MIT MANAGEMENT EXECUTIVE EDUCATION

MIT CSAIL

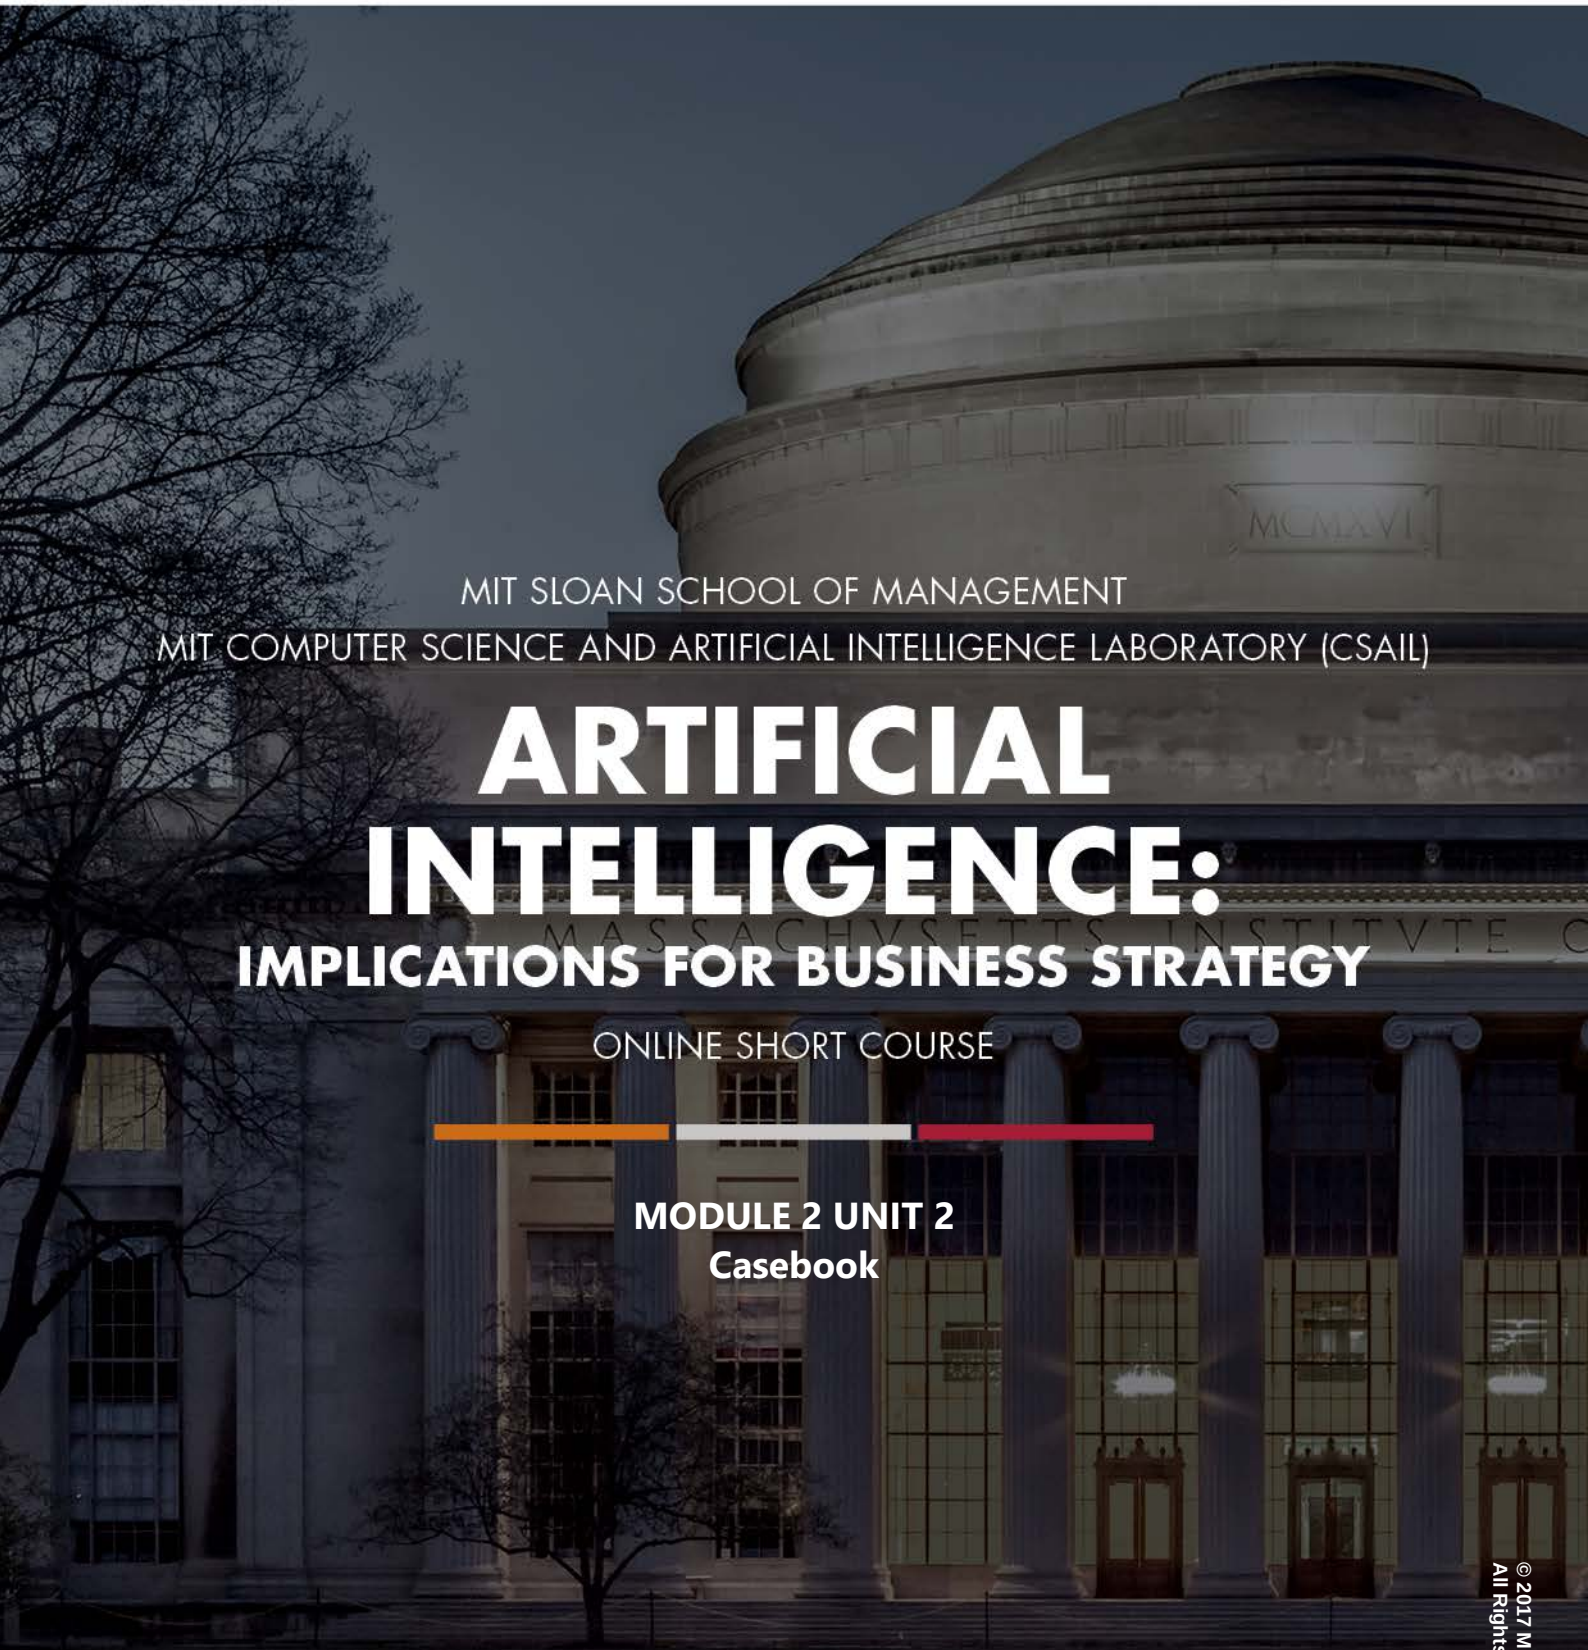

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE:

IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 2 UNIT 2
Casebook

DELIVERED IN COLLABORATION WITH getsmarter

Table of contents

Learning outcomes:

LO2: Illustrate how machine learning is currently deployed in industries and across functions.

LO3: Investigate how an organization can use machine learning to achieve cost leadership, differentiation, or focus.

LO4: Decide if an application of machine learning is appropriate in an organization.

1. Introduction

Organizations of all sizes are increasingly turning to machine learning for innovative business solutions. Machine learning is impacting many different industries and changing how people work and live. In his videos, Professor Tommi Jaakkola defined machine learning, introduced different types of prediction problems that it can address, and discussed different types of machine learning. This casebook describes applications of machine learning and illustrates how it is currently being deployed to transform industries and organizations.

2. How is machine learning used in business today?

An executive's guide to machine learning:

In an article from McKinsey & Company, machine learning is explored from an executive's perspective. The article discusses how organizations are using machine learning for insights, the importance of strategy in getting started with machine learning, and the role of senior executives in leading such initiatives.

