

Carnegie Mellon University Africa
Certificate I: Understanding AI and Machine Learning in Africa

Course AIML02: AI and Machine Learning in Africa

Module 3: AI Business Strategy

Lecture 4: The Future of AI Will Be About Less Data, Not More

Welcome back to Module 3 of AI and Machine Learning in Africa, a course which provides an overview of the relevance of AI and machine learning to Africa and their potential to solve economic and social problems.

In the final lecture of this module on AI business strategy, we look at the dependence of machine learning, especially machine learning that uses deep neural networks, on large training data sets. We will suggest that future AI systems will require less data and will instead exploit knowledge-based reasoning and probabilistic inference.

These are the main points of an article by James Wilson, Paul Daugherty, and Chase Davenport: "The Future of AI Will Be About Less Data, Not More".

We begin by looking at some of the issues with bottom-up data-driven machine learning and the alternative of human-like top-down reasoning.

We then look at the four areas where development over the next five years will lead to a likely increase in importance of top-down reasoning and artificial general intelligence. These are the use of efficient reasoning, the exploitation of expertise, the use of common sense, and the ability to perform probabilistic inference.

We finish up by summarizing what we have covered and providing the details of the article that you should read to consolidate what you have learned.

After this lecture, you should read the original article, take notes, and then go through this lecture again to consolidate the messages in the article.

We have three learning objectives, so that, after studying the material covered in this lecture, you should be able to do the following.

1. Explain why the future of AI may be based less on bottom-up data-driven deep learning and more on top-down reasoning.
2. Give examples of weaknesses in data-driven deep learning.
3. Identify four areas where developments will occur in the next five years.

Slide 1 Welcome back to Module 3 of AI and Machine Learning in Africa, a course which provides an overview of the relevance of AI and machine learning to Africa and their potential to solve economic and social problems.

In the final lecture of this module on AI business strategy, we look at the dependence of machine learning, especially machine learning that uses deep neural networks, on large training data sets. We will suggest that future AI systems will require less data and will instead exploit knowledge-based reasoning and probabilistic inference.

These are the main points of an article by James Wilson, Paul Daugherty, and Chase Davenport.

"The Future of AI Will Be About Less Data, Not More"

Slide 2 The article by Wilson, Daugherty, and Davenport opens by stating that

"Companies considering how to invest in AI capabilities should first understand that over the coming five years applications and machines will become less artificial and more intelligent."

Slide 3 Recall this diagram from the first lecture in Module 1 of AIML01 Artificial Intelligence – Past, Present, and Future, showing the three primary areas of AI:

Machine learning, including deep learning,

Knowledge-based systems, which emulate the problem-solving and reasoning skills of a human expert, and

Probabilistic Bayesian learning, which deal with uncertainty and, using probabilistic inference, draw conclusions that are the most likely to be correct when solving problems.

Slide 4 Wilson, Daugherty, and Davenport contend that future AI systems will rely less on bottom-up machine learning algorithms that require massive training data sets

and they will rely more on top-down reasoning and approaches that are more aligned with the way humans approach problem solving, i.e.,

knowledge-based systems and Probabilistic Bayesian learning.

Slide 5 This general reasoning ability will allow AI to be applied more broadly than ever and it will create opportunities in businesses where data-driven machine learning isn't well suited.

Slide 6 Recent advances in AI have been driven by machine learning and deep neural networks, building up systems from the bottom by training them on "mountains of data".

But these systems have serious limitations.

They have trouble handling "edge" cases, that is, situations for which there is little data.

Slide 7 They also have problems in unexpected or unusual situations.

A driverless car that can handle crosswalks, pedestrians, and traffic has trouble dealing with anomalies such as children dressed in unusual Halloween costumes, wandering across the street after dusk.

Slide 8 The iPhone X's facial recognition system doesn't recognize "morning faces": a user's puffy, haggard look on first awakening.

Turn an image upside down or slightly alter it and the network may misidentify it.

Slide 9 There are practical problems: not every company has the volume of data necessary to train a deep neural network.

There are privacy problems: using huge amounts of citizens' data leads to more government action such as the European Union's General Data Protection Regulation (GDPR), which imposes stringent requirements on the use of individuals' personal data.

And there are issues with transparency and explainability: these systems are black boxes and it's not always clear how the outputs – the actions or decisions – are derived from the input data.

This leaves them open to manipulation by people with bad intentions, people who might want to spread disinformation.

Slide 10 In the future, however, we will have top-down systems that don't require as much data and are faster, more flexible, and, like humans, more innately intelligent.

