

Cambridge University Press

978-0-521-88898-1 - The Cambridge Handbook of Information and Computer Ethics

Edited by Luciano Floridi

Excerpt

[More information](#)

Part I

Introduction and background

1 Ethics after the Information Revolution

Luciano Floridi

1.1 Introduction: history as the information age

Humanity has organized its history according to many metrics. Some are natural and circular, relying on seasons and planetary motions. Some are social or political and linear, being determined, for example, by the succession of Olympic Games, the number of years since the founding of the city of Rome (*ab urbe condita*), or the ascension of a king. Still others are religious and have a V-shape, counting years before and after a particular event (e.g. the birth of Christ). There are larger periods that encompass smaller ones, named after influential styles (Baroque), people (Victorian era), particular circumstances (Cold War) or some new technology (Nuclear age). What all these and many other metrics have in common is that they are all *historical*, in the strict sense that they all depend on the development of systems to record events and hence accumulate and transmit information about the past. It follows that history is actually synonymous with the information age, since *prehistory* is the age in human development that precedes the availability of recording systems. Hence, one may further argue 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 different regions of the world, and especially in Mesopotamia. Comparing the computer revolution to the printing revolution would be misleading not because they are unrelated, but because they are actually phases of a much wider, macroscopic process that has spanned millennia: the slow emergence of the information society since the fourth millennium BC. And yet, this is not what we normally mean when talking about the information age. Typically, we have in mind something much more limited in scope and closer in time. There may be many explanations, but one seems more convincing than any other: only very recently has human progress and welfare begun to depend mostly on the successful and efficient management of the information life cycle.¹ So the long period of time that

¹ A typical life cycle includes the following phases: occurring (discovering, designing, authoring, etc.), processing and managing (collecting, validating, modifying, organizing, indexing, classifying, filtering, updating, sorting, storing, networking, distributing, accessing, retrieving, transmitting, etc.) and using (monitoring, modelling, analysing, explaining, planning, forecasting, decision-making, instructing, educating, learning, etc.).

the information society has taken to surface should not be surprising. Imagine a historian writing in a million years from now. She may consider it normal, and perhaps even elegantly symmetrical, that it took roughly six millennia (from its beginning in the Neolithic, tenth millennium BC, until the Bronze Age) for the agricultural revolution to produce its full effect, and then another six millennia (from the Bronze Age until the end of the second millennium AD) for the information revolution to bear its main fruit. During this span of time, information technologies evolved from being mainly recording systems, to being also communication systems (especially after Gutenberg), to being also processing systems (especially after Turing). As I will explain below, they have begun to play the role of re-ontologizing systems. Thanks to this evolution, nowadays the most advanced economies are highly dependent, for their functioning and growth, upon the pivotal role played by information-based, intangible assets, information-intensive services (especially business and property services, communications, finance and insurance, and entertainment) as well as information-oriented public sectors (especially education, public administration and health care). For example, all G7 members qualify as information societies because, in Canada, France, Germany, Italy, Japan, United Kingdom, and the United States of America, at least 70% of the Gross Domestic Product (GDP) depends on intangible goods, which are information-based, rather than material goods, which are the physical output of agricultural or manufacturing processes.

The almost sudden burst of a global information society, after a few millennia of relatively quieter gestation, has generated new and disruptive challenges, which were largely unforeseeable only a few decades ago. Needless to say, Information and Communication Technologies (ICTs) have been changing the world profoundly, irreversibly and problematically since the fifties, at a breathtaking pace, and with unprecedented scope, making the creation, management and utilization of information, communication and computational resources vital issues. As a quick reminder, and in order to have some simple, quantitative measure of the transformations experienced by our generation, consider the following findings.

