

HOW THE
INFOSPHERE
IS RESHAPING
HUMAN REALITY

LUCIANO FLORIDI

OXFORD
UNIVERSITY PRESS

1

TIME

Hyperhistory

The three ages of human development

More people are alive today than ever before in human history. And more of us live longer today than ever before. Life expectancy is increasing (Figure 1; see also Figure 19) and poverty is decreasing (Figure 2), even if the degree of global inequality is still scandalous. As a result, disability is becoming the biggest health-related issue for humanity.

To a large measure, the lines representing the trends on Figure 1 and Figure 2 have been drawn by our technologies, at least insofar as we develop and use them intelligently, peacefully, and sustainably.

Sometimes we forget how much we owe to flints and wheels, to sparks and ploughs, to engines and computers. We are reminded of our deep technological debt when we divide human life into *prehistory* and *history*. Such a significant threshold is there to acknowledge that it was the invention and development of ICTs (information and communication technologies) that made all the difference between who we were, who we are, and, as I shall argue in this book, who we could be and become. It is only when systems to record events and hence accumulate and transmit information for future consumption became available that lessons learnt by past generations began to evolve exponentially, in a soft or Lamarckian[1] way, and so humanity entered into history.

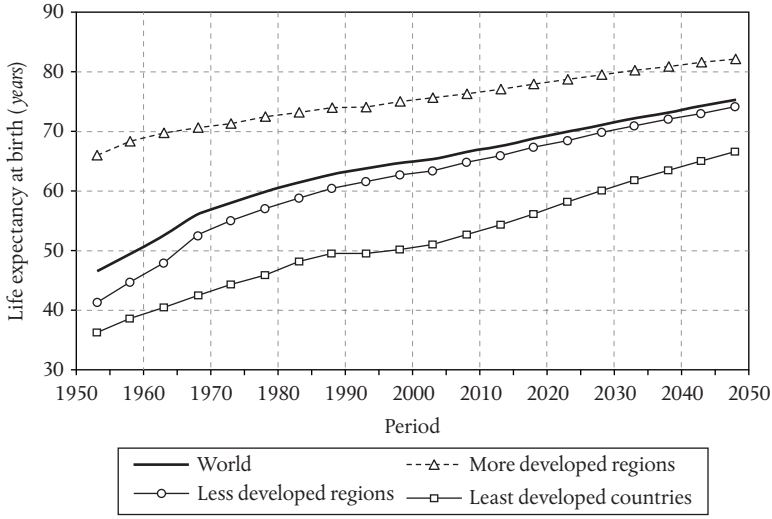


Fig. 1. Life Expectancy at Birth for the World and Major Development Group, 1950–2050.

Source: Population Division of the Department of Economic and Social Affairs of the United Nations Secretariat (2005). *World Population Prospects: The 2004 Revision Highlights*. New York: United Nations.

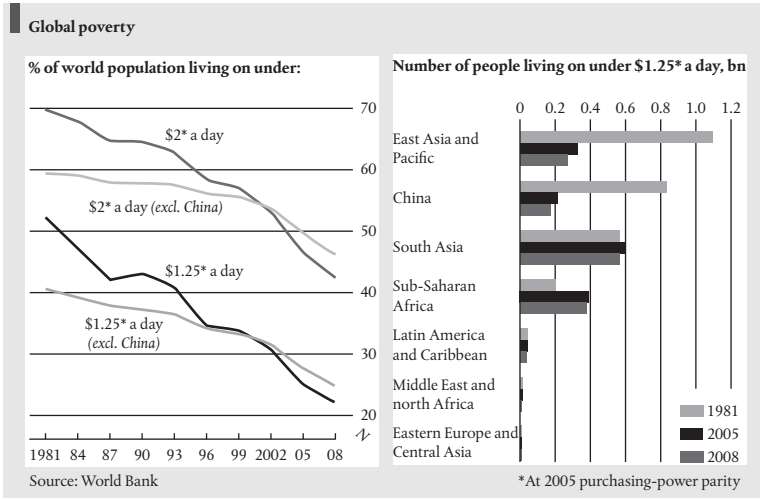


Fig. 2. Poverty in the World Defined as the Number and Share of People Living Below $1.25 a Day (at 2005 prices) in 2005–2008.

Source: World Bank. © *The Economist* Newspaper Limited, London (29 February 2012).

History is therefore synonymous with the information age. Such a line of reasoning may suggest that humanity has been living in various kinds of information societies at least since the Bronze Age, the era that marks the invention of writing in Mesopotamia and other regions of the world (4th millennium BC). Indeed, in the 3rd millennium BC, Ur, the city state in Sumer (Iraq), represented the most developed and centralized bureaucratic state in the world. So much so that, before the Gulf War (1991) and the Iraq War (2003–11), we still had a library of hundreds of thousands of clay tablets. They contain neither love letters nor holiday stories, but mainly inventories, business transactions, and administration documents. And yet, Ur is not what we typically have in mind when we speak of an information society. There may be many explanations, but one seems more convincing than any other: only very recently has human progress and welfare begun to be not just *related to*, but *mostly dependent on*, the successful and efficient management of the life cycle of information. I shall say more about such a cycle in the rest of this chapter, but, first, let us see why such a dependency has meant that we recently entered into *hyperhistory* (Figure 3).

