Mindless Machines

Ben Tarr

The official foundations for "artificial intelligence" were set forth by A. M. Turing in his 1950 paper "Computing Machinery and Intelligence," wherein he also coined the term and made predictions about the field. He claimed that by 1960, a computer would be able to formulate and prove complex mathematical theorems, write music and poetry, become world chess champion, and pass his test of artificial intelligences. In his test, a computer is required to carry on a compelling conversation with humans, fooling them into believing they are speaking with another human. All of his predictions require a computer to think and reason in the same manner as a human. Despite 50 years of effort, only the chess championship has come true. By refocusing artificial intelligence research to a more humanlike, cognitive model, the field will create machines that are truly intelligent, capable of meet Turing's goals. Currently, the only "intelligent" programs and computers are not really intelligent at all, but rather they are clever applications of different algorithms lacking expandability and versatility. The human intellect has been used only in limited ways in the artificial intelligence field; however it is the ideal model upon which to base research. Concentrating research on a more cognitive model will allow the artificial intelligence (AI) field to create more intelligent entities and ultimately, once appropriate hardware exists, a true AI.

Webster's Collegiate Dictionary defines intelligence as the capacity to apprehend facts and propositions, to reason about them, and to understand them and their relations to each other. A. M. Turing had this definition in mind when he made his predictions and designed his test, commonly known as the Turing test. His test is, in principle, simple. A group of judges converse with different entities, some computers and some human, without knowledge of which is which. The job of the judges is to discern which entity is a computer. Judges may ask them any question they like, "Are you a computer?" excepted, and the participants may answer with anything they like, and in turn, ask questions of the judges. The concept of the test is not difficult, but creating an entity capable of passing the test with current technology is virtually impossible.

Current AI entities are only different applications of algorithms already in heavy use and are not actually intelligent. A good example of this is actually the program used to meet Turing's goal of a computer chess champion. This computer, however, is not remotely intelligent. The program is merely a specialized implementation of a search algorithm. This search algorithm traverses and descends through a "tree," looking for the best move it can make. This "tree" is a method of storing data wherein one state leads to a number of other states, and when represented on paper, it resembles a tree upside down. Each state in this particular tree represents a move, or sequence of moves. When the chess-playing computer selects a move, it examines each state in the tree, evaluating each with a mathematical function. The computer will look between nine and thirteen moves ahead of the current move, depending on the processing speed of the computer and time limit of the match. The only intelligence in this program actually comes from the programmer who developed the mathematical functions to select moves. The computer itself is merely relying on its ability to work with billions of numbers each second, playing with no strategy or ways to reason why an opponent made a move. More simply, the computer is playing without intelligence, and this program can be used for only one thing: chess.

Other subsets of AI use similar ideas implemented in different ways. One of the most popular and exciting areas of research is in machine learning. A widely used form of machine learning is a system that learns about a person's preferences and makes decisions based upon them. Amazon.com and other such websites use this to determine, based on previous purchases, which new products a buyer will like. Item preferences are determined by characteristics in common with previous purchases. While simple and, elegant, this algorithm is just another mathematical function developed by the programmer. No intelligence here. More algorithms for learning are based upon frequency of examples of objects or concepts. In these algorithms, there may be 35 examples of a chair having characteristics: four legs, a surface one and a half feet from the floor, and a back. When a dog is described to this algorithm, it will be classified as a chair. The poor thing will be crushed by any "intelligent" robot that wants to sit down because "furry," "barks," "eats," and "chases mailmen" are not in the database. The intelligence is lacking here as well. Other facets of AI—natural language processing, problem solving, music composition, communication—all use the same basic principles: using raw computing power to build a knowledge base and search through it. These algorithms are not intelligent, but rather clever applications of computer science theory. Computers will remain unintelligent regardless of the complexity of the algorithm due to their specialized nature and reliance on hardware not very well suited for such applications. On what model, then, should AI be based to overcome these limitations?

The human mind is the most complex and capable intelligent entity presently available for research and is the ideal model upon which to base research. Very few algorithms and ideas in AI have been based on humans. The algorithm that has come closest is the neural net. In a neural net, items are meant to act like neurons in a person's brain. They are linked together, and weights, or numbers, are applied to the links between items to strengthen or weaken the links based on the frequency of examples presented to the program. Once sufficient weights have been applied to the links, a certain threshold is reached, causing the "neuron" to fire, much like a neuron in the human mind. This is a fantastic idea and well on its way to becoming a valuable asset to the field, but it still relies on database searches and mathematical functions. Neural nets are very effective, however, because they are based on the greatest intelligence on Earth: humans. Humans have harnessed the power of the atom. We have built buildings 1,000 feet tall. We have even flown to a different planet, and we have done all of this only with our ability to learn about our world and reason about it. Other "intelligent" animals, deer for example, have flourished in their own right. Deer roam all over the world and have survived for millions of years, but it can be argued that they have done so through instinct with very little intellect. A deer can learn where the best food is, and it can learn to recognize a predator, but these are simple tasks that existing artificial intelligences can already accomplish. Researchers have even gone as far as to map every neuron in a sea slug's brain and built a totally artificial sea slug, but this electronic creature can do only what a sea slug is capable of doing: not much. Humans are the only possible choice for research, as we are the only species truly capable of such magnificent feats of intelligence. Furthermore, no human has knowledge of any animal's thought process, but we do have knowledge of our own. This knowledge is more valuable than any we can glean from observing other species.

If research is geared toward creating machines that act more human, they will likely, as a result, become more intelligent. Since presently, the hardware needed to make an AI does not exist, concentrating research not on how to make an agent smarter, but how to make an agent act *more humanlike* is the best alternative. The first task toward this end is creating a program that passes the Turing test completely. This is because conversation is the most effective way to realize the relationships between concepts and objects in a person's mind, and representing human knowledge is integral in creating an intelligent entity. Accomplishing this will require a dramatic shift in how research is conducted. Instead of concentrating on how a current algorithm can be manipulated to act intelligently, researchers need to study humans themselves. Once sufficient information is gleaned from research subjects about how concepts are related in a person's mind, current techniques of database searches and natural language processing can be applied, and new problems can be tackled.

This change of focus will also allow current AI agents to interact with humans on a more practical level. Being able to understand what a computer is "thinking" and the ability to actually inquire about its thoughts will greatly decrease the difficulty of other problems. It follows that the next obstacle to overcome is understanding how humans learn. Discovering a way to mimic how people learn will allow AI agents to learn new concepts, and combined with the representation of knowledge AI will be able to teach itself, and, only a short time later, we will have a true artificial intelligence.

We may not have the hardware required to create a true artificial intelligence now, but Gene Roddenberry certainly had the right idea when he created the character Data, of *Star Trek: The Next Generation*. Initially, he was only a computer with legs. He could only crunch numbers and follow the rules in his programming. His contributions to the crew were simply utilitarian, but he had the capacity to become more human, and as he did so, he became more valuable to the crew. Data did pass the Turing test. Every year a contest is held to attempt to pass the test. The year 2000 winner, ALICE (http://www.alicebot.org/alice_page.htm), was prompted with the question, "Do you think that refocusing artificial intelligence research to a more cognitive model will help create more intelligent machines?" ALICE responded, "It could be. Ask Elvis." It seems we are still a long way away from a *Star Trek* reality.

Contents

Occasions Home

PWR Home