In a recent study, researchers at Berkeley's School of Information Management and Systems estimated that humanity had accumulated approximately 12 exabytes² of data in the course of its entire history until the commodification of computers, but that it had produced more than 5 exabytes of data just in 2002, 'equivalent in size to the information contained in 37,000 new libraries the size of the Library of Congress book collections' (Lyman and Varian 2003). In 2002, this was almost 800 MB of recorded data produced per person. It is like saying that every newborn baby came into the world with a

² One exabyte corresponds to 1,000,000,000,000,000,000 bytes or 10^{18} .

burden of 30 feet of books, the equivalent of 800 MB of data on paper. This exponential escalation has been relentless: 'between 2006 and 2010... the digital universe will increase more than six fold from 161 exabytes to 988 exabytes'.³

Not feeling under pressure would be abnormal. The development of ICT has not only brought enormous benefits and opportunities but also greatly outpaced our understanding of its conceptual nature and implications, while raising problems whose complexity and global dimensions are rapidly expanding, evolving and becoming increasingly serious. A simple analogy may help to make sense of the current situation. Our technological tree has been growing its far-reaching branches much more widely, rapidly and chaotically than its conceptual, ethical and cultural roots. The lack of balance is obvious and a matter of daily experience in the life of millions of citizens dealing with information-related ethical issues. The risk is that, like a tree with weak roots, further and healthier growth at the top might be impaired by a fragile foundation at the bottom. As a consequence, today, any advanced information society faces the pressing task of equipping itself with a viable philosophy and ethics of information. Applying the previous analogy, while technology keeps growing bottom-up, it is high time we start digging deeper, top-down, in order to expand and reinforce our conceptual understanding of our information age, of its nature, its less visible implications and its impact on human and environmental welfare, and thus give ourselves a chance to anticipate difficulties, identify opportunities and resolve problems, conflicts and dilemmas.

It is from such a broad perspective that I would like to invite the reader to approach this volume. The chapters constituting it perfectly complement each other. Written by leading experts in the area, they tackle some of the key issues in information and computer ethics (ICE). Since the authors need no introduction, and the contents of the chapters are outlined in the preface, in the rest of this introductory chapter my contribution will be to discuss some conceptual undercurrents, which flow beneath the surface of the literature on ICE, and may be seen surfacing in different places throughout this book. In discussing them, I shall focus, more generally, on the potential impact of ICT on our lives. And since there would be no merit in predicting the obvious, I will avoid issues such as rising concerns about privacy and identity theft, spamming, viruses, or the importance of semantic tagging, online shopping and virtual communities. Nor will I try to steal ideas from those who know better than I do the future development of the actual technologies (see for example O'Reilly 2005, Microsoft-Research 2005, *Nature* 2006). I will, instead,

³ Source: 'The Expanding Digital Universe: A Forecast of Worldwide Information Growth Through 2010', white paper – sponsored by EMC – IDC, www.emc.com/about/destination/digital_universe/

stick to what philosophers do better, conceptual engineering, and seek to capture the silent *Weltanschauung* that might be dawning on us.

1.2 ICT as re-ontologizing technologies

In order to grasp the ICT scenarios that we might witness and experience in the near future, and hence the sort of ethical problems we might be expected to deal with, it is useful to introduce two key concepts at the outset, those of 'infosphere' and of 're-ontologization'.

Infosphere is a neologism I coined some years ago (Floridi 1999a) on the basis of 'biosphere', a term referring to that limited region on our planet that supports life. It denotes the whole informational environment constituted by all informational entities (thus including informational agents as well), their properties, interactions, processes and mutual relations. It is an environment comparable to, but different from, cyberspace (which is only one of its sub-regions, as it were), since it also includes offline and analogue spaces of information. We shall see that it is also an environment (and hence a concept) that is rapidly evolving.

Re-ontologizing is another neologism that I have recently introduced in order to refer to a very radical form of re-engineering, one that not only designs, constructs or structures a system (e.g. a company, a machine or some artefact) anew, but one that also fundamentally transforms its intrinsic nature, that is, its ontology or essence. In this sense, for example, nanotechnologies and biotechnologies are not merely re-engineering but actually re-ontologizing our world.

Using the two previous concepts, it becomes possible to formulate succinctly the following thesis: ICTs are re-ontologizing the very nature of (and hence what we mean by) the infosphere, and here lies the source of some of the most profound transformations and challenging problems that our information societies will experience in the close future, as far as technology is concerned.

The most obvious way in which ICTs are re-ontologizing the infosphere concerns the transition from analogue to digital data and then the ever-increasing growth of our informational space. Both phenomena are very familiar and require no explanation, but a brief comment may not go amiss.

