Secure Coding.
Practical steps to defend your web apps.

Overview and Tutorial on Artificial Intelligence Systems

Jim Hurst

Overview and Tutorial on Artificial Intelligence Systems

Introduction

Humans have a strong fascination with the idea of machines that can think. It is a recurring theme in fiction and the movies, from *Hal* to the *Terminator* to *The Matrix*. Whether a man-made artifact can think is still the subject of debate by philosophers. This paper explores different aspects of term artificial intelligence (or AI) and provides an overview of how this technology is being used today.
Defining artificial intelligence is somewhat controversial, because there is no single definition of intelligence. Stanford scholar John McCarthy states that artificial intelligence is “the science and engineering of making intelligent machines, especially intelligent computer programs.”
The artificial intelligence community can be roughly divided into two schools of thought: conventional AI and computational intelligence. Conventional AI is based on machine learning, which is the development of the techniques and algorithms that allow machines to “learn” or at least simulate learning. Machine learning attempts to use computer programs to generate patterns or rules from large data sets. This problem is similar to the data-mining problem (and data mining is one area where AI has found commercial success). Machine learning makes heavy use of symbolic formalism and logic, as well as statistics. Key areas in conventional AI include case-based reasoning, behavior-based AI, Bayesian networks, and expert systems.
Computational intelligence, in contrast, relies more on clever algorithms (heuristics) and computation and less on formal logical systems. Computational intelligence is sometimes referred to as *soft computing*. It often involves iterative methods using computation to generate intelligent agents. Whereas conventional AI is considered to be a top-down approach, with the structure of solutions imposed from above, computational intelligence is more bottom-up, where solutions emerge from an unstructured initial state. Two areas of computational intelligence will be discussed further: neural networks and fuzzy logic.
Also, hybrid intelligent systems attempt to combine the two approaches. Some proponents claim that this is appropriate, because the human mind uses multiple techniques to develop and verify results, and hybrid systems show some promise.

Weak AI Versus Strong AI

Another distinction within the artificial intelligence community is weak AI versus strong AI. Weak AI refers to using software to solve particular problems or reasoning tasks that do not encompass fully human intelligence. Strong AI implies creating artificial systems that are fully self-aware that can reason and independently solve problems. Current research is nowhere near creating strong AI, and a lively debate is ongoing as to whether this is even possible.

Neats Versus Scruffies

Another division in the artificial intelligence community is over the best way to design an intelligent system. The Neats maintain that the solution should be elegant, obvious, and based on formal logic. The Scruffies hold that intelligence is too messy and complicated to be solved under the limitations the Neats propose. Interestingly, some good results have come from hybrid approaches, such as putting ad hoc rules (Scruffy style) into a formal (Neat) system. Not surprisingly, the Neats are often associated with conventional artificial intelligence, whereas the Scruffies are usually associated with computational intelligence.

Expert Systems

Conventional AI has achieved success in several areas. Expert systems, or knowledge-based systems, attempt to capture the domain expertise of one or more humans and apply that knowledge. Most commonly, this is done by developing a set of rules that analyze information about a problem and recommend a course of action. Expert systems demonstrate behavior that appears to show reasoning.

Expert systems work best in organizations with high levels of know-how and expertise that are difficult to transfer among staff. The experts explain how they solve problems that are incorporated into the system. The simpler expert systems are all based on binary (true/false) logic, but more sophisticated systems can include methods such as fuzzy logic.

At the heart of an expert system is an inference engine, a program that attempts to create answers from the knowledge base of rules provided by the expert. Knowledge engineers convert a human expert's "rules-of-thumb" into inference rules, which are if-then statements that provide an action or a suggestion if a particular statement is true. The inference engine then uses these inference rules to reason out a solution. Forward chaining starts with the available information and tries to use the inference rules to generate more data until a solution is reached. Backward-chaining starts with a list of solutions and works backward to see if data exists that will allow it to conclude that any of the solutions are true.

Expert systems are used in many fields, including finance, medicine, and automated manufacturing. One expert system that many people may be familiar with is the Microsoft Windows troubleshooting software, accessed by through the help section of the Windows taskbar. This system provides diagnostic advice and suggestions for common user problems.

Case-Based Reasoning

Another approach from conventional AI that has achieved some commercial success is case-based reasoning, or CBR, which attempts to solve new problems based on past solutions of similar problems. Proponents argue that case-based reasoning is a critical element in human problem solving. As formalized in computer reasoning, CBR is composed of four steps: retrieve, reuse, revise, retain. First, access the available information about the problem

(Retrieve). Second, try to extend a previous solution to the current problem (Reuse). Next, test the refactored solution and revise it if necessary (Revise). Finally, store the new experience into the knowledge base (Retain).

Behavior-Based AI

Behavior-based AI (BBAI) is the final methodology of conventional AI considered. Behavior-based artificial intelligence attempts to decompose intelligence into a set of distinct, semi-autonomous modules. BBAI is popular in the robotics field and is the basis for many Robocup robotic soccer teams, as well as the Sony Aibo.

A BBAI system is composed of numerous simple behavior modules, which are organized into layers. Each layer represents a particular goal of the system, and the layers are organized hierarchically. A low layer might have a goal of “avoid falling,” whereas the layer above it might be “move forward.” The move forward layer might be one component of a larger “walk to the store” goal. The layers can access sensor data and send commands to the robot’s motors. The lower layers tend to function as reflexes, whereas the higher layers control more complex goal-directed behavior.

