MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

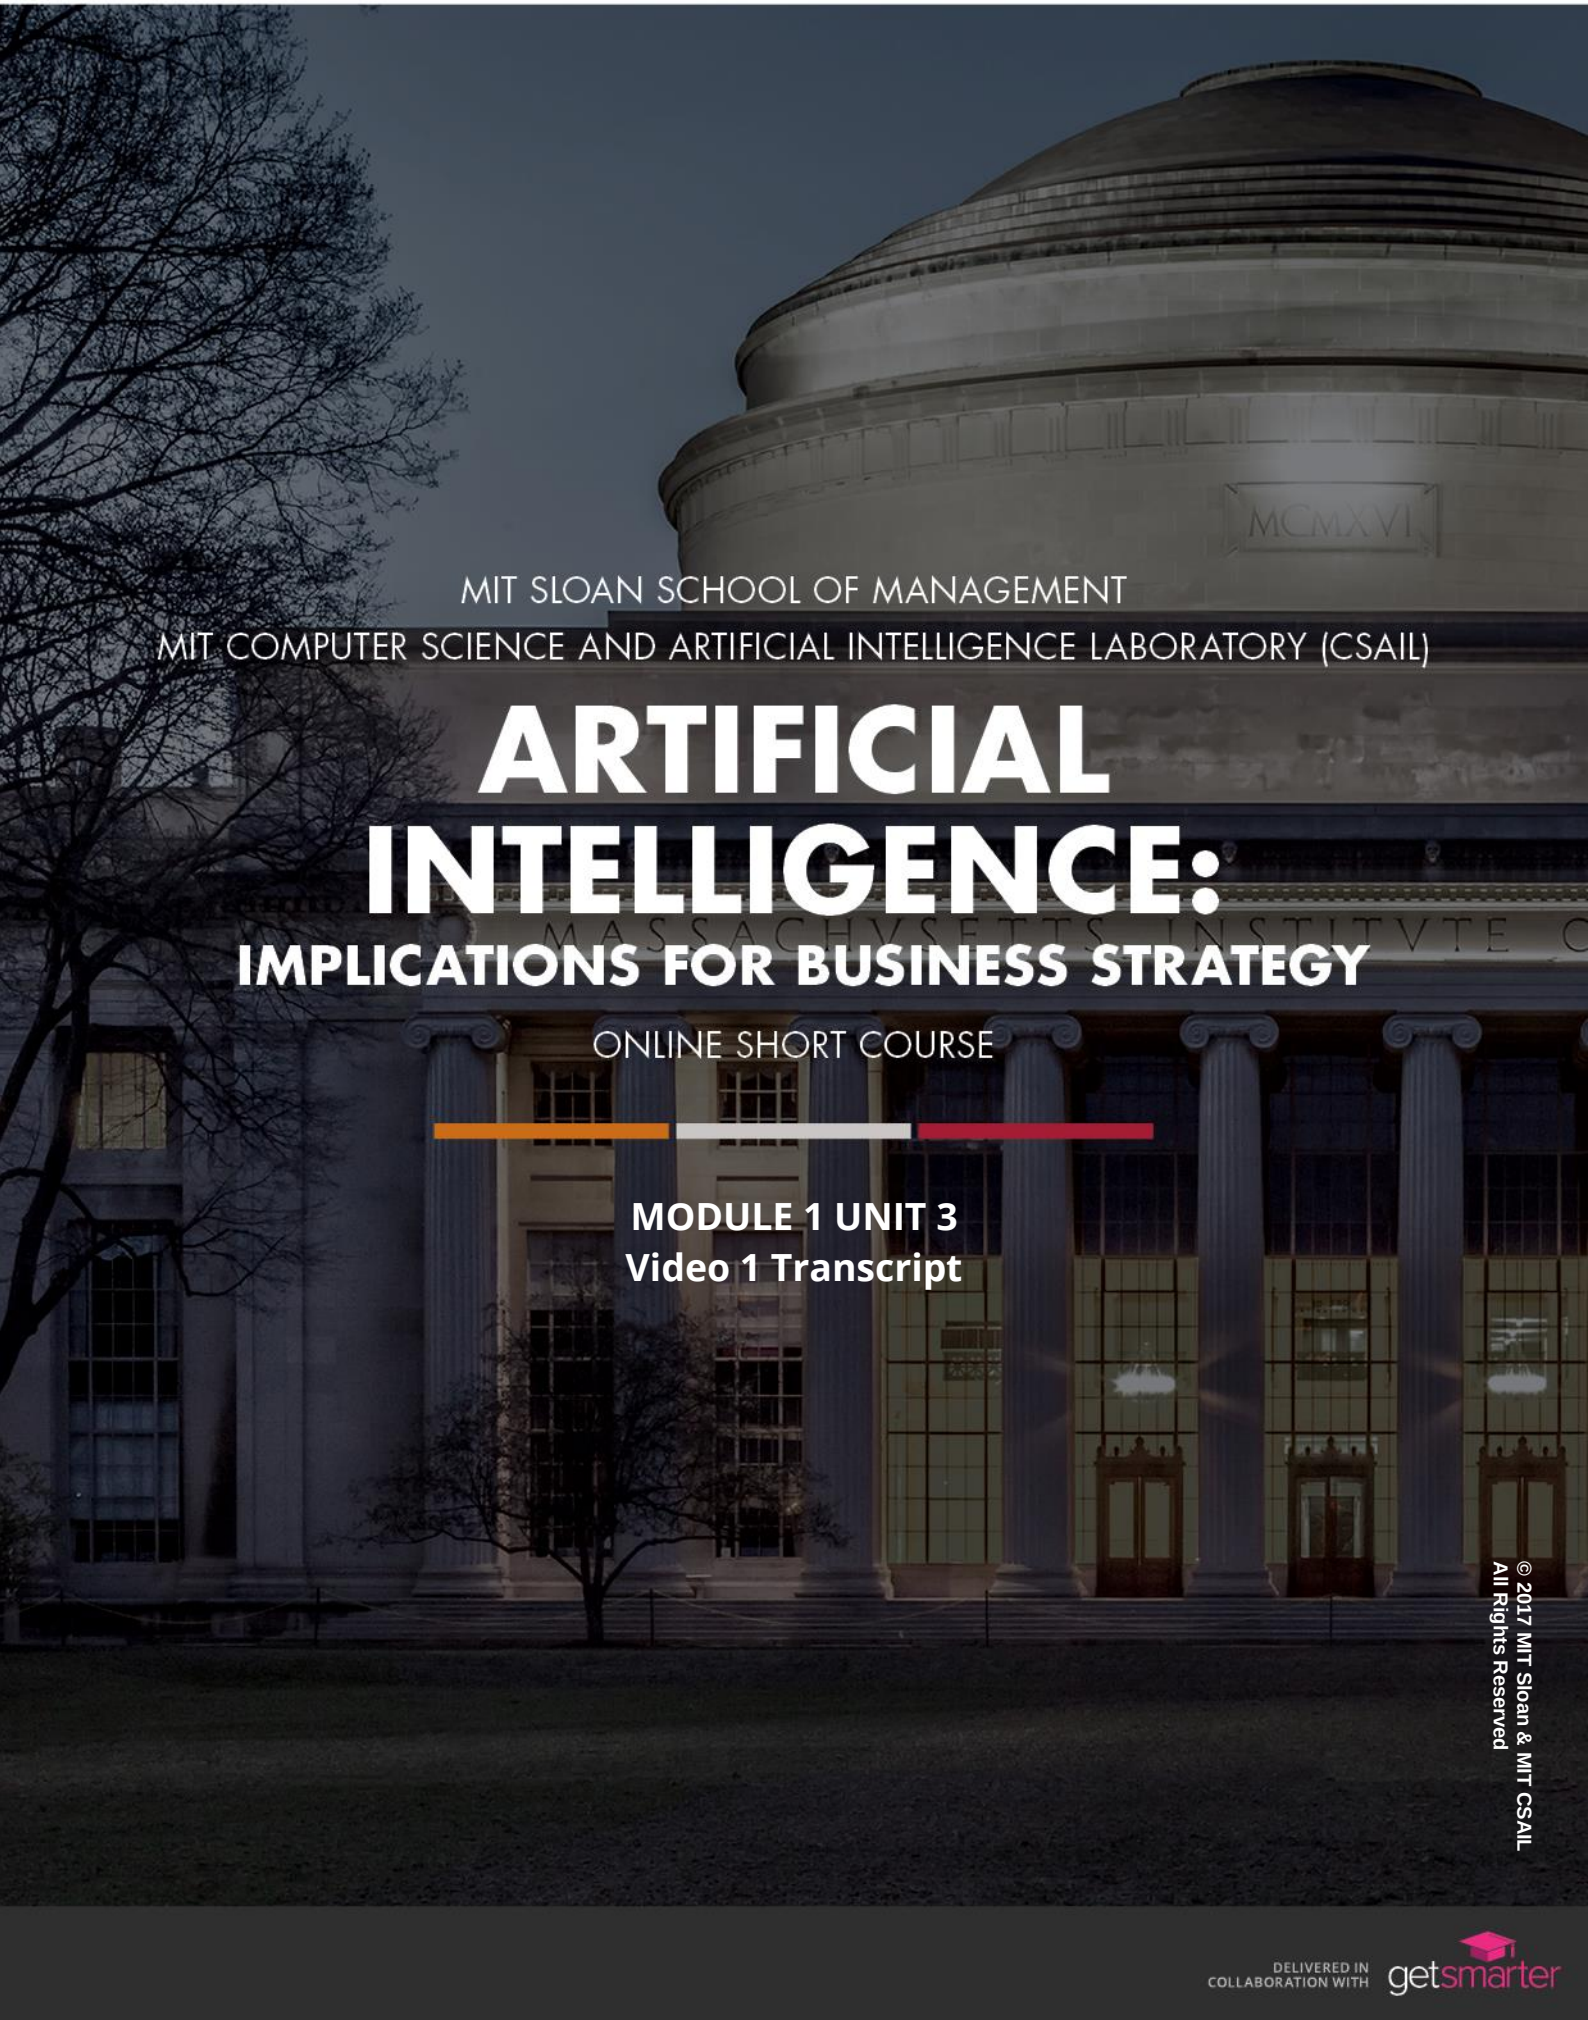

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 1 UNIT 3
Video 1 Transcript