Prehistory and history work like adverbs: they tell us *how* people live, not *when* or *where* they live. From this perspective, human societies currently stretch across three ages, as ways of living. According

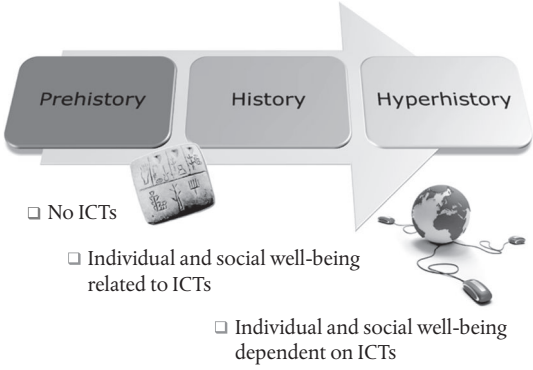


Fig. 3. From Prehistory to Hyperhistory.

to reports about an unspecified number of uncontacted tribes in the Amazonian region,[2] at the beginning of the second millennium there were still some societies that may be living prehistorically, without recorded documents. If, or rather when, one day such tribes disappear, the end of the first chapter of our evolutionary book will have been written.

The greatest majority of people today still live historically, in societies that rely on ICTs to record, transmit, and use data of all kinds. In such historical societies, ICTs have not yet overtaken other technologies, especially energy-related ones, in terms of their vital importance. Then, there are some people around the world who are already living hyperhistorically, in societies and environments where ICTs and their data-processing capabilities are not just important but essential conditions for the maintenance and any further development of societal welfare, personal well-being, and overall flourishing. For example, all members of the G7 group—namely Canada, France, Germany, Italy, Japan, the United Kingdom, and the United States of America—qualify as hyperhistorical societies because, in each country, at least 70 per cent of the Gross Domestic Product (GDP, the value of goods and services produced in a country) depends on intangible goods, which are information-related, rather than on material goods, which are the physical output of agricultural or manufacturing processes. Their economies heavily rely on information-based assets (knowledge-based economy), information-intensive services (especially business and property services, communications, finance, insurance, and entertainment), and information-oriented public sectors (especially education, public administration, and health care).

The nature of conflicts provides a sad test for the reliability of this tripartite interpretation of human evolution. Only a society that lives hyperhistorically can be threatened informationally, by a cyber attack. Only those who live by the digit may die by the digit, as we shall see in Chapter 8.

Let us return to Ur. The reason why we do not consider Ur an information society is because it was historical but not yet

hyperhistorical. It depended more on agricultural technologies, for example, than on clay tablets. Sumerian ICTs provided the recording and transmitting infrastructure that made the escalation of other technologies possible, with the direct consequence of furthering our dependence on more and more layers of technologies. However, the recording and transmitting facilities of ICTs evolved into processing capabilities only millennia later, in the few centuries between Johann Gutenberg (*c.*1400–68) and Alan Turing (1912–54). It is only the present generation that is experiencing the radical transformations, brought about by ICTs, which are drawing the new threshold between history and hyperhistory.

The length of time that the evolution of ICTs has taken to bring about hyperhistorical information societies should not be surprising. The life cycle of information (see Figure 4) typically includes the

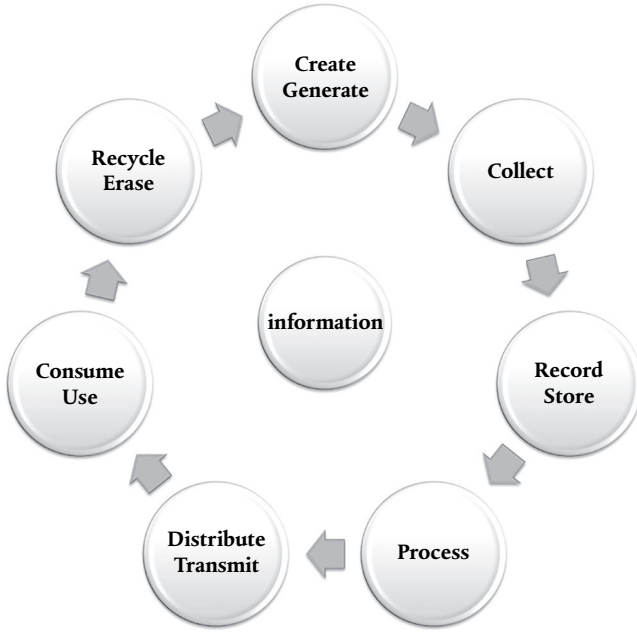


Fig. 4. A Typical Life Cycle for Information.

following phases: *occurrence* (discovering, designing, authoring, etc.), *recording, transmission* (networking, distributing, accessing, retrieving, etc.), *processing* (collecting, validating, merging, modifying, organizing, indexing, classifying, filtering, updating, sorting, storing, etc.), and *usage* (monitoring, modelling, analysing, explaining, planning, forecasting, decision-making, instructing, educating, learning, playing, etc.). Now, imagine Figure 4 to be like a clock and a historian writing in the future, say in a million years. She may consider it normal, and perhaps even elegantly symmetrical, that it took roughly 6,000 years for the agricultural revolution to produce its full effect, from its beginning in the Neolithic (10th millennium BC), until the Bronze Age, and then another 6,000 years for the information revolution to bear its main fruit, from the Bronze Age until the end of the 2nd millennium AD. She may find it useful to visualize human evolution as a three-stage rocket: in prehistory, there are no ICTs; in history, there are ICTs, they record and transmit information, but human societies depend mainly on other kinds of technologies concerning primary resources and energy; and in hyperhistory, there are ICTs, they record, transmit, and, above all, process information, increasingly autonomously, and human societies become vitally dependent on them and on information as a fundamental resource in order to flourish. Around the beginning of the 3rd millennium, our future historian may conclude, innovation, welfare, and added value moved from being ICT-related to being ICT-dependent. She might suppose that such a shift required unprecedented levels of processing power and huge quantities of data. And she might suspect that memory and connectivity must have represented bottlenecks of some kind. She would be right on both accounts, as we shall see in the rest of this chapter.

