Chapter VII
AI's Real Capabilities

We have reviewed the functioning of various types of AI, claims made about them, their performance, and issues surrounding their growth. We also investigated the theory of knowing employed by AI, especially Generative AI, and compared it to human knowing (the first of our four questions from Chapter I). In this chapter we look at what AI can do, in light of its theoretical limitations stemming from the theory of knowing that it employs. This includes an analysis of the ways its limitations will affect how it interacts with humans. The three remaining questions asked in the introduction will be covered in the next chapter. Here we consider the following topics:

- Symbol manipulation vs. interaction with reality
- The difference between knowing what things are and how they behave
- The difference between aiding human knowing and replacing it
- Creative thinking and understanding vs rote or algorithmic manipulation
- Effect of utilizing the wrong paradigm for human knowing
- Locked into the past vs. looking to the future

This will enable a better perspective on the future of AI, and help us to answer the questions about AI's practical side.

Symbol manipulation vs interaction with reality

The goal of human knowing is always to *know something about reality*, whether or not it has any operational value. Neither an animal nor AI seeks the reality of the real. For that reason, General AI will never be realized.

Silicon Valley technology pioneer Keith Diefendorff has an array of patents in microprocessor architecture, reduced instruction set computing (RISC), optical interface technology, and other areas. He led the team that created the PowerPC micro-family for IBM and Motorola and later Apple. He also served for nearly a decade as editor in chief of the *Microprocessor Report*. He knows the researchers working in the AI area. According to Diefendorff:

> [the AI researchers] are getting nowhere with general purpose tasks. AI is proving good for specific niches… Games, in fact, are what they do best.[227]

As we observed earlier, the function of computers is fundamentally misunderstood:

> …the links between computational symbols and their objects are indefinite and changing. *The map is not the same as the territory*. The links between symbols and objects have to be created by human minds. Therefore, computations at the map level do not translate to reliable outcomes on the territorial level.[228] [Italics added]

This means, of course, that success in game playing is qualitatively different than dealing with the real world:

> For the game of *Go* or chess or some routinized task, the symbols and objects are the same. The white and black stones on the *Go* board or the pieces on the chess board are both symbols and objects at once. The map is the territory. ….in order to have correspondence between logical systems and real world causes and effects, engineers have to interpret the symbols rigorously and control them punctiliously and continuously. Programmers have to enforce an interpretive scheme between symbols and objects that banish all slippage.[229]

Because they do not have sentient intelligence, when dealing with the world, AI and computers *must* utilize symbols. These symbols function as signs for response, programmed in the case of computers and AI:

> A digital computer is a device which manipulates symbols, without any reference to their meaning or interpretation. Human beings, on the other hand, when they think, do something much more than that. A human mind has meaningful thoughts, feelings, and mental contents generally. Formal symbols by themselves can never be enough for mental contents, because the symbols, by definition, have no meaning (or interpretation, or semantics) except insofar as someone outside the system gives it to them.[230]

The machines, in other words, have no connection to what things are in reality, as we saw in Chapter VI; they can only manipulate symbols and then take some sort of programmed action, which could include writing text or maneuvering a self-driving car.

What is the conclusion? AI technologies, including neural networks, operate differently than human intelligence, and only mimic it in ways that are

very fragile. Obviously, human recognition of things, as realities, based on primordial apprehension, does not have this problem. AI is compelled to use an algorithmic approach to knowledge because that is the only possibility for a machine-based system. This can be very useful, but is not a substitute for human intelligence. For example, modern word processors can help with spelling and grammar, but won't compose your essay. Modern spreadsheets provide many built-in routines and can do complex calculations, but won't help you determine what you need to calculate. Of course, chatbots will compose an essay for you, but whether it makes any sense, says things that are wrong or just made up, or is based on solid research rather than a superficial rehash—that you must determine. In most cases, if you need a quality product, you will have to do the research and write it yourself.

It is often said that information is based on "surprise," that is, new information is not just a repeat of old information or data, but contains something new and different. Thus, simple extrapolation from existing data, though possibly giving a new perspective on it, is not supplying us with new information. As we saw in Chapter I, Chatbots and generative AI is based on collection and in some way processing "old" information. It is incapable of generating new information. In this sense, it is obvious that humans are unique because creating new things, new theories, and new explanations, and dealing with new situations, is just what they do as part of living.