DELIVERED IN COLLABORATION WITH getsmarter

MIT AI M1 U3 Video 1 Transcript

TOM MALONE: As we saw at the beginning of this module, essentially all real uses of AI in business involve both people and computers. People are involved in creating the software in the first place and modifying it over time. People are involved in deciding which applications to use when and what to do when things go wrong. And, in many cases, people are involved all along in doing the parts of the process that computers can't do by themselves.

So, the systems we're really trying to optimize and design here are not just computer systems, they're human-computer systems. To do that, we need to ask two questions. The first question is: What tasks should computers do, and what tasks should people do? The second question is: How can this human-computer system improve over time. So, let's start with the first question: How are tasks divided between people and computers?

A good goal is to let machines do the things they do better than people, and people do the things they do better than machines. For instance, machines are much better than people at remembering huge amounts of information; and people are usually better than machines at interacting flexibly with other people. But we shouldn't just be trying to replace people with machines everywhere we can, we should also be trying to figure out how to let the human-computer systems do things better than any person or computer could ever do them before, and do new things that were never possible at all before.

For example, let's think about the Google search service. Depending on how you define AI, that's arguably the most widely used AI system in the world today. Here the people are creating the content, most of it, and linking it to other content.

The machines are remembering vast amounts of content and helping people find the content that's most relevant to their interests. But the Google search service didn't just replace reference librarians who did a similar function in the past, the Google service allowed vastly more creation and searching of knowledge and created many new jobs for creating, searching, and advertising.

Or think of Wikipedia. Here people create and edit almost all the content. But machines like the Wikipedia bots are better than people at quickly scanning for things like obscene words and plagiarism from other sites on the web. So, the combination of people and bots works better than either alone. But Wikipedia didn't just replace paper encyclopedias, though it did pretty much do that, it also vastly expanded the scope and timeliness of what goes into an encyclopedia and the occasions upon which we use an encyclopedia.

One more example is a CSAIL cybersecurity system where people and machines work together to detect cyberattacks. The machines are much better than the people at detecting unusual activity on the network. But the people are better than the machines at recognizing which kinds of unusual activity are just random and which are due to malicious intent.

So, it turns out, in this case, that the combination of people and computers together recognizes three times more actual malicious attacks than the machine alone. One useful way of thinking about how tasks can be divided between people and computers is in terms of four roles computers can play relative to people.

The first role is tool, where computers do tasks you give them, but you generally monitor every step of the way. Word processors with autocomplete, spreadsheets, cars with cruise control, these are all examples of computers playing the tool role. In fact, most ways computers are used today are as tools to connect people to other people. Think of email, the World Wide Web in general, Netflix, Facebook, on and on.

Lots and lots of the applications we use today are really about helping people communicate with other people and collaborate in new ways. This use of information technology is what you might call hyperconnectivity. And, I think, this may well be the most important way that information technology helps create more intelligent organizations.

But there's another very important way of creating smarter organizations. And that's creating smarter computers. That's the main focus of this course. For example, human loan officers can use artificially intelligent credit scoring algorithms to help them decide who they should give loans to. Google search algorithms also act as tools when they find results almost instantly as you're watching. But for them to work so fast, there's another part of the Google technology that's always scanning the web and indexing what it finds there.

So, that part of the Google service begins to merge from tools into the next role that computers can play. So, that role is the role of assistant. Unlike a tool, an assistant can work without your direct attention. It may take more initiative. And it may be more active in helping to formulate and solve a problem.

For instance, you can think of semi-autonomous cars, the kind that are on the roads today, as a kind of assistant, a driver's assistant. Another example would be some of the applications of IBM's Watson technology. For instance, in some cases this technology is used to help doctors diagnose medical cases. The technology has digested vast amounts of medical literature. And it takes specific information about a particular case and calculates different possible diagnoses and the reasoning paths that would lead to those diagnoses.

It presents all that information for the doctors to see. They can query as needed. And then the human doctors make the final decisions about what diagnosis to use. Another example would be the chatbots that KLM Airlines uses to respond to customer queries and problems via social media. These chatbots may seem like pure artificial intelligence, but in fact they're not because there are people in the background too.

The artificial intelligence software creates suggested replies to all the messages they receive from customers, but those replies are then evaluated by humans who in some cases decide to override the recommendation from the machine and create their own written responses. Facebook also uses algorithms to generate candidates for trending topics and for ad categories, but these items that are generated by machines are later reviewed by people to decide which ones are most appropriate.

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.