

Bits & Bytes

Arkansas' Premier Computer Club



February 2024

The Bella Vista Computer Club - John Ruehle Center

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

Website: <http://BVComputerClub.org>

Email: editor@bvcomputerclub.org

MEETINGS

Board Meeting: February 12, 6pm, in John Ruehle Training Center, Highlands Crossing Center.

General Meeting: February 12, 7pm. Program: "Artificial Intelligence", presented by Joel Ewing. AI continues to be a hot topic in the news. This will look at some of the free AI-based query tools currently available, like ChatGPT, and at the power and limitations of AI. Current AI has the ability to provide benefits; but can also cause problems if its limitations are not understood, or if it is deliberately misused.

We will meet in-person in **John Ruehle Training Center**, Highlands Crossing Center, lower level, 1801 Forest Hills Blvd, Bella Vista, or you may attend the meeting on-line via Zoom. Zoom access information is published on our website.

Visitors or Guests are welcome.

Consider attending by Zoom if you are unable to attend in-person.

HELP CLINICS

February 3, 6, 9am - noon at John Ruehle center

February 21, 9am - noon at John Ruehle center

Members may request Remote Help on our website at <https://bvcomputerclub.org> at menu path **Member Benefits ▶ Remote Help** .

Genealogy SIG: February 16, 1pm-3pm, Training Center

MEMBERSHIP

Single membership is \$30; \$15 for each additional family member in the same household.

Join on our website at <https://bvcomputerclub.org> at menu path **Get Involved ▶ Join/Renew**, by mailing an application (from the web site) with check, or complete an application and pay in person at any meeting.

CLASSES

(At BVCC Training Center)

Tuesday, February 20, 9am-11am, "Microsoft Word", with Joel Ewing.

Advance sign up required for each listed class: For reservations: email to edu@bvcomputerclub.org, or sign up at the General Meeting. Classes are **free to Computer Club members**.

Check the monthly calendar and announcements for any last minute schedule changes at <https://bvcomputerclub.org> .

BVCC MARCH 11, 2024 COMPUTER RAFFLE

Once again BVCC will be raffling of a refurbished and customized computer. This time it's a business-type, high-end workstation computer, a Lenovo P520 30BF. For pictures and detailed specs, see our website at <https://bvcomputerclub.org/raffle2024.php> . The winner will be drawn at the March 11 General Meeting. It is not necessary to be a member of BVCC or to be present at the General Meeting to win. Home setup support will be provided within Benton County, AR.

Tickets may be obtained for a \$10 donation each. For additional information on the computer or for tickets, call or text Woody Ogden at 479-966.9357, Pete Opland at 218-753-2353, or Russ Ogden at 360-789-0139.

Tickets will be available at the January, February, and March General Meetings and at other BVCC events up to the March General Meeting. Members are encouraged to participate by obtaining their own tickets and/or by helping to distribute tickets to their contacts.

A mouse and keyboard are included, but not a monitor. For those who do not have an existing monitor, BVCC may also have available an inexpensive used monitor. The computer comes with Microsoft Windows 11 Pro and with Microsoft Office 2021 Pro Plus software.

This is a highly-capable computer that would be excellent for any home use application and even for many gaming applications



ANOTHER LOOK AT ARTIFICIAL INTELLIGENCE

By Joel Ewing, President Bella Vista Computer Club
Bits & Bytes, February 2024
<https://bvcomputerclub.org>

(For a general description of artificial intelligence, also see the article "What Is Artificial Intelligence" in the April 2019 edition of the BVCC *Bits & Bytes* at <https://bvcomputerclub.org/PDF/B&B/2019-04.pdf>)

Background

Although having a background in Computer Science, I do not consider myself an expert on AI; but like many, I have seen over the last decade the improved (but not perfect) ability of smart phones to translate the spoken word into written text, chatbots on the Internet and telephone call centers that can briefly make you think you have reached a human, and the fairly sophisticated behavior of "self-driving" cars, including some experiments with autonomous self-driving vehicles in restricted environments.

