SUMMARY

HOW TO CREATE A MIND

RAY KURZWEIL

QuickRead

FREE BOOK SUMMARIES

Summary of ”How to Create a Mind” by Ray Kurzweil

Written by Nicolas Stewart

Breaks down the inner workings of brains, both human and artificial.

Introduction

Have you ever noticed how strongly our senses are connected to memories? Smell in particular seems uniquely capable of conjuring vivid and detailed memories that you may not have thought about in years.

This is because the brain doesn't store memories as singular complete pictures, but rather memories are stored in pieces. And recalling one of those pieces can trigger the recollection of the entire memory.

Techniques you can learn to help improve your memory typically focus on this principle, finding one of the key triggers of a memory that will allow it all to flood past.

Similarly it's these sorts of innocuous seeming triggers that can often cause traumatic memories to come rushing back for PTSD patients.

The pieces of a memory are often stored in a specific order as well. Presumably you could sing your favorite song verbatim with little effort. But could you sing it backwards? Or start only with the 2nd verse? Potentially, but it would require much more effort and conscious thinking.

This is because our memory is largely based on our brains' pattern recognition abilities. When we teach kids the alphabet, we don't teach them the letters in random order, we teach it in sequence, because that orderly pattern helps the brain memorize it. Memorization tricks often teach you to memorize things, like the list of US presidents or the names of all the states, by making a song out of it.

By providing an orderly, sequential pattern we exploit our brains' natural memory storage process to help us memorize things.

The Neocortex is The Center of Higher Brain Functions

The neocortex is unique to mammals and is involved in many of the higher brain functions, namely sensory perception, cognition, spatial reasoning, motor commands, and language. The usually large size of the neocortex, and the frontal lobe, are responsible for the uniquely high degree of intelligence in humans.

As we explored with memory, information in the neocortex is stored and accessed based on hierarchies and linear functions.

The knowledge of how to act out complex behaviors and actions, like using a spoon to eat cereal, are stored via cortical columns, arranged in specific patterns and when accessed do so based on heuristic functions, wherein the information is retrieved piece by piece e.g. grab the spoon, scoop cereal, open mouth, deposit cereal, chew and swallow.

Cortical columns are made up of clusters of neurons, each column containing approximately 60,000 neurons. Added together all the cortical columns in the neocortex are composed of some 30 billion neurons.

One of the primary functions of cortical columns is pattern recognition, one of the most fundamental and important elements of human cognition. When your brain sees a cereal box it utilizes the pattern recognizers in cortical columns to search your memory for the patterns you see on the box, it’s pictures, it’s letters, etc. And in doing so is able to retrieve the information necessary to recognize that you are looking at a box of cereal.

Our senses constantly receive input, and this input is compared with past experiences and patterns. To keep up, the brain’s recognizers fire continuously!

Memory, or memory retrieval, is therefore primarily based on pattern recognition. This is also the reason why you might occasionally think you recognize something despite having never seen it before, because your brain recognizes certain patterns, but can’t match them to any one specific thing in your memory.

The Neocortex is Connected to All Parts of The Brain

The neocortex is divided into regions, demarcated by the cranial sutures of the skull, these regions are the frontal, parietal, occipital, and temporal lobes. Each lobe governs a different function. The occipital lobe contains what is known as the visual cortex, which interprets the data sent from the eyes.

But the path there isn't a direct one, because of this the neocortex is directly connected to virtually all parts of the brain. So when your eyes take in light the signals from your eyes are first received in the midbrain which reads out the signals and controls where your gaze shifts, before being sent through the midbrain to the appropriate lobe of the neocortex, in this case the occipital lobe.

The neocortex then communicates with the temporal lobe, where episodic memories are stored, the part of the temporal lobe responsible for indexing memory is called the hippocampus. The hippocampus searches through your memory banks to see if it recognizes any of the data the eyes are sending, it then communicates with the neocortex as to whether there are recognized patterns or not and the neocortex interprets this information.

Cases of amnesia, such as anterograde amnesia in which the brain is unable to transfer new information from short term memory to long term memory, is often caused by damage to the hippocampus. In this case your brain is still able to access memories from before the damage occurred, but can't store new information. So you would be able to retrieve the necessary information for your brain to properly interpret that you're seeing the face of someone you've known for many years, but wouldn't be able to store the information of any new faces you see and thus would be incapable of recognizing them in the future.

Some forms of memory are stored elsewhere however. While the neocortex may govern most motor commands, movements you've practiced many times, such as riding a bike or dribbling a basketball, end up being stored in the cerebellum. The cerebellum is located at the very base of the brain and is one of the oldest and most basic sections of the brain, indeed all vertebrates have a cerebellum. The cerebellum is full of instinctual information; reflexes, basic motor functions, and instinctual emotions like fear.

When you practice a movement enough times, like a professional boxer practices blocks and dodges, that information moves beyond the hippocampus into the cerebellum where it becomes instinct. This is what is colloquially called "muscle memory". A boxer doesn't need to think about dodging a punch, his body does it reflexively.

Complex Emotions Like Love Are Governed by the Neocortex

As stated in the last chapter, the most basic and instinctual emotions such as fear and pleasure are governed by the cerebellum. However more complex emotions that are only experienced by highly intelligent mammalian species, such as love, are processed by the neocortex.

These emotions rely on complicated neuron clusters called spindle cells. Spindle cells are called such because they are very long and web-like. They connect different lobes and areas of the neocortex together like wires in a computer and use the combination of these areas' functions to interpret feelings like affection and empathy.

These complex emotional functions are experienced by humans of course, but they are also experienced by more species than you might think. Some won't be surprising, for instance it's likely fairly obvious that dogs experience feelings of affection and compassion, and our closest relatives like chimpanzees and gorillas experience virtually all the same emotions we do (albeit often on a much more basic level). But such emotions appear in one form or another by most highly intelligent social species such as elephants and dolphins.

