

Carnegie Mellon University Africa
Certificate Program in AI and Machine Learning in Africa

Course AIML01: Artificial Intelligence – Past, Present, and Future

Module 1: What is AI, where did it come from, and where is it taking us?

Lecture 2: The Early Years of AI

Welcome to Lecture 2 of Module 1 in which we will study the early years of AI.

This will give us the foundation we need to understand how AI developed and why it developed the way it did, including the increasingly important role that machine learning played. It will also help us understand how AI and machine learning are likely to develop in the future.

In this lecture, we explore the birth of artificial intelligence and its roots in both computationalist models and cybernetics, inspired by behaviorist and cognitivist psychology, respectively. The former gave rise to what would become known as GOFAI: good old-fashioned AI, focusing on human-level intelligence, symbolic representations and reasoning, and rational logic. The latter impacted strongly on research in connectionist approaches and neural networks. We will explain the difference between strong AI and weak AI. We then take a brief look at the so-called AI winter when people became somewhat disillusioned by the lack of success in achieving human-level intelligence. We say much more about this in subsequent lectures.

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

1. Identify the events that marked the beginning of AI and the people involved in the effort.
2. Compare AI's roots in cybernetics, connectionism, and computationalism, and explain what is meant by each term.
3. Explain the difference between strong AI and weak AI, and what is meant by an AI winter.

- Slide 1 Welcome to Lecture 2 of Module 1 in which we will study the early years of AI.
- This will give us the foundation we need to understand how AI developed and why it developed the way it did, including the increasingly important role that machine learning played.
- It will also help us understand how AI and machine learning are likely to develop in the future.
- Slide 2 John McCarthy, then a professor at Dartmouth College and soon to join MIT, coined the term Artificial Intelligence in 1955 in a proposal for a two-month summer project.
- It was co-written by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon
- Slide 3 The organizers based their proposal on the conjecture that intelligence and learning can, in principle, be described so precisely that it can be simulated by a machine.
- Slide 4 The scope of the proposal was very ambitious: from programming a computer to simulate an intelligent machine and getting it to use language, form abstractions and concepts, and improve itself.
- The proposal specifically identified the goal of programming a machine to "solve [the] kinds of problems now reserved for humans".
- Much of the work that ensued from this workshop focused on logical rule-based reasoning and it's sometimes forgotten that the workshop also considered "Neuron Nets", what we now call neural networks, topic we will say more about in a moment.
- Slide 5 The summer project itself became known as the Dartmouth Workshop and it was held at Dartmouth College, New Hampshire, from June 17 to August 16, 1956 (the dates vary slightly in different accounts).
- It is considered by many to be the event that founded the field of AI.
- Slide 6 Apart from the organizers, John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, the workshop was attended by several people who had a profound impact on the development of AI over the next fifty years, including Herbert Simon, Allen Newell, Oliver Selfridge, John Holland, and Ray Solomonoff, among others.

Slide 7 The proposal specifically identified the goal of programming a machine to "solve [the] kinds of problems now reserved for humans".

Slide 8 The essential position of AI at this time was that intelligence — both biological and artificial — is achieved

by computations performed on internal symbolic knowledge representations

an approach referred to as *computationalism, or the computational model of mind*,

grounded in cognitivist psychology,

and often referred to as GOFAI: good old-fashioned artificial intelligence.

It had a strong emphasis on logic and symbolic knowledge representation.

Later in the course, we will see how this position is captured formally in Allen Newell's and Herbert Simon's Physical Symbol Systems Hypothesis.

Slide 9 However, symbolic computational reasoning wasn't the only approach that researchers had taken to the problem of understanding and emulating intelligent behavior.

A field known as cybernetics, using mathematical models of self-organization, regulation, and control, had been explored since the early 1940s by researchers such as Norbert Wiener, W. Ross Ashby, Grey Walter, and Claude Shannon, among others.

Slide 10 Norbert Wiener, a professor of mathematics at MIT, is widely considered to be the originator of cybernetics.

He defined it as "the scientific study of control and communication in the animal and the machine".

Slide 11 W. Ross Ashby was a British pioneer in the fields of cybernetics and systems theory.

His books *Design for a Brain* (1952) and *An Introduction to Cybernetics* (1956) were very influential.

Slide 12 Neurophysiologist Grey Walter built two cybernetic tortoises to understand the functions of the brain

He nicknamed them ELSIE and ELMER

Slide 13 They could roam around a room, find a charging station, and recharge themselves, not unlike what a Roomba robot vacuum cleaner does, although using a completely different technology.

Slide 14 Claude Shannon, the founder of information theory, created a robotic maze-solving mouse known as Theseus

Slide 15 It is one of the first examples of machine learning.

Here we see the path Theseus took while learning a maze pattern and the direct path taken on its second trip through the same maze

Slide 16 This picture of W. Ross Ashby, W. Grey Walter, and Norbert Wiener at the 1951 Congress on Cybernetics, also features Warren McCulloch, co-author of a highly-influential paper published in 1943 entitled "A logical calculus of ideas immanent in nervous activity".

As well as being a seminal work in cybernetics, this paper is also regarded as the foundation for artificial neural networks and connectionism.