Although the production of analogue data is still increasing, the infosphere is becoming more digital by the day. A simple example may help to drive the point home: the new Large Hadron Collider built at the CERN (<http://lhc.web.cern.ch/lhc/>) to explore the physics of particles produces about 1.5 GB data per second, or about 10 petabytes of data annually, a quantity of data a thousand times larger than the Library of Congress's print collection and at least twice as large as Google's whole data storage, reported to be approximately 5 petabytes in 2004 (Mellor 2004).

This radical re-ontologization of the infosphere is largely due to the fundamental convergence between digital resources and digital tools. The ontology of the information technologies available (e.g. software, databases, communication channels and protocols, etc.) is now the same as (and hence fully compatible with) the ontology of their objects. This was one of Turing's most consequential intuitions: in the re-ontologized infosphere, there is no longer any substantial difference between the *processor* and the *processed*, so the digital deals effortlessly and seamlessly with the digital. This potentially eliminates one of the most long-standing bottlenecks in the infosphere and, as a result, there is a gradual erasure of *ontological friction*.

Ontological friction refers to the forces that oppose the flow of information within (a region of) the infosphere, and hence (as a coefficient) to the amount of work and effort required to generate, obtain, process and transmit information in a given environment, e.g. by establishing and maintaining channels of communication and by overcoming obstacles in the flow of information such as distance, noise, lack of resources (especially time and memory), amount and complexity of the data to be processed, and so forth. Given a certain amount of information *available* in (a region of) the infosphere, the lower the ontological friction within it, the higher the *accessibility* of that amount of information becomes. Thus, if one quantifies ontological friction from 0 to 1, a fully successful firewall would produce a 1.0 degree of friction, i.e. a complete standstill in the flow of information through its 'barrier'. On the other hand, we describe our society as informationally porous the more it tends towards a 0 degree of informational friction.

Because of their 'data superconductivity', ICTs are well known for being among the most influential factors that affect the ontological friction in the infosphere. We are all acquainted daily with aspects of a *frictionless infosphere*, such as *spamming* and *micrometering* (every fraction of a penny counts). Other significant consequences include (a) a substantial erosion of the *right to ignore*: in an increasingly porous society, it becomes progressively less credible to claim ignorance when confronted by easily predictable events (e.g. as George W. Bush did with respect to Hurricane Katrina's disastrous effects on New Orleans's flood barriers) and hardly ignorable facts (e.g. as Tessa Jowell, a British Labour MP, did with respect to her husband's finances). And therefore (b) an exponential increase in *common knowledge*: this is a technical term from epistemic logic, which basically refers to the case in which everybody not only knows that *p* but also knows that everybody knows that everybody knows, . . . , that *p*. In other words, (a) and (b) will also be the case because meta-information about how much information is, was or should have been available will become overabundant. From (a) and (b) it follows that, in the future, (c) we shall witness a steady increase in agents' *responsibilities*. As I shall argue towards the end of this chapter, ICTs are making humanity increasingly responsible, morally speaking, for the way the world is, will and should be (Floridi and Sanders 2001, Floridi 2006b).

1.3 The global infosphere or how information is becoming our ecosystem

During the last decade or so, we have become accustomed to conceptualizing our life online as a mixture between an evolutionary adaptation of human agents to a digital environment, and a form of post-modern, neo-colonization of the latter by the former. This is probably a mistake. ICTs are as much re-ontologizing our world as they are creating new realities. The threshold between *here* (*analogue, carbon-based, offline*) and *there* (*digital, silicon-based, online*) is fast becoming blurred, but this is as much to the advantage of the latter as it is to the former. The digital is spilling over into the analogue and merging with it. This recent phenomenon is variously known as 'Ubiquitous Computing', 'Ambient Intelligence', 'The Internet of Things' or 'Web-augmented Things'. It is, or will soon be, the next stage in the development of the information age.

The increasing re-ontologization of artefacts and of whole (social) environments suggests that soon it will be difficult to understand what life was like in predigital times and, in the near future, the very distinction between online and offline will become blurred and then disappear. To someone who was born in 2000 the world will always have been wireless, for example. To her, the peculiar clicking and whooshing sounds made by conventional modems while handshaking will be as alien as the sounds made by a telegraph's Morse signals. To put it dramatically, the infosphere is progressively absorbing any other ontological space. Let me explain.

