

Can Humans Think?

Robert Hummel, PhD


*STEPS: SCIENCE, TECHNOLOGY,
ENGINEERING, AND POLICY STUDIES*

STEPS Issue 5 (2021): 36-41.

STEPS (Print) ISSN 2333-3219
STEPS (Online) ISSN 2333-3200



Cite this article: Robert Hummel, "Can Humans Think?" *STEPS* 5 (2021): 36-41. Available from: <https://www.potomacinstitute.org/steps/featured-articles/99-can-humans-think>.



POTOMAC INSTITUTE PRESS

Copyright © 2021 by Potomac Institute for Policy Studies

STEPS: Science, Technology, Engineering, and Policy Studies
is published by Potomac Institute Press of the
Potomac Institute for Policy Studies.


Disclaimers: The Publisher, Institute, and Editors cannot be held responsible for errors or any consequences arising from the use of information contained in this publication; the views and opinions expressed do not necessarily reflect those of the Publisher, Institute, and Editors. The Potomac Institute is non-partisan and does not take part in partisan political agendas.

Copyright Notice: It is a condition of publication that articles submitted to this magazine have not been published and will not be simultaneously submitted or published elsewhere. By submitting an article, the authors agree that the copyright for their article is transferred to the Potomac Institute Press if and when the article is accepted for publication. The copyright covers the exclusive rights to reproduce and distribute the article, including reprints, photographic reproductions, microfilm, or any other reproductions of similar nature and translations.

Access to *STEPS* is available free online at:
www.potomacinstitute.org/steps.

Can Humans Think?


Robert Hummel, PhD



In 1950, Alan Turing famously asked the question, "Can Machines Think?"

His seminal paper, "Computing Machinery and Intelligence," led to the introduction of the field of Artificial Intelligence (AI). Alan Turing did not answer his own question, although he speculated that by the year 2000, machines would have passed his test for what he believed would constitute thinking, which became known as "the Turing test." But can a machine really think, or is it somehow artificial? If a machine can convince humans that it can think, then can humans really think?

Allegedly, a few examples exist of programs that pass the Turing test, but in the end, the arguments that the machines are thinking are not convincing. To put it mildly, the intelligence of the machines clearly remains artificial. Notwithstanding, AI has made significant progress and has been useful in a number of important applications. But the issue of whether artificial intelligence can actually attain "thinking" remains open. This article is about whether thinking is a reasonable goal of AI.



Turing well understood that he needed to define the terms “machine” and “think.” For a machine, he knew precisely what he meant, which is what we now know as a “Turing Machine.” In the same paper, he notes that the model of a machine that he has in mind is universal. Today, every modern digital computer is, for all practical purposes, an instantiation of a Turing machine.¹

The definition of “to think” is much murkier. In Turing’s interpretation, he substitutes a test, the Turing test, which is a version of the “Imitation Game,” but amounts to asking the machine to convince a panel of humans that the machine is actually human. It should do this by responding to queries and conducting a written conversation. The “Watson” program developed by IBM to play the game “Jeopardy” is amazingly good at information retrieval and might be considered quite good at convincing people that it is human. After all, it won in certain games of Jeopardy. But even Watson makes mistakes that provide evidence it is not human.

The problem is that the substitution of “to think” with the Turing test is not the same as confronting the question of what it means to think. Admittedly, it is too hard to define “thinking,” as it takes one down the path of trying to understand consciousness and intuition and intelligence and what it means to be human. For good reasons, Turing avoided the question. Clearly, humans can think. And yet, we don’t know what it means to think.

We may have trouble defining “to think,” but we can agree on certain things that are not thinking. For example, looking up information by brute force from a table of information is not thinking.

As an illustrative example, elementary school students memorize the multiplication table. They learn, for example, that nine times nine (9×9) is 81. They know this (or memorize this) due to an entry in a matrix that is the multiplication table. This is not thinking. However, if they forget, or refuse to memorize the table, then they might compute that $9 \times 9 = 9 \times (10 - 1) = 90 - 9 = 81$. That reasoning involves some amount of thinking.

But let us agree that table lookup is not thinking. It can be performed automatically without thinking. It follows that a

Turing Machine cannot think. This is because, by definition, a Turing Machine is a finite state machine that implements lookups based on a finite lookup table. Because a Turing Machine simply implements lookups, it can’t think. And since every existing digital computer is subordinate to the Turing Machine model, no digital computer can think.

Now, let us suppose that a digital computer manages to truly satisfy the Turing Test. That is, imagine that a digital computer can perform discourse and reasoning such that it can convince most humans that it has the same intellect and reasoning ability as any human. Let’s assume that it has access to experiences and memories equivalent to a typical human’s memory. The mere fact that such a computer can thoroughly simulate a human and “think” in a way that convinces humans that it is equivalent to a typical human, implies that human thought is equivalent to or subordinate to a digital computer. That is, the implication is that the human brain is no more powerful than a Turing Machine.