Instructions

Consider the two diagrams in Figure 5 and Figure 6. Figure 5 is famous, almost iconic. It is known as Moore's Law and suggests that, over the

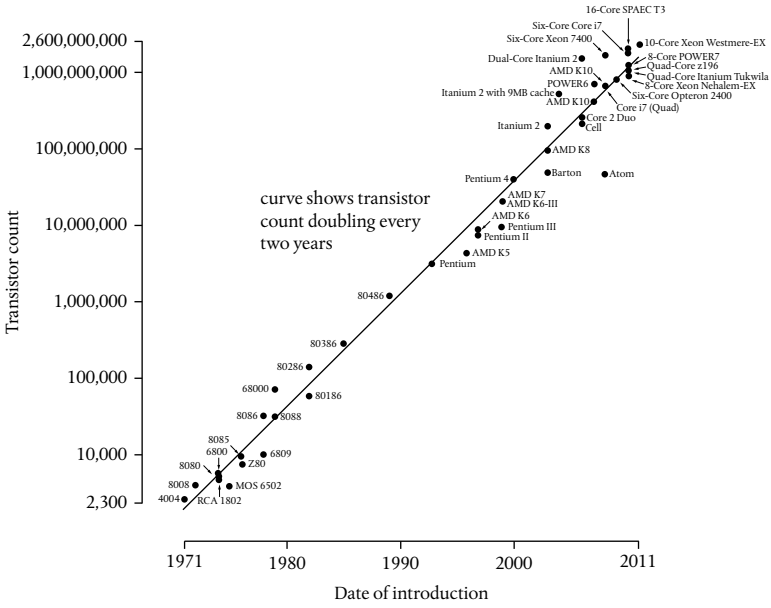


Fig. 5. Moore's Law.
Source: Wikipedia.

period of development of digital computers, the number of transistors on integrated circuits doubles approximately every two years.

Figure 6 is less famous but equally astonishing. It tells you a similar story, but in terms of decreasing cost of computational power. In 2010, an iPad2 had enough computing power to process 1,600 millions of instructions per second (MIPS). By making the price of such a processing power equal to $100, the graph shows what it would have cost to buy the computing power of an iPad2 in the past decades. Note that the vertical scale is logarithmic, so it descends by powers of ten as the price of computing power decreases dramatically. All this means that, in the fifties, the 1,600 MIPS you hold in your hands—or rather held, in 2010, because three years later the iPad4 already run at 17,056 MIPS—would have cost you $100 trillion. This is a number that only bankers and generals understand. So, for a quick comparison, consider

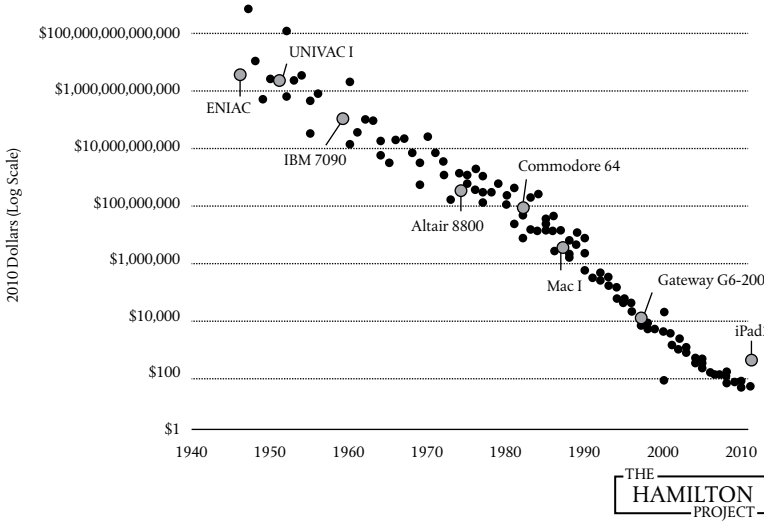


Fig. 6. The Cost of Computing Power Equal to an iPad2.
Source: The Hamilton Project at the Brookings Institution.

Qatar's GDP. In 2010, it was ranked 57th out of 190 countries in the world and its GDP would have been insufficient to buy the equivalent of an iPad2 in the fifties, for it was a mere $98 trillion.

Whether you find Figure 5 or Figure 6 more compelling, the conclusion is the same: increasingly more power is available at decreasing costs, to ever more people, in quantities and at a pace that are mind-boggling. The limits of computing power seem to be mainly physical. They concern how well our ICTs can dissipate heat and recover from unavoidable hardware faults while becoming increasingly small. This is the rocket that has made humanity travel from history to hyperhistory, to use a previous analogy. It also explains why ICTs are still disruptive technologies that have not sedimented: new generations keep teaching the old ones how to use them, although they still learn from previous generations how to drive or use a microwave.

At this point, an obvious question is where all this computational power goes. It is not that we are regularly putting people on the Moon

with our smartphones or tablets. The answer is: interactions, both machine-to-machine and human–computer ones, also known as HCI.

In machine-to-machine interactions, an ICT system, such as a meter or sensor, monitors and records an event, such as the state of the road surface, and communicates the resulting data through a network to an application, which processes the data and acts on them, for example by automatically adapting the braking process of a car, if necessary. You might have heard that there is more computational power in an average new car today than was available to NASA to send astronauts to the Moon (Apollo mission, 1969). It is true. There are more than 50 ICT systems in an ordinary car, controlling anything from satellite navigation to hi-fi display, from ABS (anti-locking brakes) to electric locks, from entertainment systems to all the sensors embedded in the engine. It is a growing market in the automobile industry, as Figure 7 illustrates. According to Intel, the connected car is already the third fastest growing technological device after phones and tablets. It is only a matter of (short) time before all new cars will be connected to the Internet and, for example, find a convenient car park space, sense other vehicles, or spot cheaper petrol prices along the journey. And of course electric vehicles will require more and more 'computation': by 2015, they will contain about twice as many semiconductors than conventional cars. Mechanics are becoming computer engineers.

