

7 Things EEs Should Know About Artificial Intelligence

What exactly are the basic concepts that comprise AI? This “Artificial Intelligence 101” treatise gives a quick tour of the sometimes controversial technology.

Just what the devil is artificial intelligence (AI) anyway? You’ve probably been hearing a lot about it in the context of self-driving vehicles and voice-recognition devices like Apple’s Siri, Amazon’s Alexa, Google’s Assistant, and Microsoft’s Cortana. But there’s more to it than that.

AI has been around since the 1950s, when it was first discovered. It has had its ups and downs over the years, and today is considered as a key technology going forward. Thanks to new software and ever faster processors, AI is finding more applications than ever. AI is an unusual software technology that all EEs should be familiar with. Here is a brief introductory tutorial for the uninitiated.

AI DEFINED

AI is a subfield of computer science that involves making computers and electronic-based products more intelligent by mimicking the human brain. Intelligence is the ability to acquire knowledge from education and experience and apply it to solve problems. AI is especially useful in analyzing and interpreting masses of data and deriving real useful knowledge from it. From knowledge comes understanding that can be applied to make decisions or initiate action of some sort. Kind of along the lines of what some of these famous said:

“We are drowning in information but starved for knowledge.”—John Naisbitt from his book *Megatrends*.

“Knowledge is power.”—Hobbes

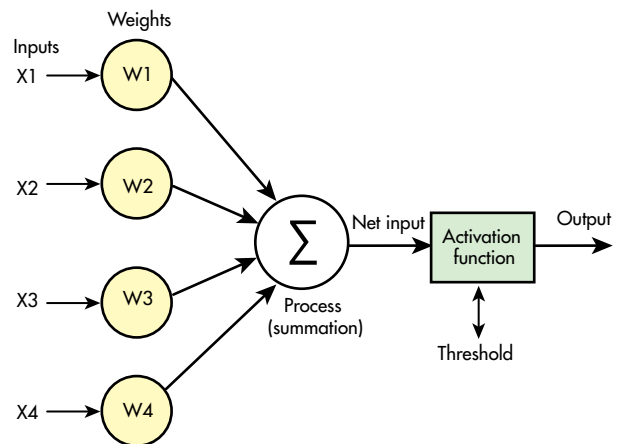
“I think therefore I am.”—Descartes

“The test of a first rate intelligence is the ability to hold two opposed ideas in the mind at the same time and still retain the ability to function.”—Francis Scott Fitzgerald

FIELDS OF STUDY

AI is a broad technology with many possible applications. It’s typically divided into special sub-branches. Here’s a brief summary of each:

- *General problem solving:* Problems with no known algo-



1. Key elements of an artificial neuron node are multiple weighted inputs, an output, and a threshold setting.

rithmic solution. Problems with ambiguity and uncertainty.

- *Expert systems:* Software that incorporates a knowledge base of rules, facts, and data derived from multiple individual experts. The knowledge base can be queried to solve problems, diagnose medical conditions, or provide advice.
- *Natural language processing (NLP):* NLP is used to derive understanding from natural text. Voice recognition is a part of NLP. Language translation is an application.
- *Computer vision:* Analyzing and understanding visual information (photos, video, etc.). Machine vision and facial recognition are examples. Used in self-driving cars and factory inspection.
- *Robotics:* Making robots more intelligent, self-operating, adaptive, and more useful.
- *Games:* AI is great at playing games. Computers already have been programmed to play and win at chess, poker, and Go.
- *Machine learning:* Procedures to allow a computer to learn from data or other inputs and make sense of the results

for application. Neural networks make up the framework for machine learning.

HOW AI WORKS

Conventional computers solve problems with algorithms. Sequences of instructions perform step-by-step procedures to reach a solution. The traditional forms of AI start with a knowledge base and an inference engine that uses various processes to access the knowledge base via a user interface. Useful results are obtained by some of these methods:

- **Search:** Search algorithms mine the knowledge base of facts or data organized into state graphs or trees. Search is a core AI method.

- **Logic:** Deductive and inductive reasoning is used to determine the truth or falsity of knowledge base statements. This involves both propositional and predicate logic.

- **Rules:** Rules are a series of “if-then” statements that can be searched to determine an outcome. Rules-based systems are referred to as expert systems.

- **Probability and statistics:** Some problems can be solved, and decisions made, by applying standard probability and statistics math.

- **Lists:** Some types of information can be organized into lists that become searchable for reasons of reaching a conclusion.

- Other knowledge forms are schemas, frames, and scripts, which are structures that encapsulate different types of knowledge. Search techniques seek answers with appropriate queries.