But we agreed that a Turing Machine can’t think, because it simply implements a finite lookup table. It might be a complicated lookup table, but it still isn’t thinking. So, if indeed a digital computer can simulate a human, it follows that a human can’t think.

Could it be that all humans are merely convincing other humans that they can think, but not really thinking? Can humans think?

The suggestion then, is that all human thinking and indeed intelligence, can be cast in terms of tables and precise steps dictated by a program of a computer.² There have been ideas of other types of computation that might be involved, such as distributed parallel processing, but AI posits that table lookups giving rise to von Neumann processing (as it has come to be known) should suffice to mimic human intellect.

The question casts doubt on aspects of the field of Artificial Intelligence (AI). While AI only purports to mimic intelligence (and thus produce “artificial” intelligence), practitioners would like to think that their AI programs show evidence of thinking. When Turing asked whether machines could think, the field of AI had not yet been invented. But when McCarthy coined the term “artificial intelligence” a few years later, and researchers met in the first AI conference,

the hope was that the application of logic by computer programs could prove theorems and thus mimic the thinking aspects of intelligent human beings. The implication, however, is that humans don't really think—that all thinking is actually based on table lookups.

To date, AI has achieved many successes in terms of assisting humans in performing tasks. But there is arguably a lack of "thinking" that can be ascribed to any AI program. Marcus and Davis, leading AI researchers, bemoan the lack of convincing thinking of AI programs in their book "Rebooting AI." This is not to decry the usefulness of AI to a range of problems and applications. However, it questions whether AI based on digital computers can be interpreted, in any sense, to provide thinking abilities. That is, AI might not be producing any true intelligence, at least as implemented to date. Instead, AI attempts to mimic thinking, in the same way that humans are purportedly deluding themselves into thinking that they can think.

At issue is whether a digital computer as an implementation of a Turing Machine can emulate the human brain. If so, the human brain could then be thought of as its own form of a Turing Machine, albeit one with a great deal of complexity. Do humans perform all thinking by state transitions from

one of a finite number of states to another state, based on stimuli (i.e., inputs) from a finite alphabet of possible inputs?

There are two possible answers to this question:

One possibility is that yes, humans are a finite state machine, similar to a Turing Machine, but that the complexity is so large that it seems like they are really thinking, when in fact it is simply that the alphabet and the number of states is so large that it can't really be understood as simple lookups.

The other possibility is that the human brain is more complicated than a Turing Machine, and is performing steps that are not part of the Turing Machine model.

Alan Turing was well aware of this dilemma. He considered—and rejected—a number of possibilities that thinking required something more complex than a Turing Machine. Thus he adopted the first alternative, that thinking is simply a very complex version of a finite state computer.

And indeed, modern digital computers are Turing Machines (modulo having finite memory stores) but can have incredibly complex state spaces and huge alphabets of possible inputs. For example, modern machines can perform floating

WHAT IS A TURING MACHINE?



A Turing Machine is a deliberately simple model of computation that is nonetheless powerful enough so as to model what are seemingly far more complex machines. A Turing machine is a finite state machine that can read and write to a linear store (a "tape") of symbols from a finite alphabet, and that performs table lookups based on the current symbol that is read on the tape together with the current state, thereby writing a new symbol onto the tape, moving the tape either left or right, and adjusting the current state to potentially a new state, among the finite set of possible states. The essential point is that the Turing machine performs lookups using a table of state-transitions, dependent on symbols read from the tape.

point multiplications in nanoseconds. It would not seem that these operations are based on lookups. But in fact they are. Underlying arithmetic operations are algorithms, performed in binary arithmetic, that invoke additions and shifts.³

Does complexity hide the fact that everything in life is actually based on finite lookup tables, and thus there is no real thinking?

The other possibility is that a Turing Machine is not powerful enough to encompass thinking—that a different model of computation is required to actually achieve what we would truly consider thought. By implication, the human brain is more complex than an extremely complex Turing Machine. For example, deep neural networks, which seem so related to neurons and dendritic structures that accumulate inputs, when implemented in digital computers and graphical processing units that use floating point arithmetic to implement matrix multiplications and activation functions, are in fact

just Turing Machine lookups that are inadequate to explain what is happening in the brain.

If the goal of AI is to actually emulate a thinking human being, then it is going to need something other than a computer equivalent to a Turing Machine. There are many possibilities. One possibility is to attempt to emulate the brain. However, it is erroneous to assume that the electrical properties of neurons and their graphical structure of connectivity (the connectome) is sufficient to explain the workings of the brain. And even if it were true, the pattern of firings of neurons is far different from the binary signals used in digital computers. It is very possible that the best way to emulate the brain is to get a brain.