In human–computer interactions (HCI), ICTs are used to create, facilitate, and improve communications between human users and computational systems. When talking about ICTs, it is easy to forget that computers do not compute and telephones do not phone, to put it slightly paradoxically. What computers, smartphones, tablets, and all the other incarnations of ICTs do is to handle data. We rely on their capacities to manage huge quantities of MIPS much less to add numbers or call our friends than to update our Facebook status, order and read the latest e-books online, bill someone, buy an airline ticket, scan an electronic boarding pass, watch a movie, monitor the inside of a shop, drive to a place, or, indeed, almost anything else. This is why HCI is so important. Indeed, since the mid-1990s, HCI does not even

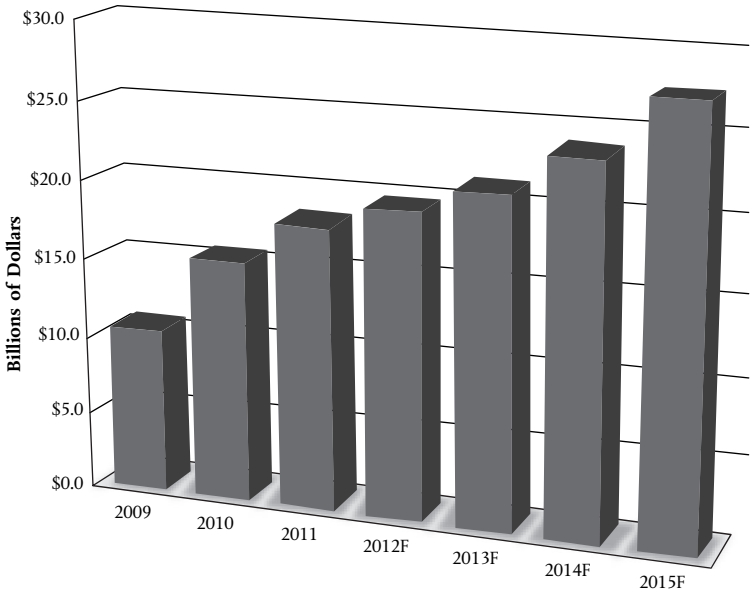


Fig. 7. Value in US Dollars of the Average Semiconductor Content in Automobiles. F=forecast.

Data source: IC Insights, 2012.

have to involve screens or keyboards. It may be a matter of a neuroprosthetic device implanted in the brain. Of course, in all human–computer interactions, the better the process, the computationally greedier the ICT in question is likely to be. It takes a lot of MIPS to make things easy. This is the reason why new operating systems can hardly run on old computers.

We know that what our eyes can see in the world—the visible spectrum of the rainbow—is but a very small portion of the electromagnetic spectrum, which includes gamma rays, X-rays, ultraviolet, infrared, microwaves, and radio waves. Likewise, the data processing 'spectrum' that we can perceive is almost negligible compared to what is really going on in machine-to-machine and human–computer interactions. An immense number of ICT applications run an incalculable number of instructions every millisecond of our lives to keep the hyperhistorical

information society humming. ICTs consume most of their MIPS to talk to each other, collaborate, and coordinate efforts, and put us as comfortably as possible in or on the loop, or even out of it, when necessary. According to a recent White Paper published by CISCO IBSG,[3] a multinational corporation that admittedly designs, manufactures, and sells networking equipment, there will be 25 billion devices connected to the Internet by 2015 and 50 billion by 2020 (see Figure 8).

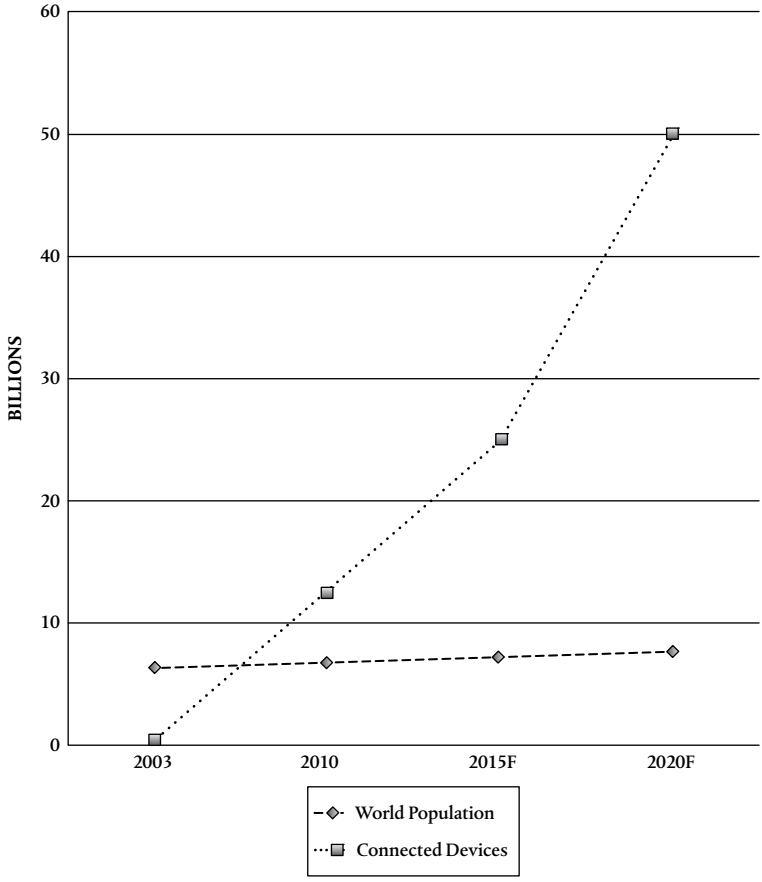


Fig. 8. The Growth of World Population and of Connected Devices. F = forecast.
Data source: Evans (2011).