In Video 1, Professor Thomas Malone discusses two main ways machine learning is used in business today: sensing and predicting.

Video 1: How machine learning is used in business.

Optional reading:

Read about how companies are using machine learning today.

Some examples of organizations using machine learning for sensing and predicting are explored in the sections that follow.

2.1 Sensing

2.1.1 Image recognition

Primer on how machine learning systems recognize images:

In these videos, the Head of AI Research at Facebook explains how computers are able to decipher images, the key mechanisms behind deep learning systems, and the principles behind convolutional neural networks.

Try a deep learning live demo, which is powered by a deep convolutional neural network trained on 1.6 million annotated images. In this demo, you can upload an image (such as one of your own photos), and the system will predict what kind of scene the photo depicts. As you'll see, the system isn't always correct, but it is often surprisingly accurate.

Example 1: Shoegazer

Shoegazer uses image recognition and machine learning to tell users the make and model of a pair of sneakers if they point their smartphone's camera at the shoes.

Key points and strategy

Applications like Shoegazer could be used to reduce marketing costs (customers point the app to a pair of shoes and click to buy a pair for themselves) or to support the development of new features for shoes (customers use the app to show features they like). This kind of machine learning can be used not only for images, but also for sounds (for example, Shazam is a mobile phone app that can identify songs).

Optional reading:

Read more about the Shoegazer app and the potential of image recognition to transform a business.

2.1.2 Image analysis

Businesses are awash in images, which range from those of security cameras and employee smartphone cameras to social media posts by consumers talking about brands or product categories.

At the same time, image analysis is also a critical subcomponent of robotics (which will be explored in Module 4). In robotics, image analysis can help the robot navigate its surroundings, recognize objects to be manipulated, and avoid dangerous conditions.

Example 1: TellusLabs

Read about how TellusLabs uses image analysis and machine learning to predict crop yields.

Key points and strategy

TellusLabs exemplifies how machine learning can be used for analysis of satellite images to make better predictions, be it to predict crop yields or the location of moving infrastructure like oil tankers. This company is pursuing a strategy of differentiation, to provide better predictions, for example, to customers who can gain an edge in domains like commodity trading because they have better information than competitors do.

Example 2: Mobileye

Mobileye makes and sells computer vision systems to vehicle makers for collision avoidance, driver assistance, and autonomous vehicles. Rather than hand-code these vision systems, the company uses neural networks-based machine learning to make sense of the large volumes of imagery taken on the road. Reinforcement learning trains the

system to recognize familiar objects such as road markings, pedestrians, and other vehicles, and then annotate them so that other parts of the Mobileye system can react appropriately by braking, steering, or alerting the driver. If the vision system cannot identify an object, it is flagged for further analysis, which leads to the addition of new images to the training set. Mobileye's system has been installed in 15 million vehicles. For more information, read about how Mobileye interprets a scene.

Key points and strategy

People play a role in improving the accuracy of the scene analysis by identifying the objects that the machines have flagged as unknown, and then feed those images into the training set for the machine, so that future versions become smarter. Mobileye's customers use the company's AI technology to differentiate their products.

2.2 Predicting

2.2.1 Reducing loss

Example 1: Consumer credit risk management and fraud prevention

In Video 2, Professor Andrew Lo from MIT talks about the use of machine learning to assess consumer credit risk. Professor Lo examines how big data can be used to predict if a consumer is likely to be delinquent or defaulting on their credit card debt. He uses the example of an experiment in which machine learning techniques were used to choose multiple features and combine them to come up with the best forecast for delinquency and default.

Video 2: Using machine learning for consumer credit risk management and fraud prevention.

Key points and strategy

As Professor Lo explained, machine learning algorithms examine different parameters from those which traditional credit-scoring models examine. The insights gained from the machine learning techniques provide banks with a new lens that enables banks to predict consumer delinquency with a higher, finer-grained accuracy. Using this type of machine learning application, banks can pursue a cost reduction strategy that reduces their losses from non-payment.

Example 2: AXA

AXA is an example of how machine learning can be used for insurance underwriting. Read about how French insurance giant AXA is using machine learning to predict which automobile drivers are most likely to cause "large loss" accidents. (For the purposes of this program, there is no need to watch the video embedded in the linked article.)

Key points and strategy

As this example shows, AXA uses machine learning to predict which customers are most likely to cause accidents that would cost AXA US$10,000 or more. Such predictions support a focus-based strategy, because they enable AXA to write policies only for lower-risk customers.

Example 3: PayPal

Financial company PayPal has been using machine learning for years in its fight against fraud. Read about how PayPal uses machine learning for early fraud detection.

