DIGITAL GENESIS

The Future of Computing, Robots and AI

The complete book is available from Amazon as a 258 page paperback or as a Kindle e-book.

ISBN: 978-1976098062

Paperback RRP: $16.99 US / £10.49 UK / €14.99 EU

For more information *see*:

ExplainingComputers.com/genesis.html

With Best Wishes,

Christopher Barnatt.

ExplainingComputers.com

PREFACE
CYBER BUSINESS

In 1984 science fiction author William Gibson published his first novel. Entitled *Neuromancer*, this envisaged a future in which people worked and spent leisure time in the virtual reality of 'cyberspace'. Gibson described cyberspace as the electronic 'consensual hallucination experienced daily by billions'. His hero, Case, even had a 'cranial jack' that allowed him to directly connect his brain to the global digital matrix.

Neuromancer inspired a lot of people to start thinking about the online realm. One of those individuals was myself, and in the Summer of 1994 I finished writing a book called *The Cyber Business* that predicted a future awash with virtual communities and e-commerce. Manuscript complete, I began ringing up publishers, and quickly discovered two things. Firstly, that few editors understood what I meant by 'doing business in cyberspace'. And secondly, that it would have been far easier to sell a book called *The Cider Business* about the apple-based alcoholic drink.

Fortunately, toward the end of 1994 my luck improved. In fact, following twelve rejections, two publishers actually ended up competing for the rights to my work. In the Spring of 1995, *Cyber Business: Mindsets for a Wired Age* was subsequently released and rode the very start of the Dot Com wave. It even got some coverage in the *Financial Times*.

Eleven books and twenty-two years later, *Digital Genesis* marks my return to writing in depth about the future of computing and the online realm. So far, most of what I predicted in *Cyber Business* has come true, with goods and services regularly purchased online, a wireless computing device in over two billion hands, and social lives orchestrated by what I then termed 'personal virtual networks'. Without doubt, over the past two decades computers and the Internet have dramatically altered many aspects of our lives. Even so, as we shall see across this book, the most radical digital transformations are yet to come.

SYNTHETIC CITIZENS

So just what are those computational watersheds that lie ahead? Well, for a start, artificial intelligence (AI) is widely expected to be the Next Big Thing. This means that, in as little as five years, a lot of dumb computing interfaces are going to be replaced with chatbots, digital assistants and other smart AIs. These will allow us to communicate with digital technologies by speaking or typing in English or any other human language, with the system able to speak or write back.

Already many people are familiar with first-generation AI interfaces such as Amazon's Alexa, Apple's Siri, Microsoft's Cortana and Google Assistant. Today these are often dismissed as a gimmick. But they really should not be ignored, as by the early 2020s the next generation of AI interfaces is going to be very useful indeed.

As chatbots and digital assistants develop, many organizations will adopt them as their main customer interface. Indeed, according to a March 2017 survey from Accenture, by 2020 most bankers expect AIs to be 'their primary method for interacting with customers'. This will, I expect, prove to be a pretty accurate prediction – and not only in banking. Just as, 20 years ago, the world-wide web was poised to

transform organizational interfaces, so today AI is set to do exactly the same thing.

In addition to becoming a common means of accessing business services, over the next decade AIs will start to take on many professional jobs. Until this point in time, accountants, doctors, lawyers and most managers have held positions that could not be automated. But in the 2020s this will no longer be the case.

Future AIs will also take control of many machines that currently require a human operator. By 2025, an increasing number of products are therefore destined to be produced in 'dark factories' manned entirely by robots. Within 15 years, the majority of new automobiles, trucks and tractors are also likely to be equipped with AIs that will allow them to drive themselves. By the early 2030s, I even anticipate that tens of millions of humanoid robots will be labouring at tasks that are no longer considered cost-effective or safe enough to be undertaken by a person.

The implications of the next wave of automation on human employment could be very stark indeed. This said, the data analysis skills of future AIs should help us to optimize our use of increasingly scarce resources. Equally beneficial will be the improvements in engineering and medicine that will arise from the cognitive abilities of future AIs. In fact, to solve challenges such as climate change and mass antibiotic resistance, we may have no choice but to turn to the promise and power of AI.

