



"It's a funny thing about artificial intelligence," mused my friend Jerry. He's an engineer who has worked for many years using neural network computing to develop systems that allow computers to speak in natural-sounding tones. "Whenever it actually comes up with something useful, they stop calling it artificial intelligence."

Language recognition and speech generation are examples of things that computers have an innately difficult time with, but which the various forms of artificial intelligence — or rather, neural networks and machine learning, as the engineers in the field prefer to call it — have managed finally to make some progress. As anyone can tell you who has tried to talk to the avatars of our smart phones, Alexa or Siri, it's gotten pretty good... but it's still far from perfect. Being understood by a machine is not the same as talking to a real person.

Back when a Jesuit novice, I worked at what is called a "sheltered workshop," a place where men with severe mental handicaps can be cared for as they earn money doing simple jobs consistent with their abilities. I was told that a typical IQ for these men would be 50 or below; recall, 100 is average on that scale. Typically, the men I was working with could not count up to three. But they all spoke English fluently. Counting is the sort of thing that even the very first computers could do well. Speech is something that still causes problems for computers. What do I conclude from this? Computers operate very differently than the human brain.

Granted, computer "neural networks" derive their name from how the structure of the computer programming was inspired—by the networks of neural cells that have been mapped in the human brain. But what these computer networks are doing is still very different from what's actually happening when a human being thinks.

Now, I am not an expert in the field of neural networks or other forms of what has been called artificial intelligence or AI. Rather, as an astronomer and physicist, I come to the topic as a user. With the advances in telescope technology, especially in electronic detectors and the ways that the signals from these detectors can be processed, astronomy is now being inundated with "big data"... so much data that our old ways of handling the results simply won't work anymore. (I am old enough to remember when we were happy to have any data at all!) Instead, we rely on clever computing algorithms — which I had no role in writing — to sift through what we have observed. The algorithm might suggest that one object could be an x-ray source, that another could be a site of planet formation, and so on.

Consider what a computer with what we call artificial intelligence is actually doing. Let's start with a simple, solvable problem. Let's say you have a body of a certain weight — a thick physics textbook,



for example — that's a certain distance above the ground. You want to determine, if you drop it, how long it will take to hit the floor, and where exactly it will land.

Well, Isaac Newton came up with an equation or two for that. We can assume that certain parameters won't change — for example: the mass of the textbook and the force of gravity acting on it. Four things that do change are the position of the textbook in three-dimensional space, and the time at which the book is to be found at any given position. Three space dimensions, one time dimension. That means four variables. So answering our question requires solving four different equations...one for each variable. In this case, Newton's laws of motion give us three of the necessary equations (one each for motion in each of the three directions) and his law of Gravity supplies the fourth. A high school student can solve these with paper and pencil. They are that simple.

But... what if the object, instead of a textbook, is instead a meteor from outer space hitting the top of the Earth's atmosphere? It is burning up, changing its weight as it falls. Friction with the atmosphere is not only burning it up, it is at the same time slowing down, as the energy of its fall is being turned into heat and light. The wind in the upper atmosphere can be pushing it in ever-changing directions. We can, of course, simply follow where the meteor traveled while it was burning up in the upper atmosphere and maybe, if it drops a meteorite, we can find the position where it hit. But in reality it's very hard to calculate that position. You still have to go out and mount a search to find the meteorite. And suppose instead we want to use the path and brightness of the meteor to determine the original speed, direction, and size of the rock before it hit the Earth?

Suddenly, the number of variables has grown faster than the number of reliable equations we can write down. It's not that the problem is insoluble; after all, nature solved it! But it's more complicated than what our high school student can calculate.

So the next step is to set up whatever equations that we do know... not just for motion, but for the friction of the air, the effect of the ever-changing winds, the rate at which different materials with different densities will burn and glow, and so forth. We can't solve backwards for a unique set of initial conditions—for example, the velocity of the meteor before entering the atmosphere. But, if we guess at different possible initial conditions, we might be able to calculate forward and see if our guesses produce a result close to what we observed. Of course, because there are so many more variables than we have equations, lots of different sets of initial conditions might work. So we start throwing dice, making random guesses for each of the starting conditions, and run the problem over and over again (that's where computers come in handy) to see which sets of variables are "winners."

Eventually, we might see a pattern in the starting conditions that gives answers close to what we observed. Of course, that's no proof that any of those conditions are actually what happened. But it hints at what might be going on.