Key points and strategy

PayPal uses machine learning to support a low-cost strategy of being efficient in customer service while avoiding fraud and customer inconvenience. Specifically, PayPal uses machine learning to determine whether a visitor to the site is a trustworthy customer. If the user seems suspicious, the system will ask for additional verification. The combination of neural networks, vast quantities of data, and deep learning have greatly improved fraud detection, but the machines cannot do it alone; people need to decide which data is relevant for the machine to use.

2.2.2 Personalizing product offerings

Example 1: Netflix

This article explores how the personalization algorithm developed by Netflix predicts what to recommend to each user. Click on the PDF icon on the webpage to download the full text.

Key points and strategy

Pursuing a strategy of differentiation, Netflix combines machine learning with human curation to tailor its offerings precisely to the tastes of each individual customer.

Example 2: Stitch Fix

Explore how clothing company Stitch Fix uses machine learning and “humans in the loop” to select clothing that a specific customer might like. The company is subscription-based and not only suggests items, but also ships them to the customer. To be successful, the recommendations must be so good that customers are willing to not just keep the clothing they receive (it would be fairly expensive for Stitch Fix if there were many returns), but to continue their subscription to Stitch Fix.

Key points and strategy

Like Netflix, Stitch Fix uses machine learning to pursue its strategy of differentiation, by personalizing its product offerings for each customer. Furthermore, human stylists and the company's algorithms work hand in hand. The algorithms augment the human stylists' productivity by doing tedious tasks such as matching client measurements to different brands and products. The stylists meanwhile read the personal notes that customers have sent and analyze their Pinterest boards to determine each customer's nuances of style and taste.

Optional reading:

Read more about Stitch Fix and how the organization uses machine learning and relies on the human-machine partnership to offer personalized product offerings for customers.

2.2.3 Improving product performance using better predictions

Example 1: Digital Cognition Technologies

In Video 3, Professor Randall Davis from MIT CSAIL illustrates how Digital Cognition Technologies (DCT) uses machine learning to improve cognitive testing. Professor Davis talks about their product, the DCTclock, which is used to screen people for potential cognitive issues.

Video 3: How machine learning is used for cognitive testing.

Key points and strategy

As Professors Malone and Davis discussed, the machine learning aspect of the DCTclock can be generalized to many different applications of a screening test to provide early indication of a potential problem. In the case of the DCTclock, the machine looks for patterns that indicate cognitive impairment. The machine was trained using an expert's knowledge (a neuroscientist) about the features that are early indicators of potential problems. The machine learned to find the patterns that indicate the early onset of impairment. Human and computer work together in that the machine can suggest which patients a clinician may want to watch or follow up with, and the clinician then provides the diagnosis and treatment therapy. The same process applies to the early screening of any kind of problem, such as a factory machine or a jet engine. Computers can alert people to an engine that may need repair, for example. The people then use their expertise to pinpoint the cause of the upcoming problem and take preventative action. Systems like these can support a differentiation strategy by allowing for more accurate predictions about future problems.

Optional reading:

Read about how machine learning techniques are applied to the DCTclock, which tests cognitive abilities.

Disclosure:

Professor Malone made a small investment in Professor Davis's company, Digital Cognition Technologies, and therefore has a small ownership share in the company.

Example 2: Cogito

In Video 4, Professor Alex Pentland from MIT talks about honest signals and how they can be used to interpret human intention and emotion. He talks about the "second language" that people use, consisting of body language and signaling behavior, as a result of some of our basic neural processes. Professor Pentland uses the example of Cogito's software, which has been used to analyze conversations between call center agents and clients in real time. Computers and people work together, as the machine learning system interprets whether the conversation is going well or not, and then provides feedback to the agent to adjust their behavior for a more engaging and effective conversation.

Video 4: Machine learning and "honest signals".

Key points and strategy

Professor Pentland's research has found that "unspoken" language – the nonlinguistic signaling of interest, attention, dominance, and so on – accounts for 40–50% of the outcome of a conversation. People tend to focus consciously on words, so it is more difficult for them to tap into these signals, but computers can help them out. Professor Pentland used supervised machine learning in creating Cogito to "listen" to these unspoken features and predict how the conversation between a customer service representative and a customer is going. Cogito's software supports a differentiation strategy by facilitating higher-quality customer service interactions by letting representatives know if the customer is paying attention or is getting angry, for example.

3. Additional optional resources

Note:

The following resource is included for your enrichment and to provide deeper insight into the topics you have learned about in this casebook.

This resource from *MIT Technology Review* provides an introduction to deep learning.

4. Conclusion

As a technology, machine learning is making daily life easier by automating tasks, sorting through data with a level of speed and accuracy superior to that of humans, and making connections between data to enhance and personalize the online experience. It is being applied to a wide range of business problems to deliver tangible business value and will continue to transform how people live and work.