FACTFILE:
GCE DIGITAL TECHNOLOGY

UNIT A2 1: INFORMATION SYSTEMS

Rewarding Learning

Artificial Intelligence (AI)

Learning Outcomes

Students should be able to:

- Explain what is meant by artificial intelligence (AI).
- Explain the significance of the Turing Test in defining what is meant by AI.
- Describe the main features of neural network modelling.

Content in Artificial Intelligence

- An outline of what is meant by AI
- Examples of applications of AI
- What is the Turing Test?
- Defining AI using the Turing Test
- Neural network modelling and AI
- Main features of a neural network
- How does a neural network learn?

What is meant by Artificial Intelligence (AI)

Artificial intelligence refers to the study of machines to model the types of intelligence or creativity normally exhibited by people. Reasons for the continued development of AI include the need to make machines more capable in addition to a need to help improve human intelligence and cognitive behaviour (cognitive science). Cognitive science is concerned with understanding the human mind and its thought processes and some scientists who study this area use computer models of information processing to help explain how the human mind functions.

Applications of AI

Some activities carried out by computer with AI include:-

- Knowledge-based or expert systems.
- Image processing and visions.
- Speech recognition and natural language processing.
- Machine learning.

The study of AI dates back to the 1950's when famous mathematician Alan Turing proposed the Turing Test, used to help determine if computers could mimic human-like behaviour.

What is the Turing Test?

The Turing Test is used to assess the ability of a machine to exhibit intelligent behaviour which could not be distinguished from that of a human. In its simplest form, a human simultaneously asks questions of a computer and human. If the person asking the questions cannot distinguish between the computers and the humans, then the machine has passed the Turing test and it can be classified as 'intelligent'.

Defining AI using the Turing Test.

The conversation between the participants is through an interface and is text only, for example via a keyboard and screen (i.e. there is no speech output). Some critics have challenged the Turing Test as a measure of intelligence as the test does not assess the correctness of the responses given only how similar they are to a human response.

For many years the Turing test was considered the standard for assessing AI in machinery however in 2014 a chatbot known as Eugene Goostman succeeded in passing the Turing Test, hence calling the reliability of the test into question. The chatbot was designed to imitate a 13-year-old boy from the Ukraine in a conversation with a panel of judges at the Royal Society in London; 33% of the judges were convinced the chat bot was human. Some experts have however still discounted these results due to the fact that the limited life experience and vocabulary of a 13-year-old boy speaking a English as their second language could not be considered a true test; especially since the chatbot Eugene was only mimicking understanding and not problem solving or learning from the interaction.

Questions asked take the form of "The football would not fit in the bag because it was too big. What was too big? A) The football B) The bag?". A human could use common sense reasoning and intelligence to answer this question whereas a machine could not as they would be able to determine that the 'it' referred to the ball as opposed to the bag.

Most AI systems used today are still very domain specific and therefore what we expect them to be able to achieve is restricted. In the future, as AI machines become more general purpose and less context specific it will become increasingly more difficult to determine if we are dealing with machines or humans. This distinction will be especially difficult to make with the increased use of voice recognition system coupled with advances in natural language processing.

Neural network modelling and AI

Neural network models are computational models (generally software simulations) used in computer science to model a primitive brain. The model is designed to simulate lots of densely interconnected brain cells to support learning, pattern recognition and the decision making process in the same way a human being would.

Main features of a neural network

A typical neural network consists of a number of artificial neurons called units. The number of units could vary from a few dozen to millions depending on the complexity of the neural network. Units will be arranged in layers. Input units will receive information from the outside world. A series of hidden units (which represent the brain) will attempt to process input from the input units. Output units provide a response from the network representing any information it has learned following the processing of input. In most cases the neural networks are fully connected meaning that each units is connected to every other unit the layers on either side.

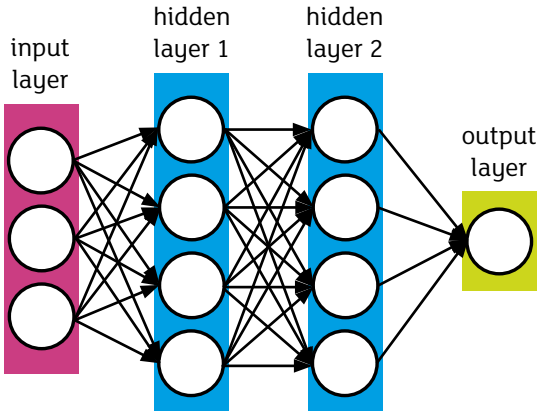

Each connection in the neural network is weighted and algorithms are used to calculate the weighted sum of any inputs into a node in order to generate an output value. A neural network does not have to be programed to learn explicitly; it learns by itself, much in the same way a human brain would learn.

How does a neural network learn?

Patterns of information flowing into a neural network will enter the network via the input units, triggering layers of hidden units until they arrive at the output units. Not all units fire all the time. Inputs into a unit are multiplied by the weightings of the connections they travelled along. The weighted inputs are added together and if the sum is above a given threshold value that unit will fire / trigger the units connected to it. This design is known as a **feedforward network**.

Feedback is important if learning is to take place in a neural network. The feedback process used to support learning in neural networks is known as **backpropagation**. During this process the output produced by the neural network is compared to the output it was meant to produce. The difference between the two values is used to modify the connections between the various units in the network (going backward from the output units, through the hidden layers to the input units). Over time this process of backpropagation reduces the difference between the actual and intended output, causing learning to occur.