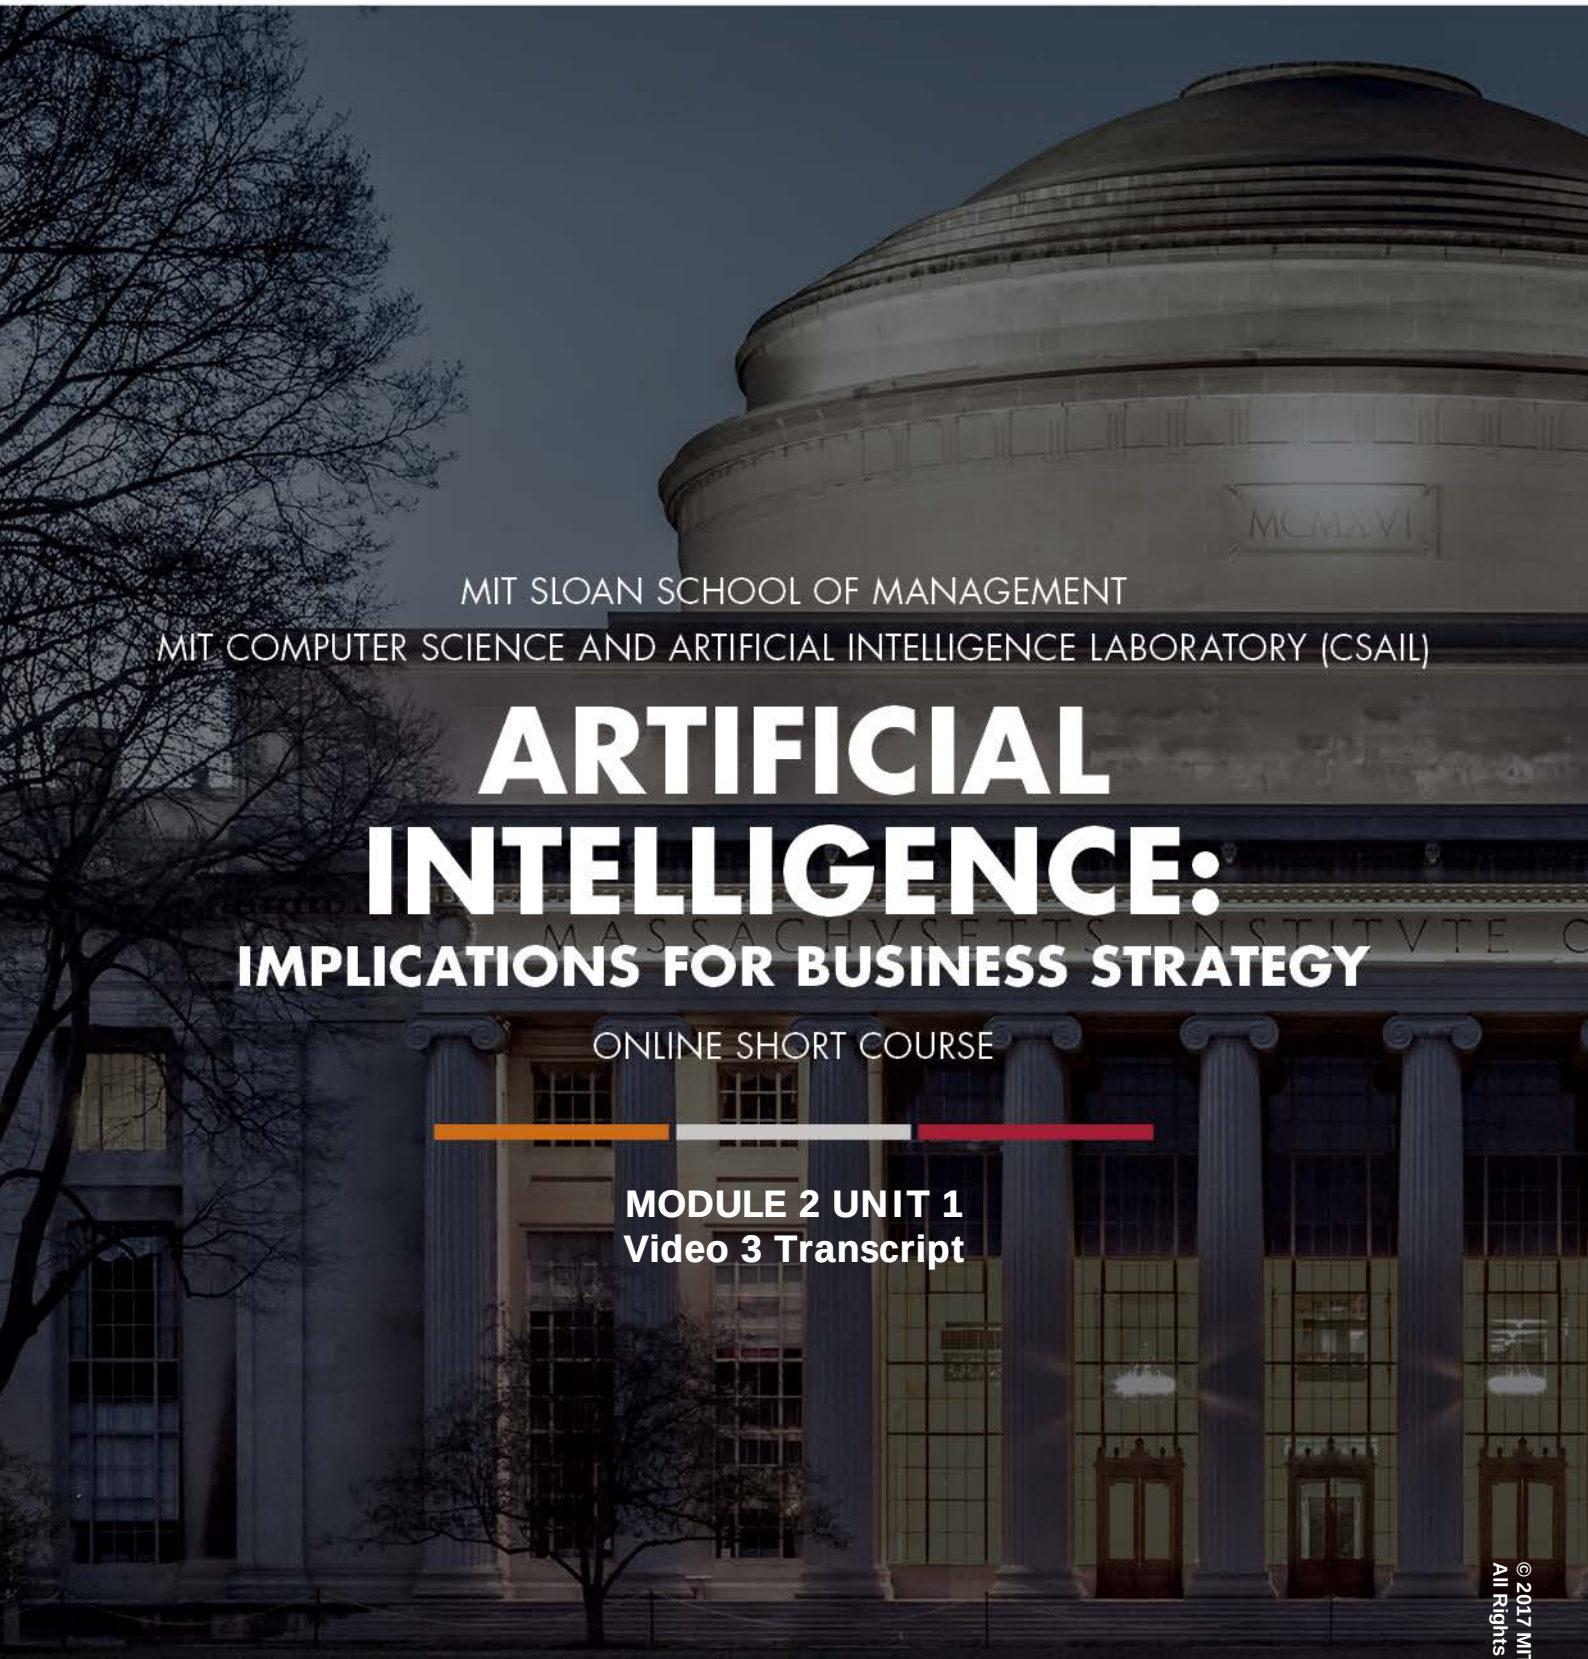

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 2 UNIT 1
Video 3 Transcript

DELIVERED IN COLLABORATION WITH getsmarter

Module 2 Unit 1 Video 3 Transcript

THOMAS MALONE (TM): So, Tommi, I thought you did a very nice job of presenting machine learning and what that all means. One question that I think a number of our students may be wondering about is, what is deep learning? It's something like machine learning, but can you give them a better understanding of what that actually means?