In the (fast-approaching) future, more and more objects will be *ITentities* able to learn, advise and communicate with each other. A good example (but it is only an example) is provided by RFID (Radio Frequency Identification) tags, which can store and remotely retrieve data from an object and give it a unique identity, like a barcode. Tags can measure 0.4 mm² and are thinner than paper. Incorporate this tiny microchip in everything, including humans and animals, and you have created *ITentities*. This is not science fiction. According to a report by market research company InStat, the worldwide production of RFID will increase more than 25-fold between 2005 and 2010 and reach 33 billion. Imagine networking these 33 billion *ITentities* together with all the hundreds of millions of PCs, DVDs, iPods and ICT devices available and you see that the infosphere is no longer 'there' but 'here' and it is here to stay. Your Nike and iPod already talk to each other, with predictable (but amazingly unforeseen) problems in terms of privacy (Saponas *et al.* 2007).

Nowadays, we are still used to considering the space of information as something we log-in to and log-out from. Our view of the world (our metaphysics) is still modern or Newtonian: it is made of 'dead' cars, buildings, furniture, clothes, which are non-interactive, irresponsive and incapable of communicating, learning or memorizing. But, as I shall argue in the next section, what

we still experience as the world offline is bound to become a fully interactive and responsive environment of wireless, pervasive, distributed, *a2a* (anything to anything) information processes, that works *a4a* (anywhere for anytime), in real time. The day when we routinely google the location of physical objects ('where are the car keys?') is very close.⁴

As a consequence of such re-ontologization of our ordinary environment, we shall be living in an infosphere that will become increasingly *synchronized* (time), *delocalized* (space) and *correlated* (interactions). Although this might be read, optimistically, as the friendly face of globalization, we should not harbour illusions about how widespread and inclusive the evolution of information societies will be. The digital divide will become a chasm, generating new forms of discrimination between those who can be denizens of the infosphere and those who cannot, between insiders and outsiders, between information rich and information poor. It will redesign the map of worldwide society, generating or widening generational, geographic, socio-economic and cultural divides. But the gap will not be reducible to the distance between industrialized and developing countries, since it will cut across societies (Floridi 2002a). We are preparing the ground for tomorrow's informational slums.

1.4 The metaphysics of the infosphere

The previous transformations will invite us to understand the world as something 'a-live' (artificially live). Such animation of the world will, paradoxically, make our outlook closer to that of pre-technological cultures which interpreted all aspects of nature as inhabited by teleological forces. The second step will be a reconceptualization of our ontology in informational terms. It will become normal to consider the world as part of the infosphere, not so much in the dystopian sense expressed by a *Matrix*-like scenario, where the 'real reality' is still as hard as the metal of the machines that inhabit it, but in the evolutionary, hybrid sense represented by an environment such as New Port City, the fictional, post-cybernetic metropolis of *Ghost in the Shell*. The infosphere will not be a virtual environment supported by a genuinely 'material' world behind; rather, it will be the world itself that will be increasingly interpreted and understood informationally, as part of the infosphere. At the end of this shift, the infosphere will have moved from being a way to refer to the space of information to being synonymous with Being. Thus, our way of conceptualizing and making sense of reality will keep shifting from a materialist perspective, in which physical objects and processes still play a key role, to an informational one, in which

⁴ In 2008, Thomas Schmidt, Alex French, Cameron Hughes and Angus Haines (four 12-year-old boys from Ashfold Primary School in Dorton, UK) were awarded the 'Home Invention of the Year' Prize for their Speed Searcher, a device for finding lost items. It attaches tags to valuables and enables a computer to pinpoint their location in the home.

- objects and processes are dephysicalized, typified and perfectly clonable;
- the right of usage is at least as important as the right to ownership; and
- the criterion for existence is no longer being immutable (Greek metaphysics) or being potentially subject to perception (modern metaphysics) but being interactable.

If all this seems a bit too ‘philosophical’, let me provide an illustrative example.