Slide 17 To summarize, then, Norbert Wiener, W. Ross Ashby, and Walter McCulloch contributed to the creation of a new field of cybernetics that would form the foundation of modern AI and machine learning, although one that differs significantly from the AI envisaged by those at the Dartmouth Workshop.

We will see later in the course that much of the recent success in AI and machine learning has more in common with cybernetic and neural network research than it does with the human-level focussed AI of John McCarthy, Herbert Simon, Allen Newell, and others. Future successes may yet differ.

Slide 18 Both Walter's and Shannon's robots built on behaviorist psychology

using associative and reinforcement learning in relatively simple neural networks

rather than focusing on internal models and symbolic computation

In associative learning, the link between two events is strengthened when they occur together.

In reinforcement learning, the link is strengthened when they yield a reward

We will say more about these types of learning later in the course

Slide 19 Let's pause here to say something about neural networks and the way they process information.

They do so by propagating it through an interconnected layered network of relatively simple processing units called artificial neurons.

These are extremely simplified versions of the neurons in biological brains: inspired by them rather than models of them.

Slide 20 These artificial neurons effectively weight the inputs, aggregate them in some manner, e.g. by summing them, modifying the result, e.g. by amplifying it or attenuating it, and then send this modified result to other processing units typically in the next layer

Slide 21 Artificial neural networks that propagate results to the next layer are called feed-forward networks

Networks that also propagate results back to previous layers are called recurrent neural networks

Slide 22 The simple one-layer networks developed in the 1960s are also called perceptrons (the input layer doesn't count as it doesn't do any processing)

One perceptron can perform binary classification; that is, it can distinguish between two classes (provided their features can be separated by a straight line)

Several perceptrons can distinguish between several classes (again, provided their features can be separated by a straight line)

They are still in use today, in a modified form

We will say more about perceptrons later in the course, highlighting the important role they played in the development of artificial neural networks and machine learning.

Slide 23 This approach is often referred to as connectionism

It progressed in parallel with the computationalist approach over the next sixty years and more

We'll say more about computationalist symbolic AI in Module 2, Lecture 1

We'll say more about connectionist AI in Module, Lecture 2.

Slide 24 We have already noted that, from the outset, symbolic computationalist AI targetted the development of machines with human-level intelligence

Just like humans, these would be

versatile (that is. able to tackle different types of task or problem)

flexible (that is, able to use different strategies)

and robust (that is, able to deal with unexpected circumstances)

This form of AI became known as strong AI

It proved very difficult to achieve.

Slide 25 Consequently, AI techniques began to be applied in more limited domains

with stronger constraints and a narrower focus

thus not requiring as much versatility and flexibility,

Targetting instead strong performance and robustness

This approach became known as weak AI

Slide 26 Even so, performance on more challenging problems was disappointing throughout the 1970s and 1980s

There were a variety of reasons for this, which we will cover in the next module.

The upshot, though, was a period called the AI winter when AI fell out of favour.

Actually, there were two AI winters, one earlier one in connectionist AI in the 1970s and one in AI generally in the 1980s, despite a resurgence in interest in connectionism and neural networks. Again, we'll say more about this in the next module.

Nevertheless, research on statistical techniques and on neural networks continued and this laid the foundations for the end of the AI Winter, the topic of the next lecture, the last in this module.

To summarize:

- John McCarthy coined the term Artificial Intelligence in a proposal for a summer project at Dartmouth College in 1956, attended by several pioneers in the field.
- This led to the birth of computational AI, focusing on symbolic representations and logic-based rational reasoning, emulating human intelligence. It is often referred to as GOFAL: good old-fashioned AI.
- The focus on human-level intelligence, exhibiting versatility, flexibility, robustness, is often referred to as Strong AI. It proved harder than anticipated.
- Alternative approaches, often referred to as Weak AI, targeted more limited goals with greater success.
- AI also has roots in an alternative approach based on cybernetics, championed by Norbert Wiener, W. Ross Ashby, and others.
- Cybernetics focusses on self-organization, regulation, and control, rather than internal symbolic representations of the world in which the intelligence agent is situated.
- This approach also targetted the use of neural networks and connectionist models of intelligent behaviour.
- Lack of early success gave rise to a decrease in interest in AI during a period known as the AI winter.

In fact, there were two AI winters, one in the early 1970s that mostly affected connectionist approaches, and one in the 1980s that impacted both symbolic and connectionist approaches.

Recommended Reading

Here is some reading. It will help you understand how the two strands of research contributed to the development of AI and how machine learning emerged as the key to its current success.

Jordan, M. I. (2019). Artificial Intelligence — The Revolution Hasn't Happened Yet. Harvard Data Science Review, 1(1).
<https://doi.org/10.1162/99608f92.f06c6e61>

References

Here are some of the references cited to support the main points in what we covered today.

McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine*, 27(4), 12.

<https://doi.org/10.1609/aimag.v27i4.1904>.

Moor, J.H. (2006). The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years. *AI Magazine*, 27, 87-91.

<https://ojs.aaai.org//index.php/aimagazine/article/view/1911>

Marsh, A. (2020). Meet the Roomba's Ancestor: The Cybernetic Tortoise, *IEEE Spectrum*, February.

<https://spectrum.ieee.org/tech-history/space-age/meet-roombas-ancestor-cybernetic-tortoise>