CONTENTS

CHAPTER 1

The Civilization Enabler

Cleverness Is Not Enough

Once on a February day some chimpanzees in Belfast Zoo staged a spectacular escape. Their home was a high-walled enclosure opened by the primatologist Jane Goodall some decades earlier. The chimps had lived there uneventfully for many years in the presence of some hawthorn bushes. Shortly before that day, storms had weakened the branches of these bushes. This offered an unexpected opportunity. Soon after the keepers had left for the evening, it took only a few seconds for the daring escape to transpire. A broken branch was leaning against the high wall of the enclosure, reaching about a third of the way up. One of the chimpanzees, using the branch as a ladder, ran up along it, and with a running jump just managed to grab the top of the enclosure wall with one hand and pull itself up and over. Other chimpanzees followed. A sheer wall had been scaled.[1]

Such acts are striking examples of how bright animals can be. In any one case, it is difficult to tease apart the constituent capabilities that the animals are using to accomplish their feats. All the relevant history is rarely available. How did the broken branch come to be leaning against the wall? How much of the skills shown by the chimpanzees was innate rather than learned? Climbing along a branch and jumping to grab hold of something with one arm, as separate capabilities, must each have substantial innate components. Did one of the chimpanzees have previous experience with climbing a dead branch conveniently leaning against a wall? The surprise is not in either act alone, but in the combined event. In this instance, there was no evidence that any one of these chimpanzees had previously had serendipitous experiences that might have given it the idea. There is no reason to believe that any human intervention was involved. If one of the chimpanzees did indeed make a plan to do something never before witnessed by it, or by any of the other chimpanzees, then at least we have to raise our hats.

Crows are also famously clever. The abilities ascribed to them in Aesop's fable *The Crow and the Pitcher* are well recognized. In the fable, a crow drops pebbles into a jug of water, causing the water level to rise. Dropping in enough pebbles allows the crow to quench its thirst when the water level has risen sufficiently.

Researchers in animal behavior have explored this phenomenon systematically through experimentation. With food floating on the surface of water in a narrow-necked container, crows have been observed putting objects in the water. When the water level has risen sufficiently, the crows retrieve the food. When offered several such containers, they choose narrow containers over wide ones, narrow containers needing fewer pebbles to raise the water level. For the objects to add, crows will choose those that sink over those that float, reducing the volume of the objects needed. One does not need a discussion of what the crow really understands or intends to marvel at this behavior.[2]

We may wonder anxiously how we would measure up if presented with tests demanding similar cleverness. But here is the rub: however clever they may be, chimpanzees or crows cannot and will not develop a civilization like humans have. They will not construct a digital computer or travel to the Moon. Even if humans were not here as competitors, neither crows nor chimpanzees would be developing technological civilizations. So, what is the difference? What capabilities do humans possess that these clever apes and crows do not?[3]

I am seeking the capability that not only sets us apart from other living animals but also *provides an explanation* for why we have been able to create the technological civilization that we have. Identifying this civilization-enabling capability is my goal. I shall call this crucial human capability the *Civilization Enabler*.

The Civilization Enabler

It is believed that anatomically modern humans (hereafter referred to as *modern humans*) emerged in Africa more than three hundred thousand years ago, having evolved from within the broader genus *Homo* over the previous two million or so years.[4] The ancestors of modern humans already had significant capabilities that we do not see in other present-day species. By four hundred thousand years ago, these prehuman ancestors were hunting animals with stone-tipped spears, had controlled use of fire, and were cooking (see table 1). Stone tools go back much further, predating the *Homo* genus altogether.

Despite these impressive capabilities already available at the beginning, cultural progress in the three-hundred-thousand-year history of modern humans appears to have been slow initially. Eventually, by around fifty thousand years ago, marked changes had occurred, sometimes called the "cognitive revolution." Art and more complex tools had appeared. Our transformative physical impact on the planet in the form of large settlements or buildings took even longer. These made their mark only a little more than ten thousand years ago. Most of what we would recognize as civilization is even more recent.

Progress toward modern civilization was therefore glacially slow at first. What prompted the more recent advances, and why were they only manifest so late in the history of our species? What caused this delay of three hundred thousand years, or more than ten thousand generations, before our civilization-enabling genes kicked in? Why did they kick in so decisively when they did?

