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

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

Module 2: The Nature of AI
Lecture 1: Symbolic AI and GOFAI

Welcome to Module 2 in which we look more deeply at the technical foundations of AI and machine learning. In this lecture, we discuss the importance of symbolic AI, also known as good old-fashioned AI, or GOFAI for short. We explain what physical symbol system are, highlighting their syntactic nature, and indicate how symbol expressions are created using logic formalisms. We go on to explain how symbol systems solve problems by using a process known as heuristic search. This leads us to the Physical Symbol Systems Hypothesis which links symbol systems and intelligence, both artificial and natural. We then discuss the strengths and limitations of symbolic AI systems. We will finish up by summarizing what we have covered and identifying the articles that you should read to consolidate what you have learned.

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

1. Explain what is meant by symbolic AI with reference to physical symbol systems and associated logic formalisms.

2. Compare propositional logic and first-order logic (also known as predicate calculus).

3. Explain the link between physical symbols systems, heuristic search, and intelligence.

4. Identify some of the strengths and weaknesses of AI.

5. Explain the relationship between symbolic AI and cognitive architectures.

Slide 1 Welcome to Module 2 in which we look more deeply at the technical foundations of AI and machine learning. In Lecture 1, we begin with symbolic AI and GOFAI: good old-fashioned AI.

Slide 2 One of the key historical, methodological, and epistemological approaches to AI is that of “Symbolic AI”, often referred to as GOFAI (Good Old-Fashioned AI).

Historical, in the sense that if forms one of the cornerstones of AI.

Methodological, in the sense that it is still an effective approach to the implementation of AI.

Epistemological, in the sense that it addresses what kinds of facts (or knowledge) about the world are available to an AI agent, how these facts can be represented, and what rules allow valid conclusions to be drawn from these facts.

Slide 3 Symbolic AI has its origins in the 1950s, specifically at the 1956 Dartmouth Workshop

It constituted the primary, classical approach in the first 30 years of AI research,

before the second AI Winter and the advent of connectionist AI and machine learning.

Slide 4 The term “symbolic” refers to the fact that AI algorithms and programs are based on a set of symbols and symbol manipulation processes.

Slide 5 Two of the founders of symbolic AI, Allen Newell and Herbert Simon, both at Carnegie Mellon University, proposed the concept of a **Physical Symbol System**.

Slide 6 Defining it as “a set of entities, called symbols,

which are physical patterns that can occur as component of another type of entity called an expression (or symbol structure)”

(Newell and Simon 1976: 116).

Slide 7 These symbols are purely formal and meaningless entities,

though in practice they are normally interpreted by the programmer to have a particular semantic content such as words, numbers, images, and actions.

Slide 8 Another way of saying this is that symbols and symbol systems are purely syntactic entities

- just like letters & words and a grammar that stipulates how thy can be combined in valid sentences -

but they have no semantic content

- like the characters written in a foreign language, e.g., Chinese, that you don't speak or understand -

The problem of attaching semantic meaning to syntactic symbols by linking them with entities in the environment is known as the symbol grounding problem.

We will return to the problem of attaching meaning to these syntactic, formal symbols and expressions later in the course

Slide 9 The symbolic expressions are created using logic formalisms,

such as propositional logic

A sentence in propositional logic is either true or false and comprises

Propositional symbols (which can be true or false)
Boolean connectives (AND, OR, NOT, ...)

Propositional logic describes facts about specific instances of things

For example, whether an object is both red and round

Slide 10 Or the more expressive first-order logic that deals with objects and relations among them

It can be used to express facts about some objects or all objects, and it can represent general laws or rules.

First-order logic is also referred to has first-order predicate calculus (or just predicate calculus).

Slide 11 The symbol expressions can also be arranged in IF-THEN rules called productions

For example, “IF apple, THEN eat".

Slide 12 In some symbolic systems, such as semantic networks, each node has a symbol (“Red”, “Apple”, “Fruit)

with links having a label for the semantic relationship between node (e.g., “IS A” or “HAS”)

and hierarchical relationship between nodes.