QUANTUM & ORGANIC TECHNOLOGY

In addition to the rise of AI and robots, another digital transformation will be the evolution of next generation hardware. Today, computers are based on inorganic microprocessors that contain millions or billions of miniature transistors. These have served us well for many years, and will remain in use for the foreseeable future. Even so, over the coming decades we should expect traditional computers to be complemented by hardware based on quantum and even organic components.

Quantum computers process data not in binary bits, but using quantum mechanical ‘qubits’. As we will see in chapter 5, IBM, Google, Microsoft and others are now investing in quantum computing research and development, with commercial systems a realistic possibility in the late 2020s.

Further into the future, developments in organic technology may allow the brains and bodies of some future computers and robots to be fashioned from living ‘wetware’, rather than semiconductors, metals and plastics. Perhaps by the 2040s, autonomous vehicles may also be organically grown. This is, after all, how autonomous vehicles were manufactured in the past (when we used to call them horses).

Already practitioners of the new science of synthetic biology (SynBio) are learning how to digitally manufacture things from modular DNA components. For at least 200,000 years, the most sophisticated computers on the planet have been the 86-billion-neuron blobs of wetware that we call ‘human brains’. As synthetic biology advances, the odds are therefore in favour of at least some of tomorrow’s computers being organic.

HUMANITY 2.0

While AI, robots and next-generation hardware offer incredible possibilities, the ultimate digital transformation will be the fusion of computers with the human body. In April 2017, Tesla and SpaceX CEO Elon Musk announced a new company called Neuralink that is ‘developing ultra high bandwidth brain-machine interfaces’. While this venture is in its very early stages, Neuralink is far from alone in its ambition to directly link people with artificial digital technology. By the 2040s, we may therefore have computers installed inside our heads.

In addition to allowing people to enter cyberspace using direct brain interfaces, future ‘cyborg fusion’ may one day allow the human body to be digitally upgraded. Already extraordinary progress in synthetic biology, genetic medicine and nanotechnology point to a future in which our own physical forms – and all other organic matter – will be able to be digitally programmed. Such programming will lie far beyond the limits of any human individual or team thereof. But it is going to be child’s play for some future AIs.

Long before AIs learn how to reprogram human biology, our race looks certain to have irrevocably melded with digital machines. Even today, millions of people find it traumatic to function in the absence of a computer. Not least the smartphone has rapidly evolved into a critical prosthesis, while both our economy and our culture have become reliant on the Internet. To suggest that we will sooner or later physically merge with artificial computing systems is therefore hardly a radical proposition. This book’s examination of the future of computing, robots and AI is therefore also an exploration of the future of ourselves.

BEYOND THE INTERNET

Around 375 million years ago our ancestors dragged themselves out of the oceans. Critical to their success in this venture was the innovation of a new, organic technology called ‘lungs’ that remains vital to our survival to this day. This said, few people would dispute that the ‘Breathing Air Revolution’ is now well and truly over. A very long time ago we not only developed lungs, but learnt how to live in a new way while using them, and so moved on to greater innovations.

To embrace the revolutions of the future we need to accept the ending of the revolutions of the past. As I mentioned a few pages back, I got in early on the whole cyberspace thing. But I firmly believe that the Internet Revolution is now over.

Granted, it is over in the same way that the Breathing Air Revolution is over, with both Internet usage and the inhaling of air remaining vital for our survival in the modern world. But just as the development of lungs is no longer a radical innovation in progress, so the continued evolution of the Internet is likely to prove relatively insignificant. This book will therefore not concern itself with the future of social media and the world-wide web.

Instead, and at a far more fundamental level, the following chapters explore the scope and scale of the most radical digital revolutions of the next few decades. Right now, we stand on the brink of a new age of digital genesis that will shake our civilization to its core. And, for better or for worse, it is going to be truly extraordinary.

Christopher Barnatt,
September 2017.

1

THE COMPUTING MACHINE

In May 2017 I visited Bletchley Park in the United Kingdom. Here, during WWII, military codebreakers made pioneering developments in computing. Today, Bletchley Park is maintained as a site of international historic importance. It is also home to the National Museum of Computing, where they have painstakingly rebuilt and restored several very early computers.

Pride of place in the National Museum is taken by a recreation of Colossus. As pictured in figure 1.1, this was one of the world's first electronic computers, and was used to help decipher encrypted radio communications. The original machine was designed and built by a man called Tommy Flowers, and attacked its first scrambled message in February 1944. A Mark 2 Colossus was then developed, and by the end of the war 10 Colossus computers were being used by the codebreakers at Bletchley Park.