It is possible that this slow progress tracked genetic changes that were happening in humans over these past three hundred thousand years. However, the search for genetic changes that closely correlate with advances in human activity has not yielded anything conclusive. One systematic search sought to find a genetic mutation that spread through the population during this period and that all humans now share.[5] Such an occurrence would be called a *complete selective sweep*. Given the genomes of a set of humans, one can apply statistical methods to estimate for each segment of the genome when the most recent common ancestor of all that genetic material lived. This study estimated that the most recent common ancestors of all the segments searched lived more than three hundred thousand years ago. This suggests that there was no genetic mutation that swept through the whole of the human population in that period. The authors note that their study would have missed mutations occurring in parts of the genome not searched and epigenetic mutations.[6]

It is not that the intervening period was devoid of known genetic events. Humans migrated between continents, and there was some interbreeding with related species, the Neanderthals, the Denisovans, and perhaps others. There have been specific physiological changes in large populations, such as increased tolerance of cow's milk and recent changes in height. Tolerance of cow's milk is due to genetic mutations that spread through only a part of the human population. This is not what we would be looking for as far as the Civilization Enabler that all humans share. Recent changes in height are believed to be *polygenic adaptations*, in which the relative proportions of a set of previously existing gene variants change, without new mutations occurring. If there have been genetic changes in the past three hundred thousand years that have affected our cognitive abilities, they may have been of this polygenic kind where the proportions of different gene variants change in the population, but no novel mutations occur.

It is therefore possible that the critical civilization-enabling mutations appeared with or even much before the emergence of modern humans, but their consequences have played out more slowly. The genetic mutations that increased tolerance for cow's milk gave rise to the selective advantage of a broadened diet that could be enjoyed immediately. In contrast, the advantages of the capability that enables civilization, at least at the most spectacular levels, may have taken thousands of generations to fully manifest themselves.

Now I come to the main proposal of this book. It is a new hypothesis, one that is respectful of these last observations. The hypothesis is that the sought-after Civilization Enabler is what I shall call *educability*. This is a new notion that I shall define more precisely in later chapters. Most basically, educability is concerned with our abilities to *acquire* beliefs, and to *apply* them in specific situations, where the acquisition process allows both the acquisition of beliefs *explicitly described* to us by others and an ability to learn beliefs by *generalizing* from our individual experience.

Table 1. A selection of events in the cultural evolution of our genus *Homo* for which supportive physical evidence has been claimed

2,600,000	Stone tools (Gona, Ethiopia)	Possibly a predecessor of *Homo*
1,700,000	Stone hand axes (Tanzania)	*Homo erectus*?
1,500,000	Control of fire (Kenya)	*Homo erectus*
500,000	Abstract markings: a zigzag engraving on shell (Indonesia)	*Homo erectus*
500,000	Stone-tipped spears (South Africa)	*Homo heidelbergensis*
320,000	Long-distance transport of obsidian for fine blades and points, and ochre for pigments (Kenya)	
200,000	Adhesive: birch tar for hafting stone tools (Campitello, Italy)	Neanderthal
170,000	Widespread use of clothing (Africa)	
142,000	Symbolic ornaments: marine-shell beads (Morocco)	
120,000	Burial of dead (Qafzeh Cave, Israel)	
100,000	Mixing and storing pigments (Blombos Cave, South Africa)	
90,000	Bone harpoons (Semliki River, DR Congo)	
75,000	Jewelry fashions: shifts in styles of threaded shell beads (Blombos Cave, South Africa)	
64,800	Symbolic cave paintings (La Pasiega Cave, Spain)	Neanderthal
50,000	Eyed needle, made from bone (Denisova Cave, Siberia)	Denisovan
48,000	Self-medication with natural pain-killer and antibiotic (El Sidrón, Spain)	Neanderthal
45,500	Representational art, a red-ochre composition of warty pigs (Leang Tedongnge, Sulawesi)	
42,000	Musical instruments: bone and ivory flutes (Swabian Jura, Germany)	
42,000	Fish-hooks, manufactured from broken shell (East Timor)	
40,000	Figurative sculpture, an ivory figurine with lion's head and human torso (Hohlenstein, Germany)	
35,000	Fully human sculpture: a mammoth-ivory "Venus" figurine (Hohle Fels, Germany)	
32,600	Food-plant processing, of dried wild oats with grindstones (Grotta Paglicci, Italy)	
30,000	Woven fabrics, made from dyed fibres of wild flax (Georgia)	
24,000	Poison arrows, with wooden ricin applicator (Lebombo mountains, South Africa)	
23,000	Fisher-hunter-gatherer brush huts (Sea of Galilee, Israel)	
23,000	Domestication: dogs from gray wolves (Siberia)	
20,000	Pottery vessels (Xianrendong Cave, China)	
15,000	String instrument: the musical bow (cave painting at Trois Frères, France)	
14,400	Baking bread: unleavened flatbread from wild einkorn and club-rush tubers (Shubayqa, Jordan)	
11,600	Monumental ritual art (Shigir, Siberia): 5-meter-tall plank carved with human forms and signs	
11,500	Cultivation of wild barley and oats around village settlements	
11,500	Monumental temple (Göbekli Tepe, Anatolia)	
11,000	Continuous settlements (southern Levant)	
8,500	Mining of metal, to heat, hammer, and grind into tools, projectile points (Great Lakes, North America)	
6,000	Earliest board games (Egypt)	
5,500	Domestication of horses (Central Asian steppes)	
5,400	Wheeled wagons (Germany, Slovenia, Near East)	
5,300	Numeral systems: pictograms of economic units (Uruk, Mesopotamia)	
5,200	Full writing (cuneiform in Mesopotamia, hieroglyphics in Egypt)	
4,650	Massive stone monuments (Egypt); contemporaneous pyramids (Peru) and megalith (UK)	