That technique, metaphorically throwing dice to guess at possible answers, is called running a "Monte Carlo" simulation. One trouble with this approach is that most of the dice throws end up with results far from what's observed. But by looking more closely at how the results shift with each changing variable, you can begin to guess which variables have the biggest effect and, eventually, how you might load the dice to come to useable answers more quickly.

All of this can be done by the computer. It can determine if changing a certain variable makes things better or worse, as well as which set of possible answers is leading us in the right direction, bringing us closer to what we observe. This is an example of the sort of processing that we can begin to call "machine learning." Of course, it's important to remember that the answers are merely the most probable, not certain to be correct. And the machine is "learning" only in the sense that its algorithms are becoming more efficient.



Now suppose we run this kind of simulation for thousands of meteors, learning from our results which parameters, and which choices, tend to come up with our observed results most often. Because these decisions are calculated by the computer — without us, the users, knowing exactly what it chose to perform the calculations — the whole process can all look rather magical. But, in fact, this sort of “artificial intelligence” is merely a buzzword for clever computer programming. It’s data and data manipulation.

Rather than imitating human intelligence, computer AI systems wind up doing the work of manipulating data that would be mind-numbingly boring for a human to undertake. We’re too intelligent to waste our time doing such work!

As my old friend from astronomy graduate school (and computer maven) Cliff Stoll once pointed out: “Data is not information, information is not knowledge, knowledge is not understanding, understanding is not wisdom.” Computers are still on step one. They are only dealing with data; they are a long way from wisdom.

*“Data is not information, information is not knowledge, knowledge is not understanding, understanding is not wisdom.” Computers are still on step one. They are only dealing with data; they are a long way from wisdom.*

Consider the flashy AI bots like ChatGPT that answer questions by pulling information out of the internet. They manage to sound more “human” than search engines, which simply report relevant web sites. Instead, these AI bots tell us a very fuzzy version of what those web sites contain. In essence, they sound human by introducing human-like imprecision and error in their results. They are still just shuffling (and blurring) the data. They sound human not with artificial intelligence, but artificial stupidity.

Intelligence itself is a loaded term. We use IQ tests—centering the average at 100—to attempt a measurement of an “Intelligence Quotient.” But these tests assume that “intelligence” is a one-dimensional thing that can be measured objectively. A cynic might say that the only thing an IQ test can measure is how well you perform on IQ tests.

At the same time as I was working with the mentally handicapped, I also had an occasion to visit NASA’s Goddard Space Flight Center. It was an opportunity to keep up with my astronomy research while I was training to be a Jesuit. I remember sitting in the cafeteria, musing that usually at this time I would be having lunch at the sheltered workshop, surrounded by men who were fifty IQ points or more away from the average, whose clothes didn’t always fit right, who sometimes talked to themselves, who lived in a world most people couldn’t follow or understand. And then I looked at the genius astronomers around me having lunch. Uh, huh. I was reminded of a comment from a visitor to Cal Tech: “Did all these guys get dressed in the dark?” (Trust me, I didn’t dress any better when I was at MIT.) There’s more than one kind of intelligence.

My first experience with artificial intelligence actually goes back to my student days at MIT. During the summer of 1973, I was busy writing computer code to model how the interior of an icy moon of Jupiter might evolve with time. My roommate, Paul, was a math major working with Seymour Papert and Marvin Minsky, pioneers in the field, at the MIT AI labs on the east end of campus. We kept typical student hours. Late at night, I would hang out at Paul’s lab, waiting for my next computer model to run, getting to know the people who were inventing AI. Around midnight, we’d head out for dinner at an all-night deli.

Those same years, Hans-Lukas Teuber taught what was then called “Introduction to Psychology” at MIT. His lectures were held in the biggest lecture hall on campus, because half the people attending, like me, weren’t taking the course for credit; we were sitting in because he was such a spellbinding lecturer. His version of psychology was very different from what was taught at other universities. Indeed, that sort of work was later renamed, far more accurately, as “brain and cognitive science.”



Today, MIT has a Department of Brain and Cognitive Science; their web site claims that they “back-engineer” the brain. Certainly, they have done amazing work to map out how the brain works. But you can’t just recreate that map in a set of silicon chips. Cognitive Science is a different department, with a different set of tasks, than AI. And while the brain scientists and the AI programmers do talk to each other, it is not at all clear that they are really doing the same thing.

