Bits & Bytes

Arkansas' Premier Computer Club

April 2019

Bella Vista Computer Club - John Ruehle Center

Highlands Crossing Center, 1801 Forest Hills Blvd Suite 208 (lower level), Bella Vista, AR 72715

Website: http://www.BVCompClub.org Email: editor@bvcompclub.org

HOW TO FIND US

All meetings are on the lower level of the Highlands Crossing Center in Bella Vista. You may use entrance A on the West side or entrance C on the South side and take the elevator or stairs to the lower level. Turn left (West) to reach the General Meeting room, right for the John Ruehle Training Center. Additional information is on our web site.

MEETINGS

Board Meeting: April 8, 6:00 pm, John Ruehle Training Center

General Meeting: April 8, (2nd Monday), 7:00 pm, Community Room 1001.
Program: "Panel of Experts Q&A", with Justin Sell and Woody Ogden. Bring your computer questions or submit them to **Q.and.A@bvcompclub.org** in advance

Bring a guest! New Members and Guests are always welcome at the General Meeting

Genealogy SIG: April 20 (meets 3rd Saturday of the month).

MEMBERSHIP

Single membership is $20; $10 for each additional family member. Join by mailing an application (from the web site) with check, or complete an application

HELP CLINICS

Saturday, April 6, 9am – noon
Wednesday, April 17, 9am – noon
Saturday, May 4, 9am – noon

Help clinics are a free service for BVCC club members, held in the Training Center

Bring your tower, laptop, tablet or smartphone for problem solving.

CLASSES

"Computer Security for Regular People, Part 1" – Justin Sell, Tuesday, April 16, 6:30 – 8 pm
Part 2 will be offered again on 3rd Tuesday in May.

"Basic Computer Knowledge"-- Joel Ewing, Tuesday, April 23, 10am - noon

Advance sign up required for classes: Contact Grace: email to edu@bvcompclub.org, text 469-733-8395, call 479-270-1643, or sign up at the General Meeting. Classes are **free to Computer Club members** and are at our John Ruehle Training Center.

Check the monthly calendar and announcements for any last minute schedule changes at http://bvcompclub.org .

WHAT IS ARTIFICIAL INTELLIGENCE?

By Joel Ewing, President, Bella Vista Computer Club
President (at) bvcompclub.org

Introduction

You may have noticed that the term "artificial Intelligence" ("AI") has of recent become popular with marketing people, appearing frequently these days in ads. The manner in which ads use the term is much more casual than in the scientific community. Advertisements seem to include as "AI" things I would consider merely large-scale automation. "AI" used to be reserved for actions by devices that appeared to in some way mimic human understanding and thought processes: the ability to interact with humans through written or spoken language, to learn from their environment, or give the appearance of some level of "understanding".

Talos

There has been a long fascination with the idea of mechanical men and artificial beings. Greek myths include robot automatons, and science fictions stories of the last several centuries include many instances of sentient humanoid robots or other mechanical devices that exhibit human-like intelligence or even intelligence superior to humans.

The historic basis of artificial intelligence was the idea that the process of human thought could be mechanized – an idea considered by Chinese, Indian, and Greek philosophers as far back as the first millennium BCE. Formal methods of reason and logic were developed over the centuries, and the 20th century study of mathematical logic by George Boole suggested that there might finally be a path to artificial intelligence – that perhaps all human reasoning could be in some fashion be "computed". This belief was encouraged by the discovery that the brain was basically a very large and complex electrical network of neurons, each one of which fired all-or-nothing pulses – the output of a single neuron is in that sense digital.

Two most significant results from the 20th century on the theoretical capabilities of computers were (1)the Church-Turing thesis – that any function computable by humans on natural numbers or symbols by following an algorithm can be computed, ignoring resource limitations (think memory and time), by an extremely simple]machine design[1], and (2)that all the more-complex computational machine designs thought up by others could be emulated by that same simple machine. The importance of this is that if artificial intelligence can be computed, you should be able to demonstrate some aspects of AI on any machine with sufficient speed and resources. Another interesting established result was that even in the realm of mathematics there are statements that may be true but which not only can't be proven but can be proven to not be provable.

Some of the philosophical critiques of AI argue that rather than algorithmic reasoning, human reasoning involves a great deal of embodied, instinctive, and unconscious knowledge. We cannot know whether these are unique attributes of our reasoning process (how can we know that which we don't think about), or whether they are just the result of very complex algorithms we unconsciously apply to the vast body of knowledge we acquire over a life span. What we regard as instincts could just be sub-conscious reasoning based on inherited unconscious knowledge wired into our genes.

---

1 A "Turing" machine, named after its creator Alan Turing. This machine is a mathematical concept designed for ease of proving a program to compute a solution exists. A physical machine approximating this design could be built, but it would be an extremely impractical and inefficient choice for doing actual computations in the real world.