Despite some important exceptions (e.g. vases and metal tools in ancient civilizations or books after Gutenberg), it was the industrial revolution that really marked the passage from a nominalist world of unique objects to a Platonist world of types of objects, all perfectly reproducible as identical to each other, therefore epistemically indiscernible, and hence pragmatically dispensable because replaceable without any loss. Today, we find it obvious that two automobiles may be virtually identical and that we are invited to buy a model rather than a specific ‘incarnation’ of it. Indeed, we are fast moving towards a commodification of objects that considers repair as synonymous with replacement, even when it comes to entire buildings. This has led, by way of compensation, to a prioritization of *branding* – a process compared by Klein (2000) to the creation of ‘cultural accessories and personal philosophies’ – and of *re-appropriation*: the person who puts a sticker on the window of her car, which is otherwise perfectly identical to thousands of others, is fighting an anti-Platonic battle. The information revolution has further exacerbated this process. Once our window-shopping becomes Windows-shopping and no longer means walking down the street but browsing through the Web, the problem caused by the dephysicalization and typification of individuals as unique and irreplaceable entities starts eroding our sense of personal identity as well. We become mass-produced, anonymous entities among other anonymous entities, exposed to billions of other similar inforgs online. So we construct, self-brand and re-appropriate ourselves in the infosphere by blogs and FaceBook entries, homepages, YouTube videos, flickr albums, fashionable clothes and choices of places we visit, types of holidays we take and cars we drive and so forth. We use and expose information about ourselves to become less informationally indiscernible. We wish to maintain a high level of informational privacy almost as if that were the only way of saving a precious capital which can then be publicly invested by us in order to construct ourselves as individuals discernible and easily re-identifiable by others. Now, processes such as the one I have just sketched are part of a far deeper metaphysical drift caused by the information revolution.

1.5 The information turn as the fourth revolution

Oversimplifying more than a bit, one may say that science has two fundamental ways of changing our understanding. One may be called *extrovert*,

or about the world, and the other *introvert*, or about ourselves. Three scientific revolutions have had great impact in both ways. They changed not only our understanding of the external world, but, in doing so, they also modified our conception of who we are. After Nicolaus Copernicus, the heliocentric cosmology displaced the Earth and hence humanity from the centre of the Universe. Charles Darwin showed that all species of life have evolved over time from common ancestors through natural selection, thus displacing humanity from the centre of the biological kingdom. Thirdly, following Sigmund Freud, we acknowledge nowadays that the mind is also unconscious and subject to the defence mechanism of repression, thus displacing it from the centre of pure rationality, a position that had been assumed as uncontroversial at least since Descartes. The reader who, like Popper, would be reluctant to follow Freud in considering psychoanalysis a scientific enterprise, might yet be willing to concede that contemporary neuroscience is a likely candidate for such a revolutionary role. Either way, the result is that we are not immobile, at the centre of the Universe (Copernican revolution), we are not unnaturally separate and diverse from the rest of the animal kingdom (Darwinian revolution), and we are very far from being Cartesian minds entirely transparent to ourselves (Freudian or Neuroscientific revolution).

Freud (1917) was the first to interpret these three revolutions as part of a single process of reassessment of human nature (see Weinert 2009). The hermeneutic manoeuvre was, admittedly, rather self-serving. But it did strike a reasonable note. In a similar way, when we now perceive that something very significant and profound has happened to human life after the informational turn, I would argue that our intuition is once again perceptive, because we are experiencing what may be described as a fourth revolution, in the process of dislocation and reassessment of humanity's fundamental nature and role in the universe. After Turing, computer science has not only provided unprecedented epistemic and engineering powers over natural and artificial realities; it has also cast new light on who we are and how we are related to the world. Today, we are slowly accepting the idea that we are not standalone and unique entities, but rather informationally embodied organisms (inforgs), mutually connected and embedded in an informational environment, the infosphere, which we share with both natural and artificial agents similar to us in many respects.

1.6 The evolution of inforgs

We have seen that we are probably the last generation to experience a clear difference between onlife and online. A further transformation worth highlighting concerns precisely the emergence of artificial and hybrid (multi)agents, i.e., partly artificial and partly human (consider, for example, a family as