Bayesian Networks

Bayesian networks are another tool in the conventional AI approach. They are heavily based upon probability theory. The problem domain is represented as a network.

This network is a directed acyclic graphic where the nodes represent variables, and the arcs represent conditional dependencies between the variables. Graphs are easy to work with, so Bayesian networks can be used to produce models that are simple for humans to understand, as well as effective algorithms for inference and learning. Bayesian networks have been successfully applied to numerous areas, including medicine, decision support systems, and text analysis, including optical character recognition.

Neural Networks

There is no widespread agreement yet on exactly what Computational intelligence (CI) is, but it is agreed that it includes neural networks and fuzzy computing. A neural network consists of many nodes that cooperate to produce an output. The system is trained by supplying input on the solution of known problems, which changes the weighting between the nodes. After training has tuned the parameters between the connections, neural networks can solve difficult problems in machine vision and other areas.

Also known as neurocomputing, or parallel distributed processing, neural networks loosely model structures in the human brain. Neural network outputs rely on the cooperation of individual nodes. Data processing in neural networks is typically done in parallel, rather than sequentially as is the standard for nearly all

modern computers. Neural nets can generalize from their training, and solve new problems, so they are self-adaptive systems. Neural networks have been criticized as “bad science” because it is difficult to explain exactly how they work. Nonetheless, neural networks have been successfully applied in areas as diverse as credit card fraud detection, machine vision, chess, and vehicle control.

Fuzzy Logic

Fuzzy logic, fuzzy systems, and fuzzy set theory are all ways to refer to reasoning that is based upon approximate values, rather than precise quantities. Modern computers are built upon binary, or Boolean, logic that is based on ones and zeros. The bit is zero or one, yes or no, with no middle ground. Fuzzy systems provide for a broader range of possible values.

Consider the question, “Are the books in the study?” Well, yes, there are books in the study. There are also books in the office, books in the bedroom, and a pile of books in the doorway to the study. Fuzzy logic provides for an answer of 72%, meaning that 72% of the books are in the study. Fuzzy sets are based on vague definitions of sets. They are not random. Fuzzy logic is not imprecise; rather, it is a formal mathematical technique for handling imprecise data.

Like neural networks, fuzzy logic is subject to controversy and criticism. But systems based on fuzzy logic have an excellent track record at certain types of problems. Antilock braking systems are based on fuzzy logic, and many appliances incorporate fuzzy logic. The artificial intelligence systems used in nonplayer characters of modern video games often used fuzzy logic.

Summary

This paper has provided an overview of artificial intelligence and its applications. Artificial intelligence is not a single thing; it is many things. A distinction is often made between conventional artificial intelligence, which is based upon machine learning, and computational intelligence, which applies iteration and computation to generate intelligent agents. Conventional AI relies on symbolic logic and formalism and is considered a top-down approach. Computational intelligence, sometimes known as soft computing, uses training based on empirical data and is more of a bottom-up approach.

Four areas of conventional AI were examined. Expert systems use rule-based inference engines to simulate reasoning. Case-based reasoning tries to apply previous experience to the current problem. Behavior-based AI uses sets of semi-autonomous agents to solve problems. Bayesian networks model knowledge by using probabilities and graph theory to draw inferences.

Two types of computational intelligence systems were examined. Neural networks consist of networks of cooperating nodes that are trained by a series of inputs. Fuzzy logic is a system of logic that allows for imprecise information and solves problems with incomplete information.

All these techniques have successful applications in use today. Artificial intelligence is improving our lives in numerous ways, even when there is no widespread agreement on exactly what artificial intelligence is.

References

Basic Questions (about Artificial Intelligence). John McCarthy. http://www-formal.stanford.edu/jmc/whatisai/node1.html.
Welcome to AI Topics. American Association for Artificial Intelligence. http://www.aaai.org/AITopics/html/welcome.html.
What Is Computational Intelligence? Computational Intelligence Group, Department of Computer Science, Vrije University. http://www.cs.vu.nl/ci/.

Last Updated: June 24th, 2017

Upcoming SANS App Sec Training

SANS London July 2017	London, United Kingdom	Jul 03, 2017 - Jul 08, 2017	Live Event
SANSFIRE 2017	Washington, DC	Jul 22, 2017 - Jul 29, 2017	Live Event
Community SANS Minneapolis DEV534	Minneapolis, MN	Aug 25, 2017 - Aug 28, 2017	Community SANS
Community SANS San Francisco DEV541	San Francisco, CA	Aug 28, 2017 - Aug 31, 2017	Community SANS
SANS Network Security 2017	Las Vegas, NV	Sep 10, 2017 - Sep 17, 2017	Live Event
Secure DevOps Summit & Training	Denver, CO	Oct 10, 2017 - Oct 17, 2017	Live Event
SANS Seattle 2017	Seattle, WA	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Cyber Defense Initiative 2017	Washington, DC	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS Cyber Defense Initiative 2017 - DEV522: Defending Web Applications Security Essentials	Washington, DC	Dec 14, 2017 - Dec 19, 2017	vLive
SANS Security East 2018	New Orleans, LA	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS OnDemand	Online	Anytime	Self Paced
SANS SelfStudy	Books & MP3s Only	Anytime	Self Paced