Slide 13 A collection of symbolic structures for a specific domain constitutes a knowledge base that is used by the system to reason about the problem.

Slide 14 In general, symbol systems solve problems by using processes of heuristic search

where the search for the optimal link between the problem definition and its solution must be guided by heuristics,

i.e., rules of thumb that are helpful in guiding the search toward the solution in an optimal way (or, at least, an efficient way).

Slide 15 AI heuristic search and planning algorithms are widely used today for scheduling and logistics, for data mining, for games, for searching the web, and for planning in robotics

Slide 16 An important aspect of the GOFAI approach is the idea that symbolic systems can model human intelligence.

In fact, Newell and Simon (1976) proposed the **Physical Symbol Systems Hypothesis**, which states that

“A physical symbol system has the necessary and sufficient means for general intelligent action”

Slide 17 The solutions to problems are represented as symbol structures

A physical symbol system exercises its intelligence in problem-solving by effective and efficient search.

Generating and progressively modifying symbol structures until it produces a solution structure.

Slide 18 A classic example of a GOFAI system is an expert system, i.e., a program that represents the knowledge of a human expert in a specific domain,

using a set of IF-THEN production rules,

and which can be used to offer advice to non-experts or to provide solutions to experts.

Mycin was one of the first expert systems to support medical doctors in the diagnosis and treatment of infectious diseases.

Today, we refer to them as knowledge representation and reasoning systems.

Slide 19 Nowadays expert systems, that is, knowledge representation and reasoning systems, have been developed in a wide range of domains, in commercial, educational, medical, and military applications,

with some capable of highly complex planning on the order of tens of thousands of search steps.

Slide 20 The major strengths of GOFAI are the ability to model hierarchical and sequential tasks, such as language processing, problem solving and games, and to represent knowledge and reason about this knowledge using logical inference.

The diagram on the right shows the structure of a knowledge base – also known as an ontology – representing the objects that robots manipulate in everyday activities.

Slide 21 Some limitations of GOFAI systems are that these AI programs can be brittle

i.e. that they can produce wrong and nonsensical decisions when there is missing or contradictory knowledge

they are subject to the frame problem

i.e. the problem of representing what remains unchanged as a result of an action or an event

Slide 22 Other limitations are that the symbol grounding problem

That is, the problem of attaching semantic meaning to syntactic symbols by linking them with entities in the environment

and they cannot learn new knowledge.

Slide 23 This, as well as the initial strong claims about the power of symbolic systems to deal with general intelligence and any problem domain,

led to an AI Winter in the 1980s,

and the subsequent development of connectionist and machine learning approaches.

Slide 24 However, significant achievements of GOFAI include

The widespread use of commercial expert systems, also known as knowledge representation and reasoning systems,

their essential role in games industry (to control the intelligent behaviour of the virtual agents)

including the historical victory of the IBM Deep Blue system in 1997 in beating the chess world champion Gasparov,

and IBM Watson's victory in 2011 over two human champions in the *Jeopardy!* TV show

Slide 25 GOFAI also made a major contribution to the field of cognitive science in the guise of Unified Theories of Cognition (UTC).

Slide 26 A unified theory of cognition is a theory that covers all the bases in modelling human cognition and intelligence – from attention to learning – from all relevant perspectives, including psychology, neuroscience, and computer science.

The concept was introduced by Allen Newell at Carnegie Mellon University in his book on Unified Theories of Cognition in 1990.

Slide 27 He also introduced the term "cognitive architecture".

Here's a picture of his proof-reading notes three years before the book was published.

A cognitive architecture is a candidate Unified Theory of Cognition.

Slide 28 This means the cognitive architecture integrates into a single system all the different core cognitive abilities

The ability to perceive
To pay attention to the things that matter
To select appropriate actions
To remember what it needs to perform these actions
To learn
To reason based on what it knows
To reason about its reasoning (we call this meta-reasoning or metacognition)
To anticipate what will happen in the future (we call this prospection)

This allows the agent to act or behave in a flexible manner, and to use the current context to guide its behavior.