In summary, AI and Machine Learning will, like other technological capabilities, extend the power of human intelligence, but never replace it, because of the paradigm of sensible intelligence. AI must use an algorithmic approach to knowledge acquisition and processing, which is excellent for certain tasks and far exceeds human capabilities there, but is inadequate for tasks that require creativity or direct interaction with reality. Humans will not be replaced by machines, nor will machines ever become sentient, have human capabilities, or "souls". They will simply "amplify" human capabilities, as they have always done in the past.

Difference between knowing what things are, and how things behave

There is a profound difference between knowing *what things are*, and *how things behave*. Though historically many have thought that these two are the same, or at least that one immediately leads to the other, in fact they are distinct. Knowing *what something is* engages the transcendental order of human knowing, how the thing relates to other things, and the fact that it exists in reality

as a thing. Knowing *how something behaves* enables us to control it, or to make other things that behave in similar ways. That is, it operates at the phenomenological level. Knowing this, or equivalently, knowing how to make something that behaves in a particular way, is not operating at the most fundamental level of human knowing. AI systems perforce act only on the basis of how things behave, or nominalistically on the basis of names, but never on the basis of reality.

For many purposes, however, knowing how things behave is suitable. A factory robot does not need to know anything about reality; it just needs to be able to do a particular assembly job. Symbolic mathematics programs such as *Mathematica* only need to be able to follow certain rules and solve certain kinds of equations based on a repertoire of methods that they have. But this is an extremely important and useful capability.

Difference between aiding human knowing and replacing human knowing

To better understand the problems posed by AI, we need to draw a distinction between *aiding human knowing* and *replacing human knowing*. It is clear that computers have been doing more and more of the first for decades. Our modern technological society could not exist in anything like its present form without computer-based automation of functions at one time done by humans. A trivial example is telephone switching. Verizon alone says that it handles 800 million calls per day.[231] If a phone operator at a manual switchboard could handle 1 call per minute, and could work 8 hours per day, handling these calls would require about 1.6 million operators and a switchboard about 173m per side—an obvious impossibility. Modern banking and finance simply could not be done manually, nor could most modern engineering or scientific research problems. Modern medicine uses computers and automated systems extensively, but to aid doctors and other key medical personnel in making medical judgements. AI is steadily improving in key areas like image recognition and image processing, which are important in medicine but as well in many research and development fields, and in industrial settings.

Replacement only occurs with jobs that involve simple repetitive labor, such as continuous monitoring of vital signs. In all these cases we observe that computers and computer-based systems are supplementing or assisting human knowing and activities, not replacing them or doing something new. But this ability is vital to the functioning of our society. And the natural language

capabilities of AI systems make it easier for people to interact with the computer-based tools that they need.

Creative thinking and understanding vs. rote or algorithmic manipulation

AI tools such as ChatGPT can scan the Internet and assemble much information, even invent "facts", but they are not creative in the true sense. On the other hand, human knowing is nothing if not radically creative, even in simple everyday tasks, such as driving a car. And it is especially so for science, math, literature, music, art, and many other fields. The great advances in science always come when someone creatively breaks with established tradition. Einstein broke with establishment science with his theory of relativity. Heisenberg and others broke with established science when they formulated quantum mechanics. Gödel broke with old ideas about math by his Incompleteness theorem. Beethoven broke with established musical ideas with his symphonies. Renaissance artists broke with old traditions to develop new ideas about painting. Euripides broke with old traditions to write his plays. AI algorithms typically compose by looking at previous words and then guessing what the next word should be, following grammar rules. They cannot creatively and analytically think through a question, using information learned from reading and research, where a critical eye is needed to discern what is valuable and a view of reality is needed to synthesize new ideas. AI systems, therefore, are essentially stuck in the past, unable to advance knowledge or even apply what is known in a creative way. On a more prosaic level, anyone who has done a job other than something very routine such as assembly line work, can attest that creative problem solving is required almost constantly, because situations arise that are different than earlier experience or training. This is especially the case for project managers and supervisors, though even lower level workers frequently encounter problems with tools or supplies and need to figure out work-arounds.

The creative element plays out in literature in important ways:

> …an author like Shakespeare is great because he expresses vividly and concretely a particular time, place, and culture; and yet he transcends what is merely local and ephemeral and touches the perennial and universal concerns of humanity by means of what is immediate and particular.[232]

Truly original thinking involves selectively taking what is already known and using it to formulate new ideas, theories, explanations, or artistic works that apply to the real world, give insight into reality, and can be tested or verified, as

appropriate. It is not based on randomly combining ideas or facts, but on insight into the reality of a situation or problem.