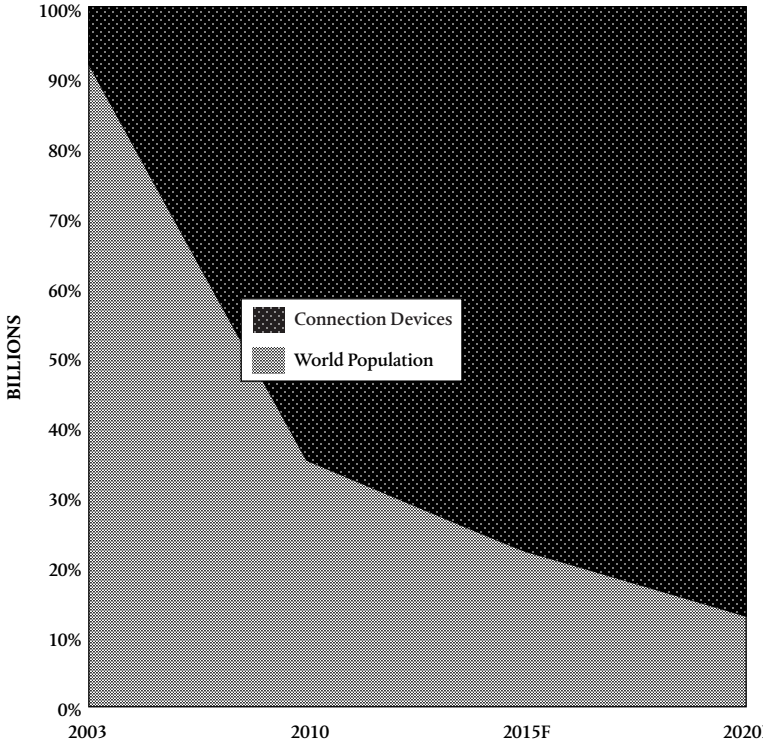


Fig. 9. The Total Space of Connectivity in Relation to the Growth of World Population and of Connected Devices. F = forecast.

Data source: Evans (2011).

The number of connected devices per person will grow from 0.08 in 2003, to 1.84 in 2010, to 3.47 in 2015, to 6.58 in 2020. To our future historian, global communication on Earth will appear to be largely a non-human phenomenon, as Figure 9 illustrates.

Almost all MIPS are invisible to us, like the oxygen we breathe, but they are becoming almost as vital, and they are growing exponentially. Computational devices of all sorts generate a staggering amount of data, much more data than humanity has ever seen in its entire history (Figure 10). This is the other resource that has made hyperhistory possible: zettabytes.

Data

A few years ago, researchers at Berkeley's School of Information[4] estimated that humanity had accumulated approximately 12 exabytes[5] of data in the course of its entire history until the commodification of computers, but that it had already reached 180 exabytes by 2006. According to a more recent study,[6] the total grew to over 1,600 exabytes between 2006 and 2011, thus passing the zettabyte (1,000 exabytes) barrier. This figure is now expected to grow fourfold approximately every three years, so that we shall have 8 zettabytes of data by 2015. Every day, enough new data are being generated to fill all US libraries eight times over. Of course, armies of ICT devices are constantly working to keep us afloat and navigate through such an ocean of data. These are all numbers that will keep growing quickly and steadily for the foreseeable future, especially because those very devices are among the greatest sources of further data, which in turn require, or simply make possible, more ICTs. It is a self-reinforcing cycle and it would be unnatural not to feel overwhelmed. It is, or at least should be, a mixed feeling of apprehension about the risks, excitement at the opportunities, and astonishment at the achievements, as we shall see in the following chapters.

Thanks to ICTs, we have entered *the age of the zettabyte*. Our generation is the first to experience a zettaflood, to introduce a neologism to describe this tsunami of bytes that is submerging our environments. In other contexts, this is also known as 'big data' (Figure 10).

Despite the importance of the phenomenon, it is unclear what exactly the term 'big data' means. The temptation, in similar cases, is to adopt the approach pioneered by Potter Stewart, United States Supreme Court Justice, when asked to describe pornography: difficult to define, but 'I know when I see it'. Other strategies have been much less successful. For example, in the United States, the National Science Foundation (NSF) and the National Institutes of Health (NIH) have identified big data as a programme focus. One of the main NSF–NIH

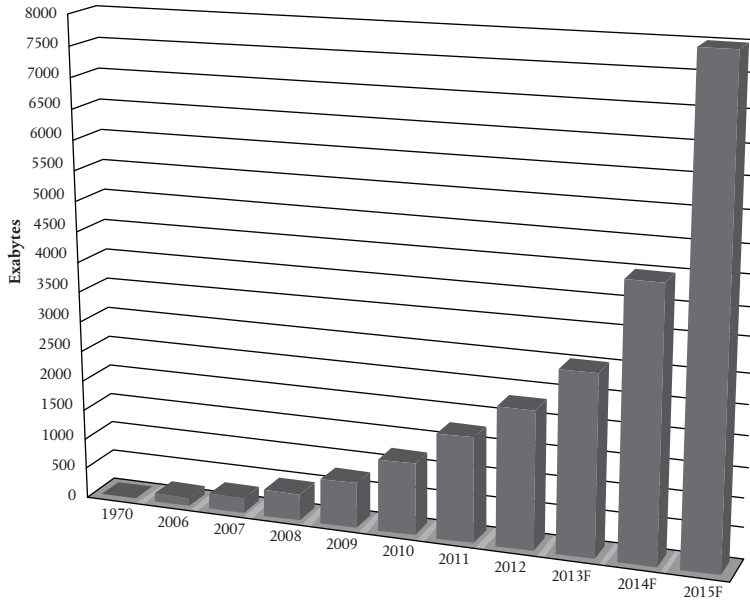


Fig. 10. The Growth of Big Data.