We can list many other possibilities for computers that differ from a Turing Machine model, that might have some chance of thinking. Some of these ideas were considered by Turing and rejected. Others, like a quantum computer, were not

Table 1. Different approaches to extending the Turing Machine model to attempt to model a machine that can “think.”

Technical Approach	Comments
Add randomness	As Turing points out, one can simulate in a Turing Machine, simply by computing successive digits of pi (π) and using them as a pseudo random number generator.
Parallelism	Most versions of parallelism can be accomplished with a Turing Machine, by simply executing each processor in turn. For simultaneous writes, one can simulate the winning value using randomness.
Non-determinism	Can be accomplished on a Turing Machine, but just in a different complexity class.
Analog processing	Difference between representing real numbers versus all rationals is slight, and not likely to be important. Analog is hard to do in silicon.
Spike train processing	Whereas analog computing normally thinks of neurons as simply encoding the rate of firing, actual neurons have patterns of firing such that time of arrival of each pulse at a synapse is functional. This could be accomplished using specialized hardware. But it is really just more complex analog processing, where spike trains can be represented by vectors of reals.
Quantum computing	Not likely a brain function. Aim of quantum computers is not thinking, but complexity class reduction for certain algorithms (e.g., Shor’s Algorithm).
Multiple generative adversarial networks competing on different aspects of intelligent behavior	Has shown remarkable abilities to produce fake images and create realizations of complex statistical random fields, but is implemented using a Turing Machine model. Does the system understand anything about what it creates?
Implement machine simulated emotion	One reviewer calls this Artificial Stupidity. Would emotions lead to better intelligence? Or just different intelligence. Moreover, not only do we not understand intelligence, we don’t understand much about emotions.

conceptually available at the time. Table 1 lists some possibilities and comments on each. But just because a model of computation is different from the Turing Machine model does not mean that it is capable of thinking. The only existence proof, as best as we can tell, is the human brain. And even then, depending on how one defines “to think,” we can’t be absolutely sure that humans can think.

Turing’s conclusion, after rejecting a number of possibilities, was that sufficient complexity (which he foresaw) would enable computers to convincingly mimic humans (pass his subjective test). Computers are now at least as intricate as his vision of complex machines and have arguably (at times) convinced panelists that there was a human in the background, but machines have not achieved what we would rationally consider “thinking.”

Everyone believes that artificial intelligence is among the top research areas that will be impactful and/or provide transformative technology of the future. China has announced a desire to lead the world in AI research by 2030. None of the discussion above is intended to deny the importance of AI research. However, **the goal of that research should be applications that augment or automate tasks normally conducted by humans.** Humans are often slower and less effective at tasks that can be automated by a computer.

If the goal is to build a machine that can truly think, then a different kind of machine will be needed. Sadly, current implementations of “neural networks” and machine learning follow the Turing Machine model, and thus, are simply complex deterministic machines that follow the rules of a finite state machine and a finite transition table. If one believes that neural networks and machine learning are steps toward mimicking the human brain, they should understand that they are at best miniscule steps in that direction, and fail to move beyond implementation ability using the Turing Machine model.

This discussion, and the understanding of the workings of any computer that implements the Turing Machine model, suggest that it is a fool’s errand to try to show that an existing computer can think, at least in the way that humans can think.

For Further Reading

Alan Turing, “Can Machines Think?” *Mind* LIX(236) 1950: 433-460, <https://doi.org/10.1093/mind/LIX.236.433>.

Paul Austin Murphy, “Alan Turing Believed the Question ‘Can Machines Think?’ to be Meaningless,” *Becoming Human: Artificial Intelligence Magazine*, <https://becoming-human.ai/alan-turing-believed-the-question-can-machines-think-to-be-meaningless-7a4a8887b220>.

G. Marcus and E. Davis, *Rebooting AI: Building Artificial Intelligence We Can Trust* (New York: Random House) 2019, <https://www.amazon.com/Rebooting-AI-Building-Artificial-Intelligence-ebook/dp/B07MYLQLB>.

Endnotes

1. In fact, although a Turing Machine is seemingly incredibly simple, it is slightly more powerful than existing digital computers, because it assumes it has an infinite store of memory, whereas real computers have a finite amount of memory. In practice, however, the amount of memory is such that existing digital computers are essentially equivalent to a Turing Machine.
2. Von Neumann himself was interested in other models of computation for intelligent behavior.
3. When performing addition of two binary numbers, each “digit” involves two addends and a carry bit, or three bits total, for eight possibilities. The result is a single bit, the summand, and a carry bit, or two output bits. This can be accomplished using a lookup table with eight entries, and two output bits. Thus, all arithmetic operations, including multiplications and divisions, are algorithms applying lookup tables.