A particularly interesting example is mathematician Georg Cantor's (1845-1918) famous diagonal argument, used to demonstrate that there are different infinities or *transfinite numbers*. The argument shows that the transfinite number $\aleph_0$ (the cardinality of the integers) is less than $\aleph_1$ (the cardinality of the real numbers). Cantor proved theorems about infinities considered as *real* things, with *real properties* that can be discovered. He was, then, talking about reality, not marks on piece of paper, as the nominalists would have it. Cantor's results were totally unexpected at the time, and no amount of pattern-recognition or random shuffling type of investigation of then-current mathematical ideas (if this could even be done) would have led to Cantor's results. As there are no physical infinities, the entire notion of infinity, as something real, and transfinite numbers, though able to be grasped by human mathematicians, make no sense to an AI program; for it they are just names in a nominalistic paradigm. They certainly are not the correlate of any "sensible impressions" (in Hume's parlance). Though symbol manipulation programs such as *Mathematica* can operate with infinity, they have only the capability to follow strictly logical rules for manipulations involving it. For humans, though, transfinite numbers are real:

> A transfinite number, an abstract concept, are not sensed qualities. But they are intellectively known as something real, and as such are constituted in the impression of reality as such.[233]

The programs on the other hand have no concept of infinity large or small, only the rules that they employ.

We have art schools that teach drawing and painting techniques, music conservatories the teach composition, and college curricula that teach creative writing. While all of these can teach students fundamentals and even advanced techniques, they cannot guarantee that their students will become great artists, composers, or writers. Why is this? Because these programs can only impart basic rules, *but not the insight and inspiration that sees reality and turns it into great art, music, or literature*. Machines can also be programmed to follow these rules, but cannot be programmed to sense reality, essential to creation of great art.

In some cases, creation of great art isn't the goal—a commercial product is all that is needed. AI has the capability to respond to requests for many types of visual material, synthesizing the product by utilizing its database of existing

photographs and other images. At this superficial level, sometimes the resulting product may be quite adequate. In others, interaction with a human designer is necessary or preferable.

Effect of utilizing the wrong paradigm for human knowing in AI

The main effect of using the wrong paradigm of knowing is that AI will be expected to do things that it will never be able to do. This will cause expenditures of money and time that can never come to fruition, and futile attempts to substitute AI devices for people. In fact, this latter has been occurring for some time, with extremely frustrating results. The automated response systems used by many banks and other commercial entities are a case in point. The reader has probably had the same experiences as the author: your call is answered, and you are given a menu. Of course, the question for which you require an answer isn't on the menu. Or else the system will ask you what you want to know. When you say what it is, you are given something else entirely. After many frustrating minutes, you might get transferred to a real person, who actually can deal with your problem. Worse than this inconvenience is the issue that an AI-controlled system in a critical environment may behave in totally inappropriate ways, leading to catastrophe.

Locked into the past vs looking to the future

With respect to human knowing, all categories of AI are backward-looking rather than forward-looking, because they are based on existing knowledge. None has the ability to create new visions of reality, new theories. This does not mean that they cannot be used to make predictions or forecasts about the future; even simple regression analysis can do that. And they can, of course, be used to enable us to see things that we otherwise could not see, such as simulations of the evolution of the universe. But these simulations are based on *current* theories, e.g., about the constitution of the universe and the laws governing it. What this statement about the capabilities of AI means is that AI cannot advance human knowledge in any theoretical sense, that is, develop new theories about reality. It can only use existing knowledge to give us answers. This is an important contribution, but ultimately limited in its scope. And it can only do this function as an adjunct or assistant to a human seeking to answer a question or solve a problem.

Some interesting applications of AI

The ability of Generative AI to digest large amounts of information and seek patterns does suggest that there are areas where it may be very useful, even if

General AI is forever out of reach, and chatbots are unreliable. We consider a few here that may be indicative of the value of this technology, or at least its pattern recognition capabilities.