Source: based on IDC white paper, 'The Diverse and Exploding Digital Universe', March 2008, and IDC white paper, 'Worldwide Big Data Technology and Service 2012–2015 Forecast', March 2012.

interagency initiatives addresses the need for core techniques and technologies for advancing big data science and engineering. However, the two agencies specify that

> The phrase 'big data' in this solicitation refers to large, diverse, complex, longitudinal, and/or distributed data sets generated from instruments, sensors, Internet transactions, email, video, click streams, and/or all other digital sources available today and in the future.[7]

You do not need to be a logician to find this both obscure and vague. Wikipedia, for once, is also unhelpful. Not because the relevant entry is unreliable, but because it reports the common definition, which describes 'big data' as a collection of data sets so large and complex that it becomes difficult to process using available management tools or traditional data-processing applications. Apart from the circular

problem of defining 'big' with 'large' (the NSF and NHI seem to be happy with it), understanding 'big data' in terms of 'small tools' suggests that data are too big or large only in relation to our current computational power. This is misleading. Of course, 'big', as many other terms, is a relational predicate: a pair of shoes may be too big for you, but fine for me. It is also trivial to acknowledge that we tend to evaluate things non-relationally, in this case as absolutely big, whenever the frame of reference is obvious enough to be left implicit. A horse is a big animal, no matter what whales may think. Yet these two simple points may give the impression that there is no real trouble with 'big data' being a loosely defined term referring to the fact that our current computers cannot handle so many gazillions of data efficiently. And this is where two confusions seem to creep in. First, that the *epistemological* (that is, knowledge-related) *problem* with big data is that there is too much of it (the *ethical problem* concerns how we use them, more on this presently). And, second, that the *solution* to the epistemological problem is *technological*: more and better techniques and technologies, which will 'shrink' big data back to a manageable size. The epistemological problem is different, and it requires an epistemological solution.

Consider the problem first. 'Big data' came to be formulated after other buzz expressions, such as 'infoglut' or 'information overload', began to fade away, yet the idea remains the same. It refers to an overwhelming sense that we have bitten off more than we can chew, that we are being force-fed like geese, that our intellectual livers are exploding. This is a mistake. Yes, we have seen that there is an obvious exponential growth of data on an ever-larger number of topics, but complaining about such over-abundance would be like complaining about a banquet that offers more than we can ever eat. Data remain an asset, a resource to exploit. Nobody is forcing us to digest every available byte. We are becoming data-richer by the day; this cannot be the fundamental problem.

Since the problem is not the increasing wealth of data that is becoming available, clearly the solution needs to be reconsidered: it

cannot be merely how many data we can technologically process. We saw that, if anything, more and better techniques and technologies are only going to generate more data. If the problem were too many data, more ICTs would only exacerbate it. Growing bigger digestive systems, as it were, is not the way forward.

The real epistemological problem with big data is *small patterns*. Precisely because so many data can now be generated and processed so quickly, so cheaply, and on virtually anything, the pressure both on the data *nouveau riche*, such as Facebook or Walmart, Amazon or Google, and on the data *old money*, such as genetics or medicine, experimental physics or neuroscience, is to be able to spot where the new patterns with real added-value lie in their immense databases, and how they can best be exploited for the creation of wealth, the improvement of human lives, and the advancement of knowledge. This is a problem of brainpower rather than computational power.

Small patterns matter because, in hyperhistory, they represent the new frontier of innovation and competition, from science to business, from governance to social policies, from security to safety. In a free and open marketplace of ideas, if someone else can exploit the small patterns earlier and more successfully than you do, you might quickly be out of business, miss a fundamental discovery and the corresponding Nobel, or put your country in serious danger.

Small patterns may also be risky, because they push the limit of what events or behaviours are predictable, and therefore may be anticipated. This is an ethical problem. Target, an American retailing company, relies on the analysis of the purchasing patterns of 25 products in order to assign each shopper a 'pregnancy prediction' score, estimate her due date, and send coupons timed to specific stages of her pregnancy. In a notorious case,[8] it caused some serious problems when it sent coupons to a family in which the teenager daughter had not informed her parents about her new status. I shall return to this sort of problem in Chapters 3 and 4, when discussing personal identity and privacy.

Unfortunately, small patterns may be significant only if properly aggregated, correlated, and integrated—for example in terms of

loyalty cards and shopping suggestions—compared, as when banks utilize big data to fight fraudsters, and processed in a timely manner, as in financial markets. And because information is indicative also when it is not there (the lack of some data may also be informative in itself), small patterns can also be significant if they are absent. Sherlock Holmes solves one of his famous cases because of the silence of the dog, which should have barked. If big data are not 'barking' when they should, something is going on, as the financial watchdogs (should) know, for example.

Big data is here to grow. The only way of tackling it is to know what you are or may be looking for. We do not do science by mere accumulation of data; we should not do business and politics in that way either. At the moment, the required epistemological skills are taught and applied by a black art called analytics. Not exactly your standard degree at university. Yet so much of our well-being depends on it that it might be time to develop a methodological investigation of it. Who knows, philosophers might not only have something to learn, but also a couple of lessons to teach. Plato would agree. What he might have been disappointed about is the fact that memory is no longer an option. As we shall see in Chapter 7, memory may outperform intelligence, but mere data hoarding, while waiting for more powerful computers, smarter software, and new human skills, will not work, not least because we simply do not have enough storage. Recall our future historian: this is the first bottleneck she identified in the development of hyperhistory, which suffers from digital amnesia.

Memory

Hyperhistory depends on big data, but there are two myths about the dependability of digital memory that should be exposed in this first chapter.

