MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

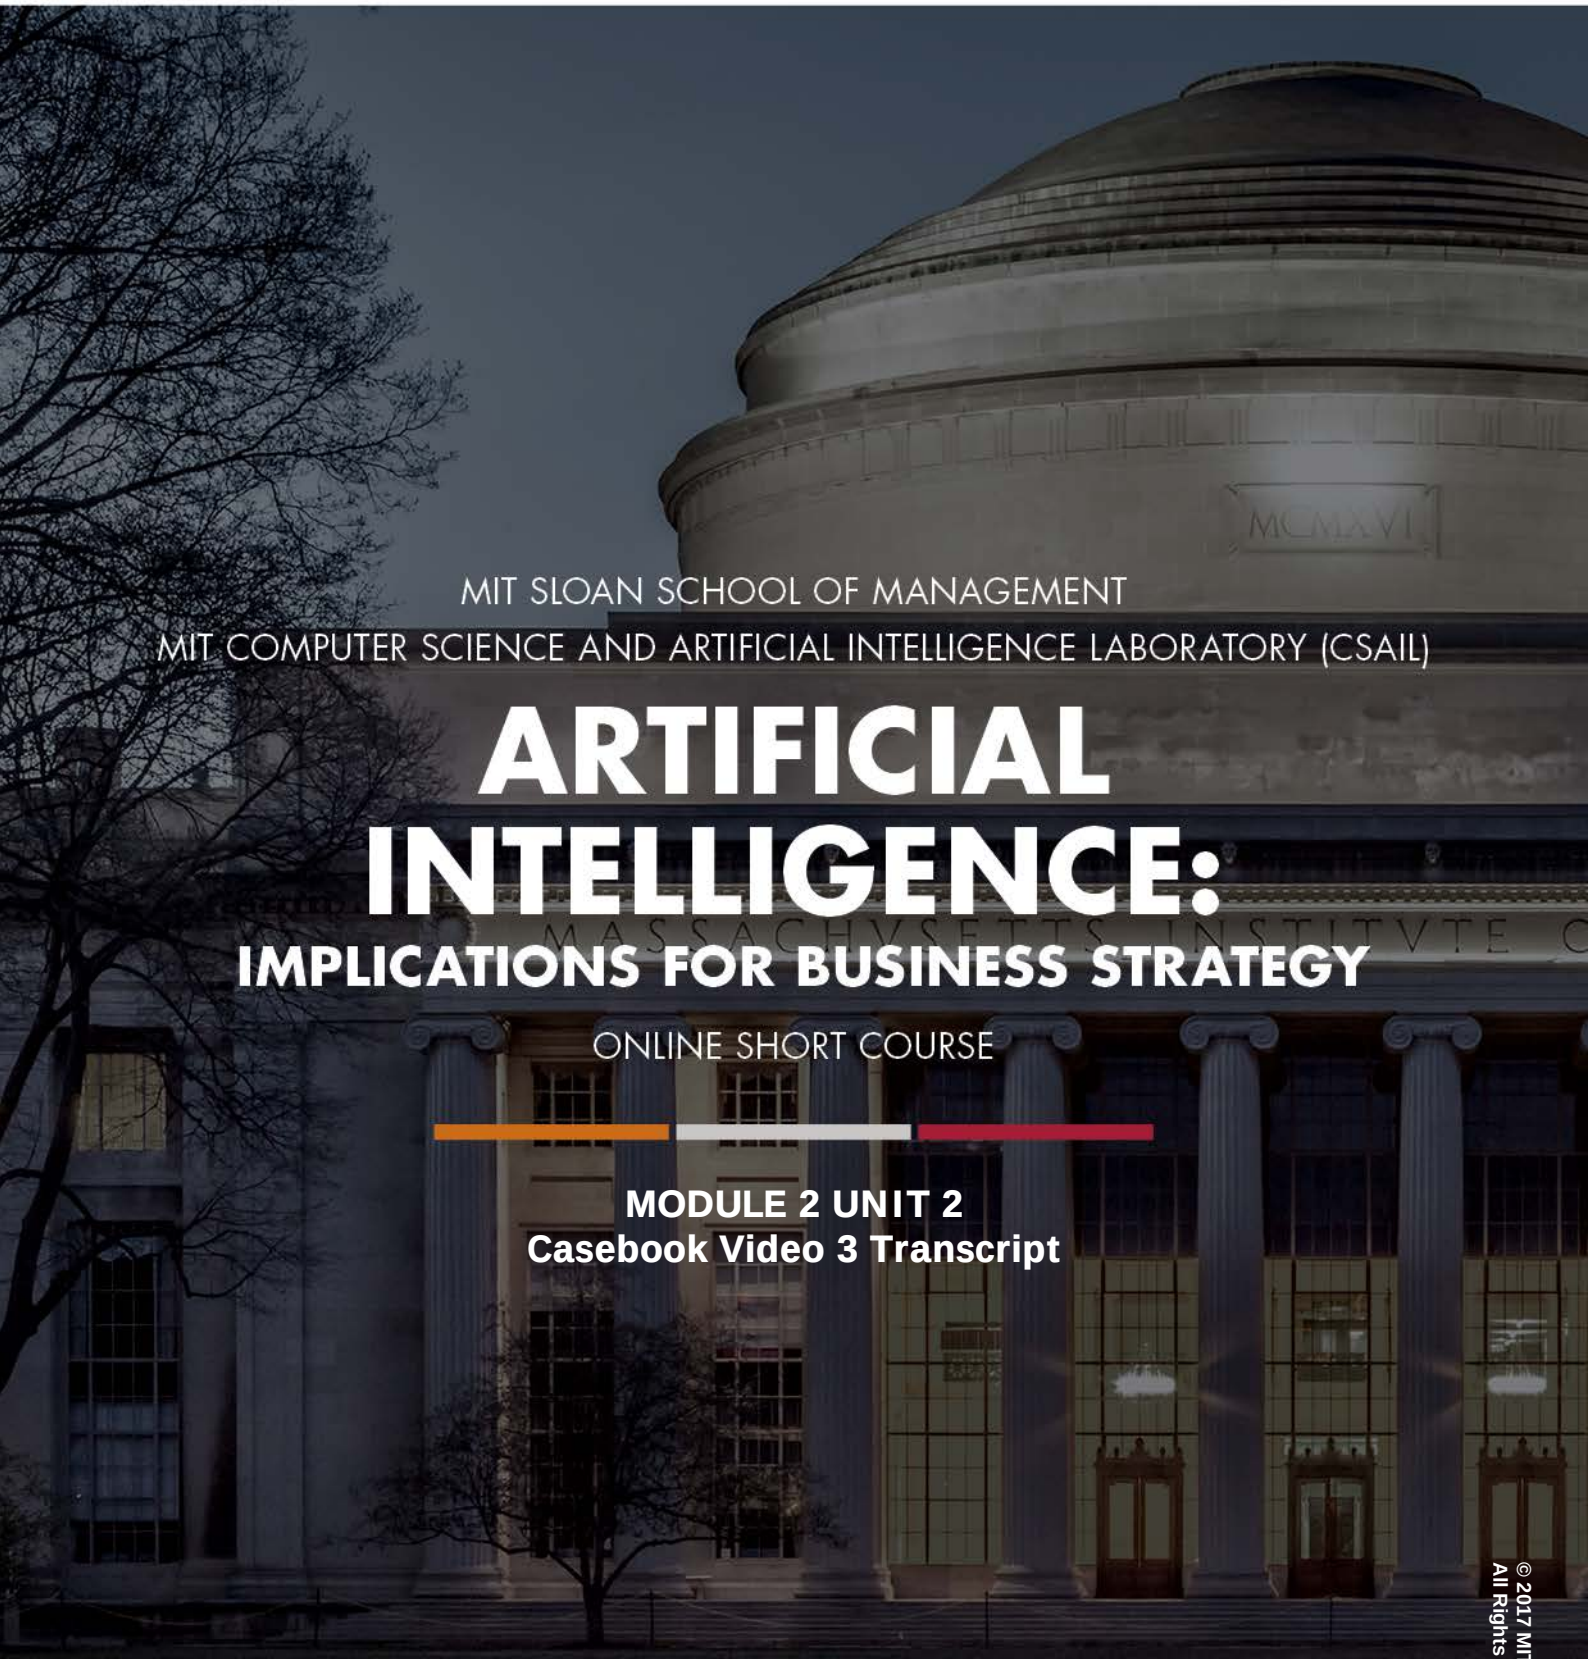

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 2 UNIT 2
Casebook Video 3 Transcript

DELIVERED IN COLLABORATION WITH getsmarter

Module 2 Unit 2 Casebook Video 3 Transcript

THOMAS MALONE (TM): So, I'm here today with Professor Randall Davis. Randy is a Professor of Computer Science at MIT, and he's also an entrepreneur, a Co-Founder of a company called Digital Cognition Technologies. Randy, it's great to have you here today.

RANDALL DAVIS (RD): Glad to be here, Tom.

TM: So, maybe we could start by you just telling us a little bit about the product that your company makes.