Abstracted from a table constructed by C. Patrick Doncaster, "Timeline of the Human Condition—Milestones in Evolution and History," https://www.southampton.ac.uk/~cpd/history.html, which also gives sources. Used with permission. *Note:* Estimated dates given are years before the present. Future findings may indicate earlier dates.

The sets of beliefs that humans can acquire I call *belief systems*. Some use the phrase "belief system" in the narrower senses, such as religious beliefs. Religious belief systems are examples of what I mean by this phrase. They are certainly among the more ancient belief systems of which we have evidence. The fact that there are many major religions and even more variants of each, I regard as evidence of the human facility for complex sets of abstract beliefs. There is no evidence that nonhuman animals have anything that corresponds. But my use of the term "belief system" here will also encompass other systems of beliefs, including myths, stories, methods of doing things, and the sciences. The commonality among them that is exploited here is the commonality in the ways we mentally acquire and apply them, and not in the degree to which they are, or we believe them to be, true.

Some groups of humans have been genetically isolated from the rest for a long time. It has been argued that the San from southwestern Africa have been largely isolated for more than one hundred thousand years.[7] However, San culture shares with humans across the globe all the basic features that I shall associate with educability. Like other groups, they have origin myths, deities, and stories about their deities, which all involve complex belief systems that individuals acquire from others. They also have sophisticated knowledge that enables them to survive in their desert environment. For example, they traditionally hunt game using arrows poisoned by extracts from the roots of certain not so easy to find desert plants. This hunting technique must have been difficult for an individual to invent, given that it uses knowledge that would be challenging to learn from experience. Once invented, however, it is easy to pass the technique down from generation to generation.

In acquiring their culture, the San therefore needed all the capabilities that are essential to educability, namely, the ability to learn from experience, combined with the abilities to acquire complex theories from others and to apply knowledge gained by either method to new situations. The San's traditional culture fully satisfies the requirements of our definition of educability. If the San have been totally isolated for one hundred thousand years, they could not have acquired mutations occurring elsewhere in the human population during that period, suggesting again that all the genetic requirements for human civilization were already in place for quite some time in the history of humans. As always, some caution is needed as there are uncertainties about how to interpret the evolutionary evidence. Proving total isolation for any group would be difficult, and, indeed, there is evidence that the San were not totally isolated genetically during this period.[8]

Fortunately, to make the case for educability, I shall not need to make *any* assumptions about the details of evolutionary history. The educability hypothesis is consistent with the idea that the development of human cognitive behavior within the past three hundred thousand years largely followed its own slow course, set on its way by our species genetically having the educability capability from the beginning. The hypothesis is minimalist in not assuming that any specific changes occurred in the genome over the course of human history.

One can make several speculative arguments for why educability may be exceptional among human capabilities in the slowness of the pace of its impact after it first emerged. A nongenetic argument is that at the beginning, when little knowledge had been accumulated, the benefit of communicating it to others was small. A genetic argument is that this capability may be the result of polygenic adaptation and dependent on what fraction of a large set of genes favor it. In that case, the fraction of relevant gene variants that favor educability may have been initially low in the population. The benefit of receiving knowledge from others is small if others have low educability and therefore little knowledge to share. For both these reasons, the selective advantage offered would be small initially. It is therefore plausible that following the emergence of the capability for educability, there would be a lengthy period in which the average educability and shared knowledge of the population both increase only slowly. Like snowballs that start rolling, their initial movement is almost imperceptible.