Early Progress in AI

Many of the early attempts at AI dealt with the playing of games, because games involve a manageable number of well-known fixed rules. Computers could be programmed by various means to make moves having a greater likelihood of a winning outcome.

The first chess-playing machine was a mechanical automaton called *The Turk*, built around 1769. It wasn't until almost a century later, in the 1850's, that it was fully revealed as a fraud, with a small human operator artfully concealed inside. A complete game of chess proved too complex to solve with mechanical devices, and would have to wait until the 1940's and 1950's before serious work would be done on methods to play a full game of chess using digital computers. It would take until 1997 before the speed of digital computers progressed to the point that IBM's Deep Blue was able to win a chess game and match against a chess master (Garry Kasparov).

Checkers is a simpler game in terms of the rules, but with more depth of moves. In 1959 Arthur Lee Samuel created a checkers playing program for the IBM 701 that learned how to play a better game by remembering which board positions had greater probability of a favorable outcome. By 2007, computer power had progressed to the point that Jonathan Schaeffer succeeded in mathematically "solving" the game of checkers, proving by exhaustive search that a perfectly played game would always end in a draw and producing the means for playing a perfect game. In the process of developing checkers playing programs, Schaeffer benefited from his program competing with and narrowly losing to the undisputed checkers genius Marion Tinsley. Tinsley was able to perceive many moves further ahead than was possible for Schaeffer's program, which in one match against Tinsley made what Tinsley immediately recognized was a "bad" move in the 10th game after 9 draw games. After Tinsley eventually won the match, it was evident that he somehow understood he could defeat the computer 64 moves later.

General Approaches to AI

Two basic approaches that have been used to create AI machines: a purely functional approach, and emulated reasoning.

The functional approach is the idea that you have many values that describe the input to a decision (like values of pixels in an image, wavelengths present in a sound, answers to a large set of questions, etc.) and many input examples with a known correct output decision response. The goal is to determine some process or algorithm that when applied to the input values will determine the correct decision result in as many cases as possible in the hope/belief that the same process applied to a new set of input values will also produce what a human would consider a correct result with high probability. A trained "neural network" that mechanically mimics some of the functions of a neurons in a human brain on a very small scale is one example of this approach. While this approach can work for some types of problems it has the disadvantage that there is no way to logically explain the reasoning behind a good or bad decision made by the machine.

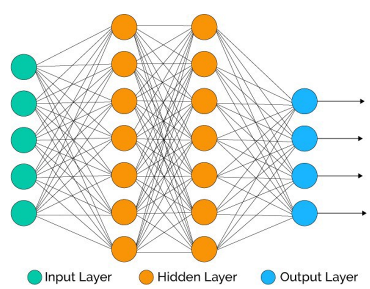

Neural Network

The reasoning emulation approach attempts to build into a machine a database of knowledge and rules similar to the way humans believe they think. In some form a machine is programmed to represent things (the "nouns" of natural language), attributes associated with those things, actions that are associated with those things, relationships between things, complex rules for inferring new knowledge and relationships, and rules for determining when an action should be taken based on known attributes of things. Even when this can be confined to a very narrow branch of knowledge the size of such a database can become very large and complex. Think about the complexity of telling another human who has never seen a dog or cat enough attributes describing a four-footed animal sufficient for them to determine reliably whether the animal is a dog, a cat, or neither, Add to that the complexity of how humans when faced with incomplete facts can still manage to reach partial conclusions ("I think it is probably a cat, but I can't be sure")

Like human perception where we just instantly "recognize" some facts without understanding the thought process involved ("This object is a tree") and reason out other facts ("the height of the tree is about 60' by comparison with the height of the house"), both of these approaches may be combined within the same AI application.

Recent Progress in AI

The field of AI research in the last half of the 20^{th} century went through repeated periods of great optimism followed by periods of failure. Things that were first assumed would be easy, proved to be difficult to achieve. What has evolved is a division of AI into many different sub-disciplines working on many separate problems. Some of these now have solutions that are good enough that many forget to credit AI research. While there have been useful partial solutions in many of these areas, we are still far from a sentient robot.

Many are now familiar with the concept self-driving vehicles and voice-activated devices, but the more difficult focus of Artificial intelligence is on how machines can "learn". Much current research is being done in finding practical applications of "Deep Learning": the idea that a machines could be designed to learn by example or store new facts and develop solutions or new responses with limited human supervision by being fed new facts or data with a large number of examples.

Systems that attempt to do meaningful "deep learning" on complex problems, especially if a large base of common-sense knowledge is involved, still require large amounts of expensive hardware, and considerable time and effort to "teach" the system.