One thing I would like to emphasize. The folks I met at that time, some fifty years ago, were delightful people. But they were not godlike creators. They were just some very clever and very dedicated folk. What they study in cognitive science, what they create in the AI labs, is a map of how the brain works: how it does what it does. But the “how” is different from the “what.”

What is intelligence? That’s a question for philosophers and, perhaps, theologians.

In fact, I think one could do well to start with Saint Thomas Aquinas, who emphasized the importance of intellect and free will in the human being. In other words, you can think of a soul as an entity that is capable of being aware of both itself and things outside of itself; and then it is free to make choices about what it does with that knowledge, and how it interacts with those other entities.

Note that this “soul” has no mass, takes up no space. Unlike our textbook example above, the soul is not subject to Newton’s laws. And yet it is very real.

To appreciate its reality, let’s consider an analogy with a computer. Imagine a pair of identical laptops: same model, same operating system, even the same physical color. The difference (besides various scratches on the surface) is found entirely in the files that are stored within each.

Those files do not contribute to the mass of the computer; nor does a computer full of data take up more space than a computer whose hard drive is mostly empty. But those files are the very thing that makes one computer different from another.

Now consider... those files are stored in some electronic fashion, for instance in the orientation of magnetic grains, that ultimately represent an array of binary choices. Without an operating system to read and interpret that array, the physical manifestation of the data in the electronic medium would be indistinguishable from random noise. It’s the operating system that converts those bits into text.

But don’t go thinking this means that the operating system is intelligent. Because all the operating system actually does is create a series of colored dots on the laptop’s screen. And if you can’t read the letters and words those dots make — indeed if you don’t recognize the alphabet or whatever system of writing used by the language displayed on the screen — it will still just look like gibberish. What can we learn from this? That what the computer contains and displays is utterly meaningless without some entity outside the computer looking at the screen and giving it meaning. And that entity must have a consciousness, an intellect, and free will... a soul. Data is not information, is not knowledge, is not wisdom. Clever programming can allow computers to perform or at times out-perform humans in certain ways. But the programmers themselves, like my engineer friend Jerry, know better than to mistake what the computers are doing for intelligence.

*...what the computer contains and displays is utterly meaningless without some entity outside the computer looking at the screen and giving it meaning. And that entity must have a consciousness, an intellect, and free will... a soul.*

Still... will it ever be the case that human beings construct something that has intelligence? Something capable of wisdom? Actually, yes. Indeed, it happens every day. Babies, gram for gram,



The Joan & Bill  
**Hank Center**  
FOR THE CATHOLIC INTELLECTUAL HERITAGE

are more powerful than any computer. And, unlike computers, they can be produced anywhere by unskilled labor.

But though we can make new people — and new computers — for all of our human cleverness we’ve never really been able to invent a new sin. Evil is just the absence of a good that ought to be there. Greed and envy and gluttony and lust and all the rest will tempt any entity that has free will. And any entity with intellect and self-awareness will always be pondering the Big Questions... and be tempted by the lures of secret knowledge, or despair at ever knowing anything for certain. The poor will always be with us because we will always be tempted to make others poor, to lord it over others. Or we will be tempted to impoverish ourselves by being afraid to trust God.

But at the same time, the virtues of faith, hope, and love will also always be possible.

Even if machines do someday become self-aware, is this any reason to be afraid? If indeed they develop the traits of the human soul, intelligence and free will, that makes them no better and no worse than any other soul, human or other. They will be capable of sin; and of love. Why should we assume the worst?



### Br. Guy Consolmagno, S.J.

Brother Guy Consolmagno, S.J., is the Director of the Vatican Observatory. A native of Detroit, Michigan, he studied planetary sciences at MIT and the University of Arizona, studying meteorites and asteroids. Along with more than 200 scientific publications, he is the author of several popular books on astronomy and the relationship between faith and science. In 2014 he received the Carl Sagan Medal from the American Astronomical Society Division for Planetary Sciences for excellence in public communication in planetary sciences.

Next Article

THE JOAN AND BILL HANK CENTER FOR THE CATHOLIC INTELLECTUAL HERITAGE · 1032 West Sheridan Road · Cuneo Hall, Room 421, Chicago, IL 60660 · 773.508.3820