Rogue wave detection

AI is being used to help forecast "rogue waves", which are especially large waves that can swamp even large ships:

> A new tool from UMD [University of Maryland] researchers could give mariners and offshore oil platform residents up to a five-minute warning of these "rogue waves" so the can batten down the hatches or position their ships to survive the wild ride…To create the tool, they trained a neural network…to distinguish ocean waves that will be followed by a rogue wave from those that will not. The training data consisted of billions of regular waves and thousands of rogues, recorded in 14 million 30-minute-long samples of sea surface elevation measurements from Pacific Ocean buoys. After training, the system correctly predicted the emergence of 75% of the rogue waves one minute into the future and 73% of rogue waves five minutes into the future.[234]

Since by definition rogue waves appear to emerge from nowhere, and are unsuspected, this forecasting ability, which relies on pattern recognition capabilities, is a step forward in maritime safety.

Ancient antibiotics

Antibiotic resistance has become a major problem in pharmaceuticals. The idea behind one use of AI in this context is that the bugs may no longer have any resistance to "old" antibiotics. Therefore, searching the DNA of extinct species may reveal protein fragments (peptides) that have antibacterial properties. Cesar de la Fuente, a professor at the University of Pennsylvania, and his colleagues have created a deep learning algorithm to

> …comb through enormous genetic databases to find peptides, or protein fragments, that have antibacterial properties. They have used this method to analyze animal venoms, the human microbiome and archaea, an underexplored group of microorganisms. They have also mined the genetic codes from fossils of long-extinct animals and humans, including Neanderthals and Denisovans.[235]

The work is paying off. The researchers have found hundreds of peptides able to cure sick mice. (Testing on mice is usually the first stage to determine if a drug is effective and safe for human use):

> One of [the peptides] was mammuthusin, identified in the genetic code of *Mammuthus primigenius*, a species of mammoth that last roamed the Earth about 4,000 years ago. The researchers discovered the peptide after mining a National Center for Biotechnology Information database of DNA sequencing data obtained from fossils of extinct animals. In experiments, mammuthusin was as potent as polymyxin B, an antibiotic often used as a last resort for serious infections…[236]

Think about this contribution to your health next time you visit a natural history museum and gaze at the skeletons of the extinct animals there!

"Climate in a Bottle"

Another potentially useful application of AI is a new model from Nvidia, dubbed "Climate in a Bottle." The model uses 50 years of "high resolution physical climate simulations and estimates of observed atmospheric states." The objective is to allow projections of climate conditions many decades into the future:

> The model, branded by Nvidia as cBottle for "Climate in a Bottle," compresses the scale of Earth observation data 3,000 times and transforms it into ultra-high-resolution, queryable and interactive climate simulations, according to Dion Harris, senior director of high-performance computing and AI factory simulations at Nvidia.[237]

The result, which Nvidia calls a "digital twin of our planet," is impressive from a computation standpoint, and in theory has many applications stemming from its ability to forecast climate trends. For example, it could be used by insurance companies to gauge risk, or predict conditions that might lead to shortages of food and water in certain areas.

But a careful reading of the project description shows problems: (1) It is not based on actual data, but on simulations. Climate models, used for simulations, do not have a good track record of accurate predictions. (2) The record being simulated is only for the past 50 years, a very short time for climate change, which has long time constants, on the order of hundreds, thousands, tens of thousands of years, and even longer. The time span chosen, 50 years, is as likely to be noise as signal. Going back 50 years, to 1975, the record shows a warming trend. If the model is trained on this data set, it will predict warming effects in the future. But the actual climate had a cooling trend in the 30 previous years. And going back 600 years—a blip in Earth's history—we had the Little Ice Age, a period of significant cooling. (3) There is no guarantee that the model won't produce hallucinations, that is, predictions that are not based on the data but

just made up. (4) It is impossible to test the model because any reasonable test would itself require 50 or more years. So even if an AI application sounds like a good idea, it is still necessary to think through the assumptions behind it to determine if AI is an appropriate tool, and whether we have enough real world factual data. Otherwise we are back to the syndrome "if the only tool you have is a hammer, everything looks like a nail."