For many years evolutionary biologists and others have been seeking to track down in detail the individual developments that occurred in the emergence of our civilization. They use DNA evidence, the evidence of artifacts from the past, and knowledge of recent human culture. My emphasis here is different. I focus on the *current result* of human evolution and aim to characterize one aspect, which I call educability. If we are to understand our defining evolutionary story, it would seem hard to evade the question of *what* is the defining cognitive capability that we have and that has evolved.

Knowledge Accumulation

The power of educability derives from the fact that the knowledge an individual can acquire if transference from others is possible is incomparably greater than what one could have discovered from one's own experience and efforts alone. Educability offers the individual the enormous power of having knowledge that took multitudes to discover over many generations. This power is not available to species that lack educability.

Language, speech, and the practice of recording information on tablets or paper accompanied the development of civilization and clearly facilitated it. But what made these abilities and technologies useful? My answer is that it is our educability. Technologies for recording knowledge and abilities to communicate have limited power in and of themselves. To individuals who are educable, however, they offer the power of using the knowledge on an unlimited scale.

In the sixteenth century Tycho Brahe observed the sky at night and collected systematic data on the positions of the planets over several decades. For this, Brahe used an observatory that is thought to have consumed for a time a sizable portion of the government expenditure of his country, Denmark. After Brahe's death, Johannes Kepler, his former assistant, used this data to deduce that the planetary orbits were elliptical. To start, Kepler needed only to read Brahe's tabulated data, without having to repeat the observations or expenditures. Kepler also had to do much more. Essential to his discovery were mathematical notions that he learned from others. He needed the notion of an ellipse, which was known almost two thousand years earlier in Greece. For his third and last law, Kepler needed to do complex calculations. For these he used the method of logarithms, about which John Napier had recently published in Scotland.

Kepler was able to exploit all this knowledge, both the data and the mathematical principles, which he had no chance of observing and deriving from scratch by himself. Through a combination of formal education and self-education, he put himself in the position of being able to use what others before him had obtained with lifetimes of effort. The understanding that planetary orbits were elliptical remains one of the crowning achievements of humanity, one that has had decisive impact on the subsequent development of science. Kepler's facility in absorbing and applying previously obtained knowledge made it possible.

To the question of what genetic changes happened between the first emergence of modern humans more than three hundred thousand years ago and the visible large-scale products of civilization around ten thousand years ago, our answer is that nothing dramatic *needed* to have happened. It may have just taken a long time for educability, which the earliest humans could have possessed, to manifest its full power. For hunter-gatherers without much physical security or much spare time to teach and learn, the opportunities for flaunting the power of educability might have been slight.

By early seventeenth-century Europe, the environment was quite different. An individual like Kepler could exploit the information he had learned from others with enormous consequence. In the present day, information is disseminated around the world at an ever more feverish rate, and scientific discoveries are being made at a correspondingly ever more rapid pace. This pace is maintained with the help of universal education and digital technology. To exploit educability to the full, both of these may be essential.

It had taken a long time before the cumulative value of the knowledge gained across the planet was great enough to become self-sustaining as a process and spawn the technological civilization we have today. Through a gargantuan multigenerational effort, the Civilization Enabler has given us a good understanding of the physical world, and a capability to transform it.

A Computational Approach

What exactly is the nature of this Civilization Enabler? Its fruits are easy enough to see: culture, the arts, knowledge, science. Here I am interested in going further. I want to understand the human *capability* that gives rise to these fruits.

The notion of educability as I define it did not arise from and cannot be defined in terms of physics, chemistry, biology, or the social sciences. Nor can it be justified by the methodologies used in these sciences. This is hardly surprising since its subject matter is information processing, which is not the focus of these sciences.

While the twentieth century saw unparalleled developments in the classical sciences, equally important and particularly in the work of Alan Turing in the 1930s, it saw the birth of the science of information processing. By that time, it was commonplace not to marvel that physical concepts that are not visible, such as energy or electric charge, could have useful meaning. The fact that the same held for notions of information processing and computation, terms that I shall use synonymously, was startling news.

The import of this news was well understood by the early pioneers of computing, namely, Turing himself and John von Neumann. They sought immediately to use computation to study biological phenomena, such as the brain, cognition, and genetics. Each of these phenomena involves the transformation of information. Focusing on the information processing rather than the physical realization became a viable and necessary approach toward understanding these once a scientific approach to information processing had come into view.