RD: Yes. The product is called DCTclock. It's a radical revision of a test, a psychological test. It's actually been around for 50 years. Traditionally, it's a pen and paper test. You give the subject a piece of paper and a pen and ask them to draw a certain object. We do the same test, but we have the person draw it using an off-the-shelf digitizing ball point pen that captures not only the drawing but the timing of the drawing. And by being able to get very accurate information about exactly what was drawn and very accurate information about when it was drawn, the speed of the pen, the delays at various places. By capturing both the product and the process and analyzing them both in detail, we were able to get a much more effective screener that helps us to determine whether someone appears to be functioning normally, cognitively, or whether they're worth following up with in order to dig deeper and seeing if there are problems.

TM: So, it's a way of screening people for potential cognitive problems. What are the business benefits, or, in this case, I guess you could say the medical benefits of using this test versus the traditional method?

RD: There are several benefits that accrue in the medical environment. One of them is reliability. There is an automated system doing the analysis, which means every time the analysis gets done, it's done exactly the same way. And it's well known that on some of these tests that are administered manually, the person-to-person variation in interpretation can be quite substantial, and that's a problem. So, we're providing the clinician with a more reliable measure. We're also providing a much more detailed and informative measure because, as I said earlier, it's analyzing both the product and the process. And it turns out there are small delays while people stop to think, that are perfectly normal in the test but, in people who are impaired, those delays can be significantly longer, and even attending to several of those is often a useful indicator that's part of our screener.

TM: So, it gives you more detailed information about the patient. My understanding, from what I've heard, is that it also allows a clinician to at least, if not conclusively diagnose, at least get very early indications of what might be serious cognitive problems, but that might not be visible until potentially years later. Is this true?

RD: Yes, that's the key potential here is the ability to spot signs of impairment before they manifest with classical clinical symptoms.

TM: So, that actually strikes me as having an interesting analogy with things that could happen in many other business environments. Instead of just recognizing symptoms of a potential medical problem in a person earlier than otherwise possible, you might be able to use an analogous approach in industrial or many other situations where, for instance, if you're doing more detailed sensing of a machine, like a jet engine or a factory machine or something like that, you might be able to detect potential problems far earlier than they would otherwise be visible, and be able to do preventative maintenance or take some other kinds of actions. Does that seem like a fair analogy?

RD: Yes. The basic idea is don't wait until it breaks. And the important point there, I think, is that I think to do it well, certainly in the medical environment, but even, I think, in the industrial environment, you need to be able to combine several things. You need good sensors. You need people who are good at interpreting those sensors. And you need advanced AI and machine learning technology. And I think you need all three of those together. If you just take all of the sensors and drop them into a machine learning system, you may or may not get a result that is sensible to you and it will be difficult to understand why this system thinks something is wrong. To the extent that you can guide the machine learning technology by having somebody who is themselves expert at making sense of the behavior of those machines, who can listen to it and say, wait, that doesn't sound quite right, let me go over here and look at this part of it, to the extent that they can help guide the information that is extracted from the sensors and then handed over to a machine learning system, we've found in our domain that that's a much more effective way to proceed and that's exactly what we do in this clinical domain. We have someone who is used to interpreting these tests, telling us what to attend to so that we can tell machine learning technology what to attend to. And that combination of human and machine learning technology seems to me to be particularly effective.

TM: Great. So, I think that's a very nice example of how it's not just a matter of the technology figuring it all out by itself; it's a matter of an interplay between the expert human, in this case, and the machine to figure out how to do a diagnosis that would never even have been possible before.

RD: Yes. And there's actually two aspects to that. The one that we've alluded to is the initial screening process – does this thing seem like it's working properly, and building an AI system to do that – and its ability to say yes or no, and if it says no, then you can make use of possibly your in-house human expertise. But now you're focusing it. Instead of having them having to do the screening, they can be asked to follow up on the situations where there really are signs that there might be something wrong that only they might know how to deal with.

TM: Great. So, it's a combination of human and computer at both places – both the development of the initial system and then the dealing with whatever the results are.

RD: Yes, that's a good way to put it.

TM: Great example. So, let me ask you to dive a little bit more into the technology itself here. I understand that your product uses AI technology, but can you tell us a little more about what kinds of AI and a little bit of an intuition for what's going on in this technology?