As its name implies, Colossus is a large and imposing piece of technology. To my surprise, it is also quite a noisy contraption. In part this is because it has no electronic memory, and hence has to be constantly fed data from a continuous loop of rapidly-circulating paper tape. Each tape contains up to 25,000 5-bit characters, and travels around a set of wheels and pulleys at up to 27 miles per hour.

In addition to the whirr of its paper tape, Colossus emits a deep, repetitive 'clunk' as its relays cycle. The aggregate sound is almost hypnotic. Indeed, when you stand next to the machine, you can almost feel its mechanical pulse resonating out from the past, and deep into our computing future.

LEARNING FROM PIONEERS

At the end of the *Preface* I stated that the Internet Revolution is over, and that to embrace the future we need to accept its passing. While I absolutely believe this to be the case, it would be strange in a book called *Digital Genesis* not to delve into computing history. As Colossus highlights, WWII codebreakers used early electronic technologies to achieve amazing things that were at the time beyond most people's comprehension. In the decades ahead, we are due for a similar round of digital innovation that many will find hard to believe. An excursion into the history of computing may therefore provide a useful reminder of just how frequently the apparently impossible has already morphed into digital reality.

Many historians and textbooks divide the history of computing into four or five generations, with each of these associated with a specific technology. So, for example, 'first generation' computing is generally accepted to have taken place between the early 1940s and the mid 1950s, and to have involved electronic computers that were based on vacuum tubes.

Computing's 'second generation' then spanned from the mid 1950s to the mid 1960s, and involved systems whose circuits were built from individual transistors. The development of integrated circuit (IC) technology next allowed multiple transistors to be manufactured on a single chip, so giving rise to IC-based 'third generation' computers from the mid 1960s to the early 1970s.

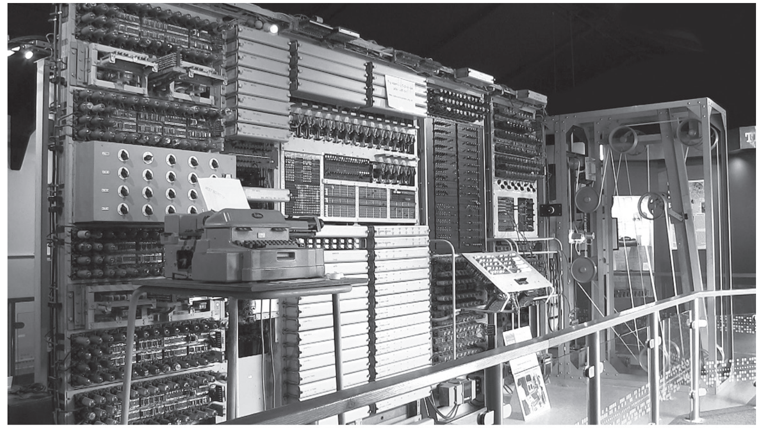

Figure 1.1: The Reconstructed Colossus at the National Museum of Computing.

The early 1970s heralded a 'fourth generation' of computing based around microprocessors, and which arguably continues to this day. However, in 1982, the Japanese Ministry of International Trade and Industry proposed a 'fifth generation' focused on AI and natural language systems. Quite when and if this began continues to be debated. Which all means that today, we are either at the end of the fourth generation of computing, or at the start of the fifth.

THE FIVE AGES OF COMPUTING

Until the early-to-mid 1990s, the four/five generations model of computing history served its purpose as a reasonable analytical tool. Yet since that time, it has become evident that any categorization of the evolution and impact of computing needs to be based on more than a measure of its underlying hardware. Within this book, I have therefore chosen to distinguish five 'ages of computing'. As illustrated in figure 1.2, these inevitably and rightly overlap. It is also important to understand that each of the five ages is distinguished by its focus on *additional* digital capabilities, rather than the replacement of what has gone before.

At the top of the figure we have the Early Computing Age. This began with the invention of the earliest mechanical calculators, was catalyzed by the construction of the first electronic computers in the early 1940s, and began to draw to a close in the mid 1970s. During all of this period computing was a minority, elite activity, with no personal computers in existence, and only large organizations able to invest in non-human computational resources.