Some of the areas in which AI progress has been made involve

- Recognizing human speech
- Forming natural language responses
- Learning from natural language
- Understanding natural language – structure is easier than meaning
- Recognizing words and sentences from images of printed or hand-written text
- Recognizing objects in still and live images and reaching conclusions about objects
- Mimicking animal mobility
- Learning from the environment (mimicking human senses)

- Correlating information from many different sources

There has been much improvement in the last several decades – but it is still far from equivalent to human intelligence. We can talk to a smart phone to compose text messages, not with 100% accuracy but usually more accurately than I can type with a tiny keyboard. Still, at times, one gets some very strange misinterpretations, so to be safe you need to proofread before sending. We can verbally ask or type questions to Cortana on Windows 10 and get responses to a query – but the response is not always what you want, and you certainly wouldn't think you are conversing with another human or expect good answers to questions on a random topic. Optical character recognition is much better than ten years ago, but a human can still read imperfectly-formed text in an image that OCR software can't come close to reading, and bad hand-written text is even worse. Image recognition has evolved to the point that human faces can be isolated and compared with some degree of accuracy; and that lanes in a road can be detected and followed, and most obstructing objects can be detected and avoided.

A machine that is able to correlate and organize related information is getting close to mimicking the concept of "understanding"; but, there is an incredible amount of common-sense knowledge that we unconsciously use every day. For example, if you were to attempt to program a self-driving vehicle to minimize the number of human deaths in an emergency situation where some death can't be avoided, how do you add in the rules that unrecognizable moving objects along the road on the evening of October 31 may be human children in Halloween costumes?

One of the more useful results from AI is the development of "Expert Systems" where the object is to solve problems only about a specific domain of knowledge using logical rules derived from the knowledge of experts. These systems are manageable because confinement to a narrow domain eliminates having to deal with the full breadth of common-sense knowledge. There are successful examples of Expert AI Systems in such widely-different areas as stock trading, medical diagnosis assistance, corporate sales prediction, and planetary-exploration-vehicle navigation systems.

AI and Morality

If a human through action or inaction causes serious and un-correctable harm or injury to others, he can be held accountable. Who is accountable if the harm is caused by a decision made by machine using AI technology? This is especially troubling if it is obvious a competent human acting under the same circumstances could have made a better decision. Recent events of planes seemingly choosing to override human pilots and dive into the ground catastrophically come to mind; although, to be fair it must also be remembered that on rare occasions human pilots have done similarly self-destructive acts.

Science fiction (novels, movies, TV series) are full of examples of AI machines either running amok, or working as designed with undesirable consequences, and causing great threat to living beings. All machines eventually wear out or can just fail, so one of the lessons we should learn from literature is that lethal force, especially the power of mass destruction, should never be placed in the hands of autonomous AI machines. In fiction, even when the goal of the machine was purely defensive or even peaceful in nature, the lethal force eventually became a threat to the creators ("Forbidden Planet", "Colossus"/"The Forbin Project", "War Games", etc.). Humans can make bad and lethal decisions also, but their slower reaction time improves the odds that really bad decisions that effect thousands of lives can be circumvented by other humans.

The harm done by AI applications can also be more subtle. Many humans tend to assume decisions suggested by a machine using some form of AI as being more accurate and unbiased. It is important to always keep some level of skepticism, because AI applications and machines are designed by fallible humans who don't always recognize their own biases. It has been found that some large banks have been using AI applications to help loan officers decide whether loans should be approved, but the applications unintentionally discriminated against minorities because the questions and data designed into the decision making process did not take into consideration the different circumstances and culture of minority communities.

Automated AI stock trading applications are known to make the markets less stable. A particular AI trading application, when used by many traders, may correctly or falsely cause many traders to buy or sell a stock at the same time, amplifying a small market movement into a larger one, de-stabilizing the market in ways that hurt investors. Also, traders with higher-speed access to on-line trading have been known to use automated AI trading systems to micro-manipulate a stock price by picking opportune times to issue back-to-back buy-sell orders to nudge the price slightly higher and pocket the difference to the detriment of other investors with slower access. These practices may not necessarily be illegal, but the ethics are certainly questionable.

Cylon

It is challenging to speculate whether technology will ever advance to the point that it is possible to develop a machine so complex that it is for all practical purposes a sentient being. It is a serious scientific debate as to whether such an outcome would benefit or doom humankind. Such a machine could solve problems that have eluded humans, but it could also decide that humans are the problem. SciFi literature is also divided as to whether the end result of sentient machines will be negative ("Battlestar Galactica") or positive (works of Isaac Asimov) for humanity.

R. Daneel Olivaw

The only reason Asimov's robots ultimately succeed in preserving humanity is that they were somehow infallibly constrained by robotic laws that required them to act in ways that preserved human life with which they come in contact. Asimov's robots eventually realized on their own that the only way they could properly follow their mandate to preserve individual human lives long-term was to mandate as a higher priority that they act in ways that preserve humanity as a whole. Asimov is quoted as saying a rational person would not build a robot without implementing the robotic laws, but then he followed that with "humans are not always rational". I personally hope the possibility of sentient robots is far in the future, because I doubt we currently have the wisdom to follow Asimov's guidance.