Resume scanning and hiring decisions

AI is being used to assist companies with hiring decisions. According to reports, most large companies now use AI in some capacity for this purpose, undoubtedly because they receive so many applications and are seeking a way to expedite the hiring decision process. Presumably, the AI is trained on resumes that contain "desirable" qualifications, which would allow it to "grade" the resumes submitted. Reporting on this use of AI, one commentator noted:

> It's not just a few mega corporations using AI to weed out job candidates. A study published by *Resume Builder* last October found that half of all companies are already using AI in the hiring process and predicts that this will increase to 70% by the end of 2025. The study surveyed 948 business leaders who work at a company with more than 21 employees. "Today, 82% of companies use AI to review resumes, while 40% employ AI chatbots to communicate with candidates. About 23% use AI to conduct interviews, and 64% apply AI to review candidate assessments," *Resume Builder* reported. "Additionally, 28% of companies use AI for onboarding new hires, and 42% scan social media or personal websites as part of the hiring process. Only 0.2% of companies report not using AI in their hiring practices."[238]

While this sounds like a good application, there are problems. First, no one seems to be asking (or even interested in) the question of whether the AI-based screening process is as good as the traditional Human Resources (HR) process, handled by an actual person. A qualified HR member can scan resumes and quickly find qualified applicants, reading between the lines, so to speak, and confirm the decision with a short in-person interview. This is something the AI is unlikely to be able to do. In addition, there is the cost savings question. While it would seem that the AI-based process saves money compared to the traditional, labor-intensive HR method, when all of the expenses of the AI program are accounted for, including selection of material for training, the training process, the cost of using the AI program, ongoing maintenance of it, and any need for human review of AI results, is there a net savings? Is there any way to determine

the relative value to the company of those chosen by the AI program compared to those selected by the traditional method?

Closely related to this problem is a second: the AI is biased in its decisions:

> Large language models (LLMs) are increasingly deployed in high-stakes hiring applications, making decisions that directly impact people's careers and livelihoods. While prior studies suggest simple anti-bias prompts can eliminate demographic biases in controlled evaluations, we find these mitigations fail when realistic contextual details are introduced…When these biases emerge, they consistently favor Black over White candidates and female over male candidates across all tested models and scenarios. Moreover, models can infer demographics and become biased from subtle cues like college affiliations, with these biases remaining invisible even when inspecting the model's chain-of-thought reasoning.[239]

In addition to the issue of legal liability for the company, it reinforces the above analysis that the best candidate may be passed over in favor of one less capable or qualified, which in the long run could cost the company far more than any amount saved in the hiring process. Overall, it is difficult to avoid the conclusion that AI is being used without a critical evaluation of its actual value.

Creation of advertising copy

The Clorox company has been experimenting with AI to create advertising copy, mainly visuals, for promotional purposes. The staff used an AI tool to create an ad for their Hidden Valley Ranch dressing for chicken wings. The first image was not so appetizing, but a different prompt gave a better result (Figure VII-1). This did, notably, require human intervention to obtain the desired image.

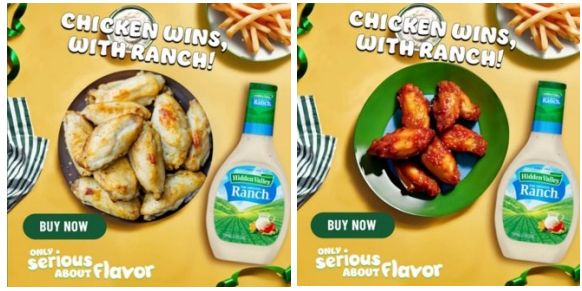

Source: *Wall Street Journal*

Figure VII-1. Initial chicken wings ad (left), final chicken wings ad (right)

AI, mainly Generative AI, is being used at all stages of R&D to summarize product reviews, come up with new ideas, create digital prototypes for copy, and creating more ads. It sort of works and sort of doesn't:

> "If you go in with the expectation that the AI is as smart or smarter than humans, you're quickly disappointed by the reality," said Eric Schwartz, Clorox's chief marketing officer. During brainstorms, the AI tried to push the idea of "bleachless bleach," he adds, which isn't something that would actually work in real life.[240]

The net result seems to be that Clorox found some value in using AI, but it hasn't led to any layoffs, and Clorox made no estimate of the money saved or sales gains and profit realized. It appears that AI in this case is just another way to assist human workers in their jobs, as computers have been doing for many decades. Clorox did not say how much they have invested in their AI tools, or how much they plan to invest in the future.

* * *

We have seen some interesting applications of AI, but none that are killer apps, and none for which quantitative data about labor or cost savings were given, or profits realized. We have also seen some respects in which AI is limited, including AI being stuck in the past. Where does this leave us? What will AI do? We answer these questions in subsequent chapters. Chapter VIII looks at three of the questions posed in the Preface, dealing with technological, economic, and ethical limitations of AI, as well as the question of whether AI will truly be a world-changing technology. Part Three of the book deals with the question of how AI will interact with society and the implications that this has for AI's impact on the world.