From the mid 1970s onwards we had the arrival of the Personal Computing Age. During the following couple of decades most forms of media started to go digital, and computing became a mainstream human activity. The use of mainframes and other large computers did advance considerably during this period. Nevertheless, it was the rise of personal computing that was most significant in the 1980s and 1990s, with over a billion PCs entering commercial and domestic use.

Just before the turn of the millennium, the arrival of the public and business Internet heralded the dawn of the Network Computing Age. This current slice of computing evolution continues to be characterised by mass digital interconnection and widespread interpersonal communication. To this end it has seen a growing addiction to mobile computer hardware, including laptops, tablets and smartphones.

Looking ahead, I predict that the period between about 2020 to 2040 will be the Cognitive Computing Age, with the next two decades characterised by the mass application of AI. In particular, during the Cognitive Computing Age we will witness a transition away from dumb computer systems, and toward smart 'attentive computers' that will anticipate human needs. Almost certainly, many of these attentive computers will be robots that will work and play with us in the physical world.

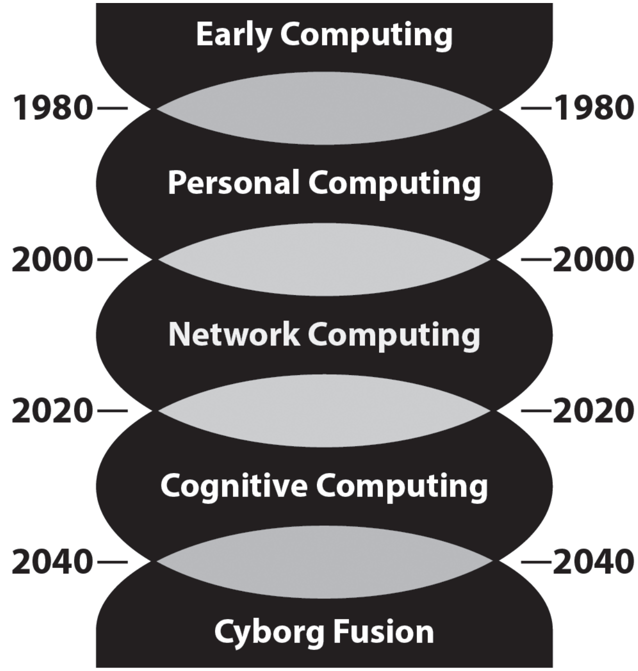

Figure 1.2: The Five Ages of Computing

Finally, speculating even further into the future, I predict that around 2040 we will enter the era of Cyborg Fusion. This means that, a little over 20 years from now, it will become increasingly common for human beings to be directly augmented with digital technology. Beyond 2040, it will therefore become difficult to separate some computers from their users.

To provide a foundation for the rest of this book, the remainder of this chapter explores each of the Five Ages of Computing. Along the way it also charts the origins of key players in the computing industry, including IBM, Apple, Microsoft, Amazon, Google and Facebook. All of these companies are likely to play an important role in shaping our digital future, and I will report on their endeavours throughout this book. It is therefore wise for us to gain some understanding of their early beginnings.

THE EARLY COMPUTING AGE

A 'computer' is technically any device capable of undertaking computation. For thousands of years human beings have sought to make such contraptions, with the abacus in use by the Mesopotamians as early as 3000 BC. However, it was not until 1642 that the first mechanical adding machine was constructed by French scientist and tax collector Blasé Pascal.

The next significant development came in 1823, when Cambridge professor Charles Babbage designed his famous 'Difference Engine'. This was intended to assist in the calculation of astronomical, ballistics and engineering tables to an accuracy of up to twenty decimal places. But unfortunately, due to problems with funding and engineering accuracy, Babbage never managed to construct operational hardware. In fact it took until 1990 for Babbage's vision to be vindicated, when a team of researchers at the London Science Museum used his blueprints to build a working Difference Engine. The completed machine has over 4,000 iron, bronze and steel parts and weighs three tonnes.

By 1837, Babbage had designed an even more ambitious mechanical computing device called the 'Analytic Engine'. This used punched cards for data and program storage – a technique which had been pioneered in 1804 by Joseph-Marie Jacquard in the automation of mechanical weaving looms.

The use of interchangeable punched cards to store instructions for the Analytic Engine resulted in the development of computer programming. In fact Lady Ada Augusta – who worked on the Analytic Engine with Babbage – developed several coding techniques that remain in use today.

FULL BOOK CONTENTS

ExplainingComputers.com