RD: Sure. In the particular test that we're talking about – it's called the clock drawing test – we give the subject a blank sheet of paper and ask them to draw a clock showing 10 minutes after 11. When they do it with the digitizing pen, the result is a whole bunch of pen strokes that the system then has to interpret. Now, one of those pen strokes is the three, another one is the hour hand, another one is the clock circle. But the artificial intelligence in the system, its job is to analyze those pen strokes and make sense of them, where make sense of them simply means to classify them. This is the three, that's the seven, that's the 12, that's the clock circle. So, the first step is analyzing the pen strokes to classify them. The next step is to extract hundreds of features from those pen strokes, both geometric and temporal features. So, what was drawn and at what speed and where were the delays? And that combination of both geometric and temporal features are then what goes into the machine learning aspect of the system.

TM: And then what does the machine learning learn?

RD: The machine learning has been trained on a set of previous examples to help us to screen. It looks at the hundreds of features that we have decided are likely to be informative, and it selects the ones that it thinks are going to produce the best results. It optimizes a combination of those features to make the best decision between infirm and normal cognition.

TM: Okay, good. So, there's at least two kinds of AI in your application, it sounds like. One is the sketch understanding that's not really machine learning, it's another kind of AI, and then there's the machine learning capabilities that recognize or figure out how to determine which are the most indicative features, either spatial or temporal, in the data that's been gathered.

RD: Yes, that's right. Two very different capabilities there.

TM: Okay, great. So now, let's try generalizing this a little bit. It sounds like a very interesting application. Let me now ask me to put yourself in the position of somebody at an insurance company or in a pharmaceutical company who is considering buying a product like yours, or potentially this could apply to many other kinds of products. If a vendor comes to you and says, hey, we've got an AI program, it's just what you need, what should a potential buyer of such a product think about in deciding whether to buy this software application? Is there anything, kind of, unusual or different about AI applications from any other kind of software product?

RD: Yes, to first order, I think the answer is no. The obvious question is, does it work? And on what grounds are you telling me it works? Show me the test that you've run, show me the performance. And it's also very interesting to ask, show me the failures, show me the cases that it didn't handle properly. There's no shame in that. No program is perfect. And

it's often quite revealing to see not only what the program can do but what it can't do yet. And you want to know that for a number of reasons, one of which is it will keep you from applying that system in a circumstance where it might not be appropriate. So, to first order, it's just software and you ought to ask the obvious questions about it. A second order issue is if you are going to integrate this into your IT department and you're going to be responsible for extending it and maintaining it, then you have to think carefully about having the right kind of expertise in-house to do that. And those people, for the moment at least, are hard to come by, so that's at least a consideration. Make sure you've got the appropriate expertise in-house to continue to develop and improve the system and maintain it, as all software has to be maintained.

TM: Okay. And I guess this last consideration would also apply to a company that's thinking about not just buying an AI application from some outside vendor, but developing their own AI application for some particular problem they have in their own company. Is that true too?

RD: Yes. Just the same idea. You want to be careful about picking the problem appropriately, making sure you've got folks who have the appropriate expertise. One of the standard promises of machine learning technology is that it can improve its performance over time. As you gain more data, it can be retuned with that data in mind. And that's largely true, but it still takes some expertise and it still takes some knowledge of how to use that technology to make that promise come true.

TM: Okay. So that's a very good point about machine learning. You know, I think many people who've had any experience with software at all know that software often takes a long time and usually, or at least often, longer than you think it will. Is there anything unusual about AI in this regard? Is it unusually hard to get it right? Is there anything else that's different about the AI development process compared to any other kind of software?

RD: It's unfamiliar, the first order. So that's the first problem. But even once you're a little more familiar with it, it's not programming as usual. There's a different mindset to it, there's thinking about how people might think about the problem. It's not just writing a whole lot of code. And you have to get used to that aspect of it as well.

TM: Especially the machine learning part where you're writing the code, but then the main thing that happens is that the code learns from the data and you have to help it do that.

RD: Yes, and you have to understand how it learns in order to guide it. Because ultimately, you do have to guide it. It will optimize, but you have to set up the optimization properly – selection of the right features and so forth.

TM: Great. So, any other advice you have for somebody who's thinking about buying a software application that has AI in it, or potentially developing their own application, not just in the medical world, but in manufacturing or any other place where people might use machine learning or other kinds of AI? Any other advice you have?

RD: Yes, there's a generic issue with respect to new technologies, of which AI and machine learning, I think, are one good example. And that's to keep in mind a principle I call the no-magic principle. It's really very simple. There's no magic. And as much as we would love to find a magic bullet that solves problems, they're not out there. So, maintain a healthy level of skepticism. Ask to see why it is that this thing works. And as I said earlier, ask to see the situations where it doesn't work. Get an understanding of what it's actually capable of. And then you can make a sensible decision.

TM: Great. That sounds like great advice, Randy. Thank you very much.

RD: Glad to do it.

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.