When planning their AI strategy, companies should look for developments in four areas:

1. More efficient robot reasoning.
2. Ready expertise.
3. Common sense.
4. Making better bets.

Slide 11 When robots have a conceptual understanding of the world, as humans do, it is easier to teach them things, using far less data.

By drawing on computational neuroscience, Vicarious, a startup based in California, is working to develop artificial general intelligence for robots, enabling them to generalize from few examples.

Slide 12 Modeling what a human expert would do in the face of uncertainty and little data, top-down artificial intelligence can beat data-hungry approaches for designing and controlling many varieties of factory equipment.

Slide 13 For example, Siemens is using top-down AI to control the highly complex combustion process in gas turbines, where air and gas flow into a chamber, ignite, and burn at temperatures as high as 1,600 degrees Celsius.

Slide 14 Many organizations are working developing AI systems with common sense: the ability to understand everyday objects and actions, communicate naturally, handle unforeseen situations, and learn from experiences.

Slide 15 What comes naturally to humans, without explicit training or data, is extremely difficult for machines.

so the Allen Institute for Artificial Intelligence (AI2) is developing a portfolio of tasks against which progress can be measured.

Slide 16 In DARPA's Machine Common Sense (MCS) program, researchers will create models that mimic core domains of human cognition, including "the domains of *objects* (intuitive physics), *places* (spatial navigation), and *agents* (intentional actors

Slide 17 Researchers at Microsoft and McGill University collaborated on the development of a system that can deal with the ambiguities in natural language, a problem that requires different forms of inference and knowledge.

Slide 18 Humans routinely, and often with little apparent effort, consider probabilities of different options and act on the likeliest, even with relatively little previous experience.

Slide 19 Machines are now being taught to mimic such reasoning through the application of Gaussian processes: probabilistic models that can deal with extensive uncertainty, act on sparse data, and learn from experience.

They are sometimes referred to as Bayesian networks, Bayes nets, belief networks, decision networks, and probabilistic graphical models. We met this approach in AIML01, Module 1, Lecture 1, and Module 2, Lecture 3.

Slide 20 Alphabet, Google's parent company, launched Project Loon to provide internet service to underserved regions of the world through a system of giant balloons hovering in the stratosphere.

Slide 21 Since the movement of the balloons will depend on stratified and highly variable winds, their navigational systems use Gaussian processes to predict the altitude the balloons need to go to adjust their position.

Recall we met Gaussian processes earlier in the previous module of this course, in the application case study in agriculture (Lecture 3) and in the application case study on conservation (Lecture 6).

Slide 22 Each balloon then moves into a layer of wind blowing in the right direction, arranging themselves to form one large communication network.

Slide 23 Gaussian processes don't require massive amounts of data to recognize patterns; the computations required for inference and learning are relatively easy, and if something goes wrong its cause can be traced, unlike the black boxes of neural networks.

Slide 24 Although these advances are relatively recent, they hark back to the very beginnings of AI in the 1950s, when the pioneers of AI began to develop models that mimic human intelligence.

Recall from AIML01, Module 1, Lecture 2, the proposal for the Dartmouth Workshop that had the goal of programming a machine to "solve [the] kinds of problems now reserved for humans".

Today, top-down human-level AI is being reborn. As its great promise begins to be fulfilled, smart companies will put their money where the mind is.

To summarize, we borrow from the "Takeaways" in the target article.

1. "As AI develops, it will rely less on bottom-up big data and more on top-down reasoning that resembles the way humans approach problems and tasks."
2. "This will enable us to apply AI more broadly than ever, creating opportunities for early adopters even in businesses and activities to which AI previously seemed unsuited."
3. "Until recently, most AI advanced through deep learning and machine learning, building up systems by training them on mountains of data. But these data-hungry networks have serious limitations and difficulty handling situations where little data exists."
4. "To craft a vision of where AI is heading in the next several years, and plan investments and tests accordingly, companies should look for the following developments: more efficient robot reasoning ..., ready expertise ..., common sense ..., making better bets ..."

Here is the article on which this lecture is based. Please take the time to read it and then review this lecture again.

Wilson, H. J., Daugherty, P., and Davenport, C. (2019). The Future of AI Will Be About Less Data, Not More, in Insights You Need from Harvard Business Review - Artificial Intelligence, Harvard Business School Publishing Corporation.

<https://hbr.org/2019/01/the-future-of-ai-will-be-about-less-data-not-more>