I, and many others, consider it self-evident that if we are to understand how the brain works, we will need to understand it in terms of information processing. Some skeptics have suggested that information processing is only a metaphor. They point out that there is historical precedent for comparing the brain with the most complicated and prized machines of earlier eras, such as the camera. Perhaps once again we are just comparing the brain with the most complex machine that we happen to have, the computer. Perhaps computation is just one of many possible metaphors for the brain and there is nothing much to choose among them.

I would say that this view is mistaken. Physics is more than a metaphor for the physical world. It relates to falling apples and the motion of the planets more than as a mere metaphorical narrative. It seeks to *explain* falling apples and the motion of the planets. Whenever a gap appears between the description offered by physics and the observed behavior of falling apples or moving planets, every effort is made to update the physics. Physics is more than a metaphor because its *ambition* is more than that. It is always willing to change and improve so as to be more useful in explaining reality. It has usually succeeded.

The case is the same for computer science as it relates to the world of information processing. Long before humans, life on Earth was processing information. The copying and mutations of the DNA that occurred between successive generations of bacteria, animals, and plants is information processing. With the evolution of nervous systems, information processing came to be carried out in brains on a yet different and even more massive scale.

Computer science is more than just a metaphor for the world of information processing. Its ambition is all-encompassing in aiming to explain *every* kind of information processing that is possible, whether in biology, silicon, or some other realization. As long as gaps are found between the description that is offered by computer science and a real-world information processing phenomenon, efforts will be made to update computer science in order to resolve the gap. It is willing and able to change and improve.

Why Computation?

In a nutshell, computation is used here to provide *concrete* descriptions of processes, such as of learning and education. There are two senses in which I will need this concreteness. The first sense is that of *precisely specifying the outcome* of a process—defining, for example, exactly what the outcome of a learning process needs to be if we are to declare that learning has been successfully achieved. The second sense is that there is a precise *step-by-step description*, possibly as a computer program, of how this outcome *can* be achieved using reasonable resources—one cannot ascribe to the brain or other nature phenomena capabilities that one does not know how to realize concretely by any means in this universe.

My intention is to define a notion that is good for more than a coffee table discussion. I am claiming that educability is a useful scientific concept. To make that case, I will need to specify the nature of educability quite precisely.

I regard both educability and education as phenomena of computation. In the course of education, information is presented, whether as the description of a specific situation or as an explicit description of a general belief. The result of the presentation will be to make a difference in the student's subsequent behavior as compared with the past. The change in behavior will be attributable to a change of the *state* of the student, realized as some physical change in the brain that persists for some time.[9]

Computation is about changes of state that can be realized by step-by-step processes. Physical systems also change state—if you boil water there is a change of state. In a computer or a brain, there is extreme flexibility in the realizable state changes and in their possible effects. In both computers and brains, there are billions of parts that at any instant have some state. Each state arises as the cumulative effect of past experiences. Each state can influence future behaviors. In computers certainly, the state of a single one of the billions of elements, namely, whether it is a 0 or a 1, can have a decisive effect on later behavior. The way in which the effect depends on the state can be arbitrarily complex.

Computation encompasses all ways experience can cause changes of state and, in turn, state can influence later behavior. It is this expressiveness and flexibility of computation that is so useful for describing cognitive capabilities.

The idea that many of the unresolved secrets of biological phenomena, such as cognition, lie in the world of computation needs some discussion. I will try to break this down further in chapter 5. Progress in unraveling these secrets has been slow, and the difficulty of treading this path has often been underestimated. In a talk at a conference banquet in 1957, the artificial intelligence pioneer Herbert Simon made four bold predictions.[10] The first was that “within ten years a digital computer will be the world’s chess champion, unless the rules bar it from competition.” As it happened, it took forty years for computer chess to approach the needed level, and the International Chess Federation rules have been barring its participation.

Simon’s fourth prediction is the most relevant to the discussion here. He predicted that “within ten years most theories in psychology will take the form of computer programs, or of qualitative statements about the characteristics of computer programs.” Well, even an extra half century has not been enough for this to happen. I believe the reason is that the necessary connection between psychology and computation is more subtle than Simon’s statement suggested. This book is about some of the nuanced connections between the two. Explaining these connections, and why they work, is one of the tasks I take on here, and one I will keep coming back to.

INDEX

Some of the terms listed are used in the text in a technical sense. For those, the page numbers in italics indicate where the terms are defined.