Other species however, like snakes, do not experience these emotions since they lack these higher sections of the brain. Species like that are only governed by the most instinctual feelings, like fear or hunger, that are stored in the cerebellum.

One of the inherent qualities of the pattern recognizers that make up the neocortex is the ability to create metaphors and to interpret multiple meanings from symbols, a key element of poetry and art.

Humans have a larger neocortex than any other species, so it should come as no surprise that behaviors that are wholly unique to humans such as creating art or telling stories, are controlled by the neocortex.

Visual art and forms of storytelling like writing make use of symbolism and metaphor. Both of which are only understandable because of our ability to recognize patterns. Metaphor requires you to be capable of drawing connections between what a person is saying, and what deeper meaning they are trying to convey. This is only possible because of the neocortex.

Computer Scientists Are Using Our Understanding of the Brain to Develop Artificial Intelligence

Let's say you want to create a computer capable of learning new information, or new skills, or even a computer capable of self-awareness and thought. How would you go about doing it? Well it would probably make sense to begin by looking at the only machine in existence that is already capable of those things. That "machine", of course, being the brain.

The brain is quite similar to a computer already,it processes and computes information in much the same way. Neurons are in essence biological circuits, and so with this in mind the brain, like a computer, is the sum of a complex system of circuits. The only real difference being the degree of complexity, and the means by which information is transmitted, the brain being electro-chemical while a computer is electro-mechanical.

And so it brings up an interesting question, if we fully understood the brain and we were able to design a computer mimicking the structure of it, using circuits instead of neurons, could we create a computer that can learn? Or perhaps even think?

This is a question computer scientists have been exploring for decades, and have been getting closer to solving. Believe it or not AI, if we define it as a computer capable of learning, has existed for years. Learning computers certainly haven't been on par with the human brain, or even the brains of much smaller and less intelligent mammals. But scientists have been building algorithms and hierarchical models similar to the hierarchical system of the neocortex for years. One such model was created all the way back in the 80s, a statistical model known as the Hierarchical Hidden Markov Model.

The HHMM utilizes a hierarchy of probabilistic models to allow computer programs the ability to recognize patterns. Have you ever wondered how your phone is able to autocomplete words for you? This is due to hierarchical models like HHMM. When you start typing a word your smartphone follows a hierarchy of probability calculations where it searches the word processing program's list of words and calculates which word you're likely attempting to type.

This is the same way other technologies, like translation programs, or voice recognition programs function. And as you've probably noticed this sort of software has been improving.

So while your phone's autocomplete might still make mistakes, it won't be long before it will seem to know what you're trying to type even before you do!

Self-aware Computers Are Closer To Reality Than You Think

While computers being able to recognize what words your saying may be interesting, it likely isn't what you think of when you hear the term Artificial Intelligence.

The term AI more typically conjures up images of sentient robots, possessing consciousness and free-will. While this is still science fiction at the moment, real technology is actually a lot closer to it than you might think.

The author specifically references IBM's computer system known as Watson. Watson is a question-answering computer system designed by IBM's DeepQA project, and is capable of recognizing and answering questions posed in natural language.

Watson is composed of a cluster of 90 servers, with a total of 2,880 different 8-core processors, and a staggering 16 terabytes of RAM. With this hardware Watson is able to process 500 gigabytes per second. That's the equivalent of a million books worth of information per second!

And more importantly than simply containing immense amounts of information Watson is capable of hearing spoken questions, recognizing the meaning of the words, interpret what the question is asking, search it's memory for patterns related to the question, and answer the question. But not just answer the question, answer it correctly!

Watson demonstrated its incredible capabilities by competing on the quiz show Jeopardy, where it competed against humans and beat out champions Ken Jennings and Brad Rutter to win a million dollar prize.

Watson's process of answering questions is shockingly similar to the way our brains accomplish the same task. But it still isn't sentient, nor does it have free will.

But consider that modern PCs have only existed for 40 years, and a computer from the 1980s only had 64 Kilobytes of RAM. For comparison Watson has 250 million times the processing power.

So imagine where computers will be in 10 years, or 20. 30 years from now might not even be imaginable, it might very well be beyond what we see in SciFi today.

10 Years From Now AI Will Be An Indispensable Element of Society

Moore's Law is the name of an observation in computer science that the number of integrated circuits on a transistor tends to double every 2 years. What this means in layman's terms is that computers essential double in processing power every 2 years. So far this has proven accurate since 1975.

And the implications of it are immense. Computer advancement has been exponential, meaning the longer we go on the quicker advancements happen, and the larger the advancements become.

The neocortex stores the equivalent of 20 gigabytes of information, or 20 billion bytes. As stated in the last chapter Watson has 16 trillion bytes of Random Access Memory (RAM). The number of calculations necessary for a neocortex to process per second is estimated to be around 1016, we already have computers capable of this.

As it is the average video game console is capable of storing several times the amount of information storage necessary to create an artificial neocortex. So why don't we have sentient computers yet?

Because we still need to better understand the structure of the brain, and how to design a computer mimicking it. But given that 30 years ago we had no computers capable of any of the necessary functions listed, and today we have computers capable of all of them except one, it should be no surprise if we have genuine examples of basic sentient computers by the beginning of the 2030s.

Final Summary

The mind is the greatest biological mystery, and unlocking it will provide opportunities for technological advancements incomparable to anything in human history. As we begin to understand more and more about how it functions, about how memory is stored and accessed, how information is processed, and how consciousness arises from collections of neurons and electrical signals, we will also begin to understand how to one day create artificial minds ourselves.