The first myth concerns the *quality* of digital memory. ICTs have a kind of forgetful memory. They become quickly obsolete, they are volatile, and they are rerecordable. Old digital documents may no

longer be usable because the corresponding technology, for example floppy drives or old processing software, is no longer available. There are millions of abandoned pages on the Internet, pages that have been created and then not updated or modified. At the beginning of 1998, the average life of a document that had not been abandoned was 75 days. It is now estimated to be 45 days. The outcome is that so-called link decay (links to resources online that no longer work) is a common experience. On 30 April 1993, the European Organization for Nuclear Research (CERN) announced that the World Wide Web it had created would be free to everyone, with no fees due. Twenty years later, to celebrate the event, a team at CERN had to recreate the first web page (with its original URL etc.), because it no longer existed. Our digital memory seems as volatile as our oral culture was but perhaps even more unstable, because it gives us the opposite impression. This paradox of a digital 'prehistory'—ICTs are not preserving the past for future consumption because they make us live in a perennial present—will become increasingly pressing in the near future. Memory is not just a question of storage and efficient management; it is also a matter of careful curation of significant differences, and hence of the stable sedimentation of the past as an ordered series of changes, two historical processes that are now seriously at risk. Ted Nelson, for example, a pioneer in ICTs who coined the terms 'hypertext' and 'hypermedia', designed Xanadu so that it would never delete copies of old files. A website constantly upgraded is a site without memory of its own past, and the same dynamic system that allows one to rewrite a document a thousand times also makes it unlikely that any memory of past versions will survive for future inspection. 'Save this document' means 'replace its old version', and every digital document of any kind may aspire to such an ahistorical nature. The risk is that differences are erased, alternatives amalgamated, the past constantly rewritten, and history reduced to the perennial here and now. When most of our knowledge is in the hands of this forgetful memory, we may find ourselves imprisoned in a perpetual present. This is why initiatives aimed at preserving our increasingly digital cultural heritage for future

generations—such as the National Digital Stewardship Alliance (NDSA) and the International Internet Preservation Consortium (IIPC)—are vital. The job of information curators is bound to become ever more important.

There is then the potentially catastrophic risk of immense quantities of data being created simultaneously. We saw that most, indeed almost all, of our data have been created in a matter of a few years. They are all getting old together, and will reach the threshold of system failure together, like a baby-boom generation retiring at the same time. To understand the problem, recall the old debate about your collection of music CDs and how they would all be unusable within a decade, as opposed to your vinyl records. According to the Optical Storage Technology Association, the shelf life of new, unrecorded CDs and DVDs is conservatively estimated to be between 5 and 10 years. And according to the National Archives and Records Administration,[9] once recorded, CDs and DVDs have a life expectancy of 2 to 5 years, despite the fact that published life expectancies are often cited as 10 years, 25 years, or longer. The problem is that after a few years the material degrades too much to guarantee usability. The same applies to our current digital supports, hard disks and memories of various kinds. The 'mean time before failure' (MTBF) figure indicates an estimate of a system's life expectancy.[10] The higher the MTBF, the longer the system should last. An MTFB of 50,000 hours (5.7 years) for a standard hard disk is not uncommon. This short life expectancy is already a problem. But the real issue that I am stressing here is another. Contrary to what we experienced in the past, the life expectancies of our data supports are today dangerously synchronized. This is why you may think of this as a sort of 'baby boom': big data will age and become dead data together. Clearly, huge quantities of data will need to be rerecorded and transferred to new supports at regular intervals. Indeed they already are. But which data are going to make it to the other side of any technological transition? For a comparison, consider the transition of silent movies to new kinds of support, or of recorded music from

vinyl to the CD. Huge quantities of data were left behind, becoming lost, unavailable, or inaccessible.

According to a 2012 Research Report by IBIS World, the data recovery industry saw its overall revenue over the five years to 2012 fall at an annualized rate of 0.9 per cent to total $1 billion, with a decline of 0.6 per cent in 2012.[11] This may seem counter-intuitive. Big data is growing and so are the problems concerning damaged, corrupted, or inaccessible files and storage media. The industry that takes care of such problems should be flourishing. The explanation is that cloud or online storage has expanded the options for data recovery and data loss prevention. If you use Dropbox, Google Docs, Apple iCloud, or Microsoft Skydrive, for example, and your computer is damaged, the files are still available online and can be easily recovered, so you will not need a data recovery service. Yet, this seems to be just a question of transition and hence time. Cloud computing has put pressure on an industry specialized in computers at a consumer level. The more our gadgets become mere terminals, the less we need to worry ourselves about the data. But the storage of those data still relies on physical infrastructures, and these will need increasing maintenance. The data recovery industry will disappear, but a new industry dedicated to cloud computing failures is already emerging. It is not a matter of relying on the brute force of redundancy (having more than one copy of the same file). This strategy is not available at a global level, because of the second myth about the dependability of digital memory, the one concerning the *quantity* of digital memory.

Since 2007, the world has been producing many more data than available storage.[12] This despite the fact that, according to Kryder's Law (another generalization), storage density of hard disks is increasing more quickly than Moore's Law, so that it is predicted that in 2020 a disk of 14 terabytes will be 2.5 inches in size and will cost about $40. Unfortunately, this will not be enough, because even the growth projected by Kryder's Law is slow when compared to the pace at which we generate new data. Think of your smartphone becoming too full because you took too many pictures, and make it a global

problem. In history, the problem was what to save: which laws or names were going to be baked in clay or carved in stone, which texts were going to be handwritten on papyrus or vellum, which news items were worth printing on paper. In hyperhistory, saving is the default option. The problem becomes what to erase. Since storage is insufficient, something must be deleted, rewritten, or never be recorded in the first place. By default the new tends to push out the old, or 'first in first out': updated web pages erase old ones, new pictures make old ones look dispensable, new messages are recorded over old ones, recent emails are kept at the expense of last year's.