When in full "self-driving" mode a Tesla car can recognize enough of the environment around it to sense and follow lane markings, to sense other vehicle types and their relative speeds on all sides, interpret speed limit signs, stop signs, stop lights, a variety of warning signs, speed limit signs, to adjust speed for sharp curves and slower vehicles, to make decisions on passing slower vehicles. When navigation mode to a destination is engaged, an extensive road map database allows it to know when to turn and what lane it needs to be in for the next turn, to suggest the optimal super-charger stations for long trips, and to automatically pre-condition the battery temperature for optimal charging speed before arriving at a charging station to minimize charging time. Yet with all that sophistication the Tesla's most important safety features are (1) monitoring the human driver for alertness, and (2) alerting the human driver to take over if the road conditions go outside the bounds that the automation can reliably handle. This is great for removing much of the tedium of long-distance driving on consistently marked interstate highways; not so useful in city driving where road repairs, obstructions, and local driving customs tend to demand more human interaction.

The topic of AI was recently raised in a professional computer-related on-line discussion list that I follow. I have also been involved in a church-related Religion & Science group meeting on Zoom over several years that has discussed, among other topics, whether humans are the only life forms on Earth that are sentient or aware, AI, and whether machines could ever cross the thresholds of self-awareness and become sentient and direct competitors of humans, possibly even a threat. Some of the following thoughts come from those discussions.

AI Current Capabilities

AI may easily surpass humans in the ability to analyze and deduce conclusions from a large amount of data and propose actions – it already has done that in a number of areas that are sufficiently limited in scope.


AI is already being utilized to assist in medical diagnosis, such as interpreting medical scans, where large scans may need to be analyzed tediously to look for small anomalies.

Many Internet search engines utilize some AI techniques to improve the quality of the results.

Some of the more interesting recent free Internet search additions are ChatGPT, BardAI, and BingAI, which can accept queries in natural language and return paragraph responses that appear appropriate, almost human-like. It is possible to request responses that are in the form of specific types of poetry (ChatGPT), or to request creation of an image featuring certain objects (BingAI).

Without getting too much into details, there are at least two main techniques that are involved under the covers for the deep learning of advanced AI: trained digital neural networks, and large databases created by analyzing many natural language sources of information.

Digital neural networks attempt to emulate the design of the brain on a small scale. The idea is that you have a large number of nodes (neurons) that are interconnected to nearby layers of the network, but that each of the interconnections can be adjusted by giving a "weight" to the each of the connection paths to determine what strength of incoming signals is sufficient to activate the neuron. To "teach" a neural network, you must supply many different combinations of input signals (on the left), for which the desired output (on the right) is known, and modify the weights so that the correct outputs are obtained for as many as possible, hopefully all, input value combinations. This is a computationally expensive and thus energy intensive process, but once the best weights are obtained, it is easy to replicate the neural network. The illustration on the right only has 23 nodes. A useful neural network may have 1000's of nodes.



Neural networks are particularly useful for cases where no clear written rules exist for arriving at the outputs, like pattern recognition of images and sounds, interpreting the ambiguous parts of natural language speech and text, the prediction of large social system behavior like the stock market. While it can be hoped that a new combination of inputs that was not used in the training process will also produce reasonable outputs, this is not guaranteed and heavily dependent on how the training data was chosen. Another drawback of neural networks is that there is no way to explain how it arrived at any specific output from some specific input.

It is clear from the behavior of ChatGPT that it has access to a very large database of contemporary and historical literature. Whether that does or should include those materials that are still under copyright protections is still under debate as to whether that violates fair-use restrictions. To be of greatest use the natural language must be parsed to determine what is being discussed (nouns) and related descriptions (adjectives) and actions (verbs) and deduce further relationships. From that it is possible to produce indexes and statistical probabilities of relationships. Applying a similar parsing to information requests, it is possible to either find sentences or paragraphs that are relevant to the query, or (more difficult) possibly construct "original" relevant sentences based on known associations with the topic of the query.

AI Limitations

One problem that needs to be addressed is that there is a considerable amount of data loose in the real world that is misinformation, and too much garbage-in still produces garbage-out.

Many humans have problems telling the difference between fact and fiction, and between reality and conspiracy theories. Humans have biases on their choices of data; yet it is humans that will be choosing the sources of data to use to program future AI. This will inevitably cause some problems.

Every year you find people, sometimes prominent people, that stumble across a parody or satire written in the style of a news article, mistake it for real news, and make a fool of themselves by calling others to take action against the alarming "news" they have just uncovered. Such humor is usually so over-the-top as to be easily

recognized as humor by normal people aware of current events, but still there are always some that take it as factual, even if it is located on a website like "The Onion" that is well-known for only publishing imaginary news releases.