TOMMI JAAKKOLA (TJ): Sure. Maybe I'll start with shallow learning, just as a contrast. If you think of classifying images, an image is just a bunch of pixels, and you can feed that into a simple classifier and try to see whether it would learn what the images are about. So, there's no sort of construction of features, it just focuses on the pixels. That would be shallow learning. Deep learning means that you do layers of processing of the images, trying to represent the images in terms of features, combinations of features, parts of an object, arms, legs, head, eyes and so on, and then combining them ultimately to answer the question.

TM: So, shallow learning would be where you have just one or two layers between the pixels of the image and the answer of whether this is a person or a cat, and deep learning would be where you have multiple layers, and the lower layers would be things like, is this an edge or an angle or a leg or an eye? And the higher levels would then be, kind of, is it a person or a cat?

TJ: That is correct. So, deep learning tries to understand how to represent the images, as well as giving you the answer to that specific question that you asked.

TM: So, the simple answer is, deep learning is just more layers.

TJ: Correct. It is a very loose term in that sense.

TM: I think another big question our students may be wondering about is, what's easy and what's hard with today's machine learning algorithms. Sometimes they do things that seem kind of like magic, sometimes they do things, or they can't do things that any four-year-old human could do. Can you give us some more intuitions about what's easy and what's hard?

TJ: Sure. I would break that question into three parts – whether I can formulate it as a machine learning problem, whether I have tools already available to address that formulation, and whether it's likely to work well in terms of as a solution to the problem. So, in terms of the formulation, many tasks can be formulated in terms of posing it as, say, a simple supervised learning problem of giving an input and the correct output, and lots of examples like that. So, the posing part is really what has, sort of, spread machine learning all over the place, because that is very easy, and many tasks can be broken in that way. You don't have to know the solution to pose it as a problem. Now, increasingly, there are tools available to address many of the standard formulations of the problem – whether you're classifying images, whether you're interpreting natural language, whether you're interpreting speech, or, say, making business decisions or trying to predict how markets evolve and so on. So, many of the tools are available, not for all tasks, but for many tasks they are. Then the last part is whether it's possible to solve the problem. If you're trying to predict a chaotic system, how it evolves, however you formulate it, whatever tools you have

available, it's simply unpredictable. So, it will not work, even if you can address it as a machine learning problem.

TM: What would be an example of something that's probably unpredictable?

TJ: Like planetary motion, sometimes they can be chaotic.

TM: Or, what's the stock market going to do tomorrow?

TJ: Well, that is less chaotic. People are driven by news articles, company releases, judgments about what's hot, what's not. So, there are lots of signals that correlate in terms of how people invest, and therefore the valuations in the market. So, at least, sort of, a medium to long-term investing, there seems to be a lot of signal left to exploit. High-speed trading is quite exhausted in terms of automated systems.

TM: So, some things are hard to formulate, but nearly everything could be formulated as supervised learning. Some of those things might or might not have pre-canned or pre-written programs to use them, and some might not still work even though you can apply the techniques. You just wouldn't get very good answers.

TJ: Right. So, the fact that you can solve it as a machine learning problem, oftentimes it works. You get something out of it, but it could be that the task is inherently not possible to solve, and as a result, a machine learning approach wouldn't work either. Not every task can be processed as a supervised learning task. There are other types of learning from lessons, like unsupervised learning and reinforcement learning and things like that.

TM: So, can you give us high-level intuition for what unsupervised learning and reinforcement learning mean?

TJ: Sure. Unsupervised learning is trying to understand, based on observation, how things work. What's their regular structure in terms of images? If you are just watching videos, you learn to understand how objects move and relate to each other and so on, based on observation alone. And that's very powerful to learn those representations that we discussed in the deep learning context, purely in an unsupervised learning setting.

TM: So, you don't need a person teaching you that this is a person, this is a cat, and things. You can't move things unless they're on top of a pile. You could learn those things by yourself potentially.

TJ: You wouldn't learn the labeling. You wouldn't know that we call it a human, we call it a cat, but you would know that it exists, and it is a regular object and it has this and this constraints and so on.

TM: So, if you had already learned that there are these things in the world that kind of look humanoid and other things that look cat-oid, and then somebody said, this is a human and that's a cat, you might be able to attach the label with only one more example.

TJ: There lies the power of unsupervised learning. You learn to represent things in the right way, so that when the feedback comes you can really associate it with, sort of, the highest-

level category that you already know how to represent, and therefore learn with very little feedback.

TM: Great. And then you also said reinforcement learning is another kind.

TJ: Right. Almost anything is a reinforcement learning, but typically it's used to learn to act. When you actually take action in the world, the world changes its state in response to your actions, so you need to formulate the task slightly differently. If a robot reaches an object and pushes it forward in doing so, it changes the state of the world. So, it has to learn to take that into account in learning to perform in the world.

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.