Hyperhistory ran out of memory space in which to dump its data many years ago. There is no name for this 'law' about the increasing shortage of memory, but it looks like the gap is doubling every year. Barring some significant technological breakthrough in physical storage or software compression, the process will get worse, *quantitatively*. The good news is that it does not have to be as bad as it looks, *qualitatively*. Rephrasing a common saying in the advertising industry, half of our data is junk, we just do not know which half. You are happy to take ten pictures because you hope one will come out right, and the other nine can be discarded. They were never intended to be saved in the first place. This means that we need a much better understanding of which data are worth preserving and curating. This, in turn, is a matter of grasping which questions are, or will be, interesting to ask not only now, but also in the future, as we saw in the previous section. And this leads to a slightly reassuring virtuous circle: we should soon be able to ask big data what data are worth saving. Think of an app in your smartphone not only suggesting which of the ten pictures is worth keeping, but also learning from you, once you have taken a decision (maybe you prefer darker pictures). Then new challenges will concern how we may avoid poor machine-based decisions, improve so-called 'machine learning', or indeed make sure machines relearn new preferences (later in life you may actually like brighter pictures). More information may help us to decide which information to save and curate. Our future historian may well

interpret the zettabyte age of hyperhistory as the time of transition between blind and foresighted big data.

So much for the first bottleneck: shortage of memory. To understand the other, concerning connectivity, we need to look first at some features of networks.

Connectivity

Computers may be of limited use if they are not connected to other computers. This was not always obvious. Sometimes it is still questionable, as when your computers need to be hacker-proof because they control the launch of nuclear missiles, for example. But, in general, the observation is rather trivial today. In the age in which tethering no longer means restraining an animal with a rope or chain (the tether) but actually connecting one ICT device to another, the question is no longer whether connectivity has any value, but how much value it actually has. Many theories and laws have been proposed: Reed's Law, Sarnoff's Law, Beckstrom's Law . . . but the most famous remains Metcalfe's. Like the laws just mentioned and Moore's Law, it is a generalization ('this is how things tend to go, more or less') rather than a scientific law, but it is, nevertheless, enlightening. It states that the value of a network is proportional to the square of the number of connected nodes of the system (n^2). So a network of two computers has only a value of $2^2=4$, but doubling the number of connected computers already means quadrupling the value of their network to $4^2=16$. Figure 11 illustrates what happens after 20 iterations. The idea is simple: the more nodes you have, the more useful it is to be connected and expensive to be unconnected. Indeed, the point to keep in mind is even simpler, for there is an even more inclusive generalization. Any growth bigger than linear (a linear growth is when you multiply x by a fixed number, like your salary by 12 months), e.g. squared, like Metcalfe's, or cubic (n^3), or exponential (e^x), after a few iterations looks like a straight perpendicular line, like a capital L which has been rotated 180 degrees on its axis:

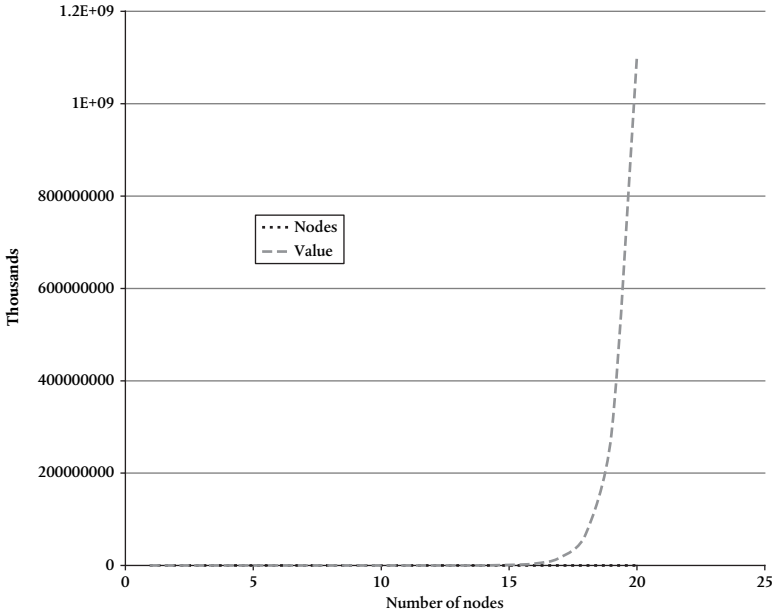


Fig. 11. Metcalfe's Law: The Value of a Network of n Nodes = n^2.

This 'L law' is all one needs to remember. It is the shape of growth that any business would like to achieve. It is the shape of hyper-connectivity. According to a report by the International Telecommunication Union (ITU),[13] in 2013 more than a third of the world population was online. No wonder the value of the network has skyrocketed, straight like a rod. So what is the problem? Any L Law does not really address the communication within the network, but rather the value of its complexity (how many links are possible among how many nodes). Communication requires a link but it comes with a speed. Think of a road, and the difference it makes whether it is a small street or a motorway, with or without traffic. This is the bottleneck our future historian identified. It is known as Nielsen's Law.

Some years ago, Jacob Nielsen noticed that, in general, the speed of network connections for home users like you and me increases approximately 50 per cent per year, thus doubling every 21 months

or so. This is impressive, but not as impressive as the speed identified by Moore's Law. It is also already insufficient to cope with the faster growing 'weight' (number of bits) of the files we wish to transfer. As a result, for the foreseeable future our online experience will be constrained by our bandwidth.

Conclusion

The living generation is experiencing a transition from history to hyperhistory. Advanced information societies are more and more heavily dependent on ICTs for their normal functioning and growth. Processing power will increase, while becoming cheaper. The amount of data will reach unthinkable quantities. And the value of our network will grow almost vertically. However, our storage capacity (space) and the speed of our communications (time) are lagging behind. Hyperhistory is a new era in human development, but it does not transcend the spatio-temporal constraints that have always regulated our life on this planet. The question to be addressed next is: given all the variables we have seen in this chapter, what sort of hyperhistorical environment are we building for ourselves and for future generations? The short answer is: the *infosphere*. The long answer is provided by Chapter 2.