It also allows it to use prospection to identify what it should do to achieve its goals and guide the action as it's doing it.

Slide 29 Cognitive architectures are often depicted as diagrams showing how the various modules are connected.

However, a cognitive architecture must also identify the formalisms – or computational theories – that provide the foundation of its representations and the processes that operate on those representations.

Slide 30 There are three paradigms of cognitive science:

Cognitivist or symbolic

Emergent or sub-symbolic, often implemented using connectionist neural networks

and Hybrid, combining both symbolic and sub-symbolic processes.

We cover connectionist neural networks in the next lecture AIML01-02-02

Slide 31 While the term cognitive architecture was invented by one of the pioneers of AI and cognitive science, Allen Newell, it has

[click]

been subsequently been borrowed by the borrowed by the other paradigms.

Slide 32 Allen Newell's cognitive architecture is called Soar, one of the most influential and oldest cognitive architectures.

Today, it is categorized as a hybrid architecture, but it started out as a cognitivist or symbolic cognitive architecture.

Slide 33 The development of Soar began in the early 1980s as a collaboration between John Laird and his Ph.D. adviser, Allen Newell, at Carnegie Mellon University. Paul Rosenbloom joined the effort after completing his Ph.D. at CMU in 1983, the same year John Laird completed his thesis.

Both John Laird and Paul Rosenbloom continued to lead the development of Soar after they left CMU in 1984.

John Laird published the definitive book on Soar in 2012.

Soar is now maintained and developed by John Laird's research group at the University of Michigan.

Slide 34 One of the exercises in this lecture is to watch a video about Soar, presented by John Laird.

Slide 35 ACT-R is another iconic cognitive architecture.

Slide 36 Like Soar, ACT-R has been around for a long time, and it too is still being developed by John Anderson and Christian Lebiere, both at Carnegie Mellon University, among others.

It is also a production, rule-based, system but in this case, it is modelled on the high-level structure of the human brain.

To summarize:

- Symbolic AI, also known as good old-fashioned AI, is an important branch of AI, both historically and in contemporary terms.
- Implemented in different logic formalisms, it provides the basis for symbolic knowledge representation and reasoning
- Problems are solved using symbol systems by heuristic search
- According to the physical symbol systems hypothesis, all intelligent systems – natural and artificial – are physical symbol systems
- Symbolic AI gave rise to many important developments, including expert systems, now referred to as knowledge representation and reasoning systems, and symbolic cognitive architectures.

Here is some recommended reading. It will help you understand the importance of good old-fashioned AI, historically and today.

Boden M. (2014). GOFAI. In The Cambridge Handbook of Artificial Intelligence Frankish, K and Ramsey, W, Editors. https://www.cambridge.org/rw/academic/subjects/philosophy/philosophy-mind-and-language/cambridge-handbook-artificial-intelligence

Cangelosi A and Vernon D (2022). "Artificial Intelligence: Powering the Fourth Industrial Revolution", in *EPS Grand Challenges: Physics for Society at the Horizon 2050,* coordinated by the European Physical Society. http://vernon.eu/publications/2022_Cangelosi_Vernon.pdf

Vernon, D. (2022), Cognitive Architectures, in Cognitive Robotics, Cangelosi, A. and Asada, M. (Eds.). Cambridge, MA: MIT press, 2022. https://direct.mit.edu/books/oa-edited-volume/5331/chapter-standard/3809838/Cognitive-Architectures

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

Franklin S. (2014). History, motivations, and core themes. In The Cambridge Handbook of Artificial Intelligence Frankish, K and Ramsey, W, Editors. https://www.cambridge.org/rw/academic/subjects/philosophy/philosophy-mind-and-language/cambridge-handbook-artificial-intelligence

Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search (ACM 1975 Turing award lecture), Communications of the ACM, 19, 3, 113-126. https://dl.acm.org/doi/abs/10.1145/360018.360022