Traditional or legacy AI methods like search, logic, probability, and rules are considered the first wave of AI. These methods are still used and are good at perceiving knowledge and reasoning from it, especially for a narrow range of problems. Missing are the human traits of learning and abstracting decisions. These qualities are now available in the second wave of AI, thanks to neural networks and machine learning.

NEURAL NETWORKS

Most AI research and development today is based on the use of neural nets or artificial neural networks (ANNs). These networks are made up of artificial neurons that simulate the neurons in the human brain, which are responsible for our thinking and learning. Each neuron is a node in a complex interconnection that links many neurons to many others by way of connections call synapses. An ANN mimics this network.

Each node has multiple inputs that are weighted, along with an output and a threshold setting (Fig. 1). Such nodes are usually implemented in software, although hardware emulation is possible. A typical arrangement is three layers; an input layer; a hidden, processing or learning layer; and an output layer (Fig. 2). Some arrangements use back propagation to provide feedback to change input weights of some nodes as new facts are gleaned.

ANNs must be trained usually one layer at a time to provide



2. The illustration shows how artificial neurons are arranged for machine learning. Multiple hidden layers are used in deep learning.

a starting point. Initial weights are given as an example. Then, as new inputs occur, different results cause weights or thresholds to change as the ANN zeros in on a conclusion. The ANN learns from its inputs and stores the result.

MACHINE LEARNING AND DEEP LEARNING

Machine learning is a method of teaching a computer to recognize patterns. The computer or other hardware is “trained” with an example and then special programs are run to compare the input to the trained value. Typically, massive amounts of data are needed to train the software. Machine-learning programs are designed to learn automatically as they gain more knowledge and experience through new inputs.

Neural networks are commonly used for machine learning; however, other algorithms may be used. The software can then modify itself by improving its recognition ability based on new inputs. It’s now possible for some machine-learning systems to recognize patterns on their own without training, then modify themselves to further improve.

Deep learning is an extended case of machine learning. It, too, uses neural networks, but incorporates extra hidden layers of computing to further refine its ability. Massive training is required. Programmers can boost performance by playing with the weights of the interconnections.

Both machine learning and deep learning are widely used to analyze big masses of data, as well as in computer vision and voice recognition. And they can be applied to other areas such as medicine, law, and finance.

SOFTWARE OF AI

Almost any language can be used to program AI, but some languages are better than others. Special languages designed specifically for AI include LISP and Prolog. LISP, one of the oldest higher-level languages, processes lists. Prolog is based on logic. Today, C++ and Python are popular. Special software for developing expert systems is available.

Several large users of AI provide development platforms, including Amazon, Baidu (China), Google, IBM, and Microsoft. These companies offer pre-trained systems as a starting point for some common applications like voice recognition.

HARDWARE OF AI

Running AI software on a computer usually requires high speed and lots of memory. Some of today's processors are more than adequate, and multiple parallel processors can be an ideal solution for some applications. In addition, special processors have been developed for some applications.

Graphics processor units (GPUs) represent an instance of focusing the architecture and instruction set to a given use to optimize performance. One example comes from Nvidia, with its special processors for self-driving cars. Google has developed its own processors to optimize its search engines. In some cases, special logic in an ASIC or FPGA can implement a specific application.

CURRENT STATUS AND ACTIVITY

AI was once considered exotic software allocated to specific special needs. The requirement for high-speed computers with lots of memory once limited its use. Today, thanks to super-fast processors, multiple cores, and cheap memory, AI has become more mainstream. Google's search engines we all use daily are AI-based.

The current emphasis is clearly on neural nets and machine learning. While voice recognition and self-driving cars are in the spotlight, other key applications are emerging, such as facial recognition, drone navigation, robotics, medical diagnosis, and finance. Advanced military and defense applications (e.g., autonomous weapons) are also in the works.

The future of AI looks promising. Orbis Research says that the global AI market is expected to grow at a compound annual growth rate (CAGR) of greater than 35% by 2022. The International Data Corporation is also positive, saying that spending on AI is expected to balloon to \$47 billion in 2020, up from \$8 billion in 2016.

As for major issues, will AI replace jobs? The response is vague, as the answer is "maybe and some." More likely, AI-based computers will aid and complement workers to improve productivity, decision-making, and efficiency. A common response from the industry is that some jobs will be lost, but new different jobs will be created. That's what the industry said about robots, and it's essentially true.

Another issue raised by some is the possibility of AI being a threat to humanity. AI is good, but not that good. Its primary

uses will be data analysis, problem solving, and decision making based on available information and distilled knowledge. Otherwise, humans still dominate, especially when it comes to innovation and creativity. However, it's difficult to predict what the future may hold. Scary super-intelligent robots are not here... yet.