One needs to understand that while AI can provide useful guidance within the realm for which it has been trained, that it is not infallible. If you ask a neural network to categorize an unexpected combination of inputs it may make an incorrect identification. If you ask an AI system trained from analyzed text about a topic for which it lacks accurate information, it may make incorrect inferences from data it thinks is similar because it appears statistically related. Even if it has processed relevant information, it may not be able to accurately interpret context, implicit relationships, and time ordering of relationships. If you just ask an ambiguous question ("show me something like X"), the AI may not understand the distinction between creating something fictional that looks like "X" versus finding an actual fact like "X".

There have already been at least two instances reported in the news of a judge berating attorneys for submitting supporting briefs that referenced legal precedents that didn't exist, because some AI system created a plausible fictional reference to support their legal argument.

The general public is already displaying an indifference to the limitations of AI and abusing AI: asking questions of AI apps in areas for which the AI is poorly trained and trusting without verification fallible responses as being authoritative because they are produced by a machine. Some newsworthy cases of this abuse have been reported when fatalities result after owners of cars with AI autopilots requiring supervision find ways to entrust the AI with unsupervised autonomous driving, and a crash occurs when the AI encounters situations it is not designed to handle.

Could AI Be Dangerous

There are some who theorize that if we could build a neural network of enough complexity – many orders of magnitude larger than possible today – that at some point a threshold might be passed and that creation might become self-aware and possibly worthy of being called sentient. What discoveries such a tool could make and create! The problems are multi-fold: If it has more knowledge than any human, can exercise creative thought faster than any human, its analysis of human history may lead it to the conclusion it should be the master, not the slave, of humans. This has long been a persistent cautionary theme of science fiction writers. In the fictional Dune prequel "The Battle of Corrin", where sentient machines and the humans have long been at war, there is a fictional quote from the sentient robot Erasmus: "Humans were foolish to build their own competitor – but they couldn't help themselves."

If we were to succeed in creating machines that could truly approach the complexity and creativity of the human mind, might they not also be subject to some of the same dangerous mental instabilities as human minds, when confronted by conflicting "facts" or subtle design deficiencies? Fatal problems like those afflicted the fictional HAL 9000 in "2001: A Space Odyssey" and the M-5 Multitronic computer in "The Ultimate Computer" Star Trek episode.

I don't think we are at any risk within my lifetime of finding a way to cross that threshold, but that doesn't mean that we shouldn't be concerned that AI in its present state has the capability of causing harm.

AI-based tools used by financial institutions to evaluate whether to approve or deny credit requests have been found to exhibit bias against granting loans to those in minority groups. The AI tools were programmed using data about loans which had been approved subject to historical biases against minorities. Without great care in programming, AI can be used to justify that there is no bias, when the reality is that the bias was included in the AI programming and is just more hidden.

Another way we are already seeing AI being abused is by its use to create and spread misinformation on the Internet. While not perfect, it is already good enough to produce deep-fake video and audio and false news reports and control bots on social media networks that spread messages to promote and amplify division and hate. There is some hope that it may be possible to employ AI on the other side to detect and suppress such abuse, but that does require corporations to accept that responsibility.

Another area that should be a serious concern is the possibility that AI could be given the power to make and implement life and death decisions without adequate human oversight, or just to propose actions that humans might blindly follow in a life or death situation. The science fiction genre is full of plots that suggest how badly this can go – possibly human enslavement or human extinction.

When you grant AI the power to make decisions and take actions in the real world, most of us want not just "intelligence" involved, but also wisdom and morality. Most of us want decisions made by others that impact our lives to be constrained by a moral compass of what is right and fair. We have no idea how to build that awareness into AI; and since there is no universal agreement on whose standard of morality to use, that may be an impossible task. Even if there were agreement on how to define morality and an understanding of how to embed that concept into AI, do we trust fallible humans to consistently implement that design without error?

A legitimate concern is that if future AI regulation or its enforcement is inadequate, some individual or corporation with more power and money than wisdom may go rogue, ignore whatever AI restrictions are in place or commonly accepted, and create a dangerous AI device because of a belief it will enhance their short-term profit.

It would seem we need to be very cautious for the short and long term about what kinds of decisions and actions we entrust to AI. We also need to research how best to enforce that caution. Perhaps future AI, appropriately applied, might even become part of the solution.