

Why Information Matters

Luciano Floridi

When we use a computer, its performance seems to degrade progressively. This is not a mere impression. Over the years of owning a particular machine, it will get sluggish. Sometimes this slowdown is caused by hardware faults, but more often the culprit is software: programs get more complicated, as more features are added and as old bugs are patched (or not), and greater demands are placed on resources by new programs running in the background. After a while, even rebooting the computer does not restore performance, and the only solution is to upgrade to a new machine.

Philosophy can be a bit like a computer getting creakier. It starts well, dealing with significant and serious issues that matter to anyone. Yet, in time, it can get bloated and bogged down and slow. Philosophy begins to care less about *philosophical* questions than about *philosophers'* questions, which then consume increasing amounts of intellectual attention. The problem with philosophers' questions is not that they are impenetrable to outsiders—although they often are, like any internal game—but that whatever the answers turn out to be, assuming there are any, they do not matter, because nobody besides philosophers could care about the questions in the first place.

This is an old problem. In the sixteenth century, the French scholar and doctor François Rabelais satirized scholastic philosophy in his *Gargantua and Pantagruel*. In a catalogue of 139 invented book titles that he attributes to the library of the Abbey of St. Victor, he lists such titles as “The Niddy-noddy of the Satchel-loaded Seekers, by Friar Blindfastatis” and “The Raver and idle Talker in cases of Conscience.”

Centuries later, we seem to be back to the same problem. This is how philosophy speaks today: “The Failure of Class: Postcapitalist narrative and textual precapitalist theory” and “Deconstructing Lyotard: Cultural narrative and premodern dedeconstructivism.” Or: “As Lewis taught us in a classic series of articles, trope theories Gettierise zombie arguments”

Luciano Floridi is professor of philosophy and ethics of information at the University of Oxford, and chair of the Data Ethics Group of The Alan Turing Institute, the British national institute for data science. He is the author, most recently, of The Fourth Revolution: How the Infosphere Is Reshaping Human Reality (Oxford, 2014).

WINTER 2017 ~ 7

and “While the contextualist disagrees, we still hold that supposed mind/body ‘problems’ cannot generate an unacceptably Russellian picture of the world.”

Do not try to understand these lines. I produced the first two using a “Postmodernism Generator,” and the second two using an “Analytic Philosophy Generator.” They sound like real examples of contemporary scholasticism—philosophy talking about itself to itself in its own jargon. Such scholasticism is the ultimate freezing of the system, the equivalent of a Windows computer’s “blue screen of death”: so many resources are devoted to internal issues that no external input can be processed anymore, and the system stops working. The world may be undergoing a revolution, Rome may be burning, but the philosophical discourse remains detached, meaningless, and utterly oblivious. Time for an upgrade.

A New Focus for Philosophy

How should philosophy be upgraded so that it can address real-life issues again? In a talk in April 2010, Bill Gates asked “whether the brightest minds are working on the most important problems.” The problems he was thinking of were some of those affecting the world’s poorest people, such as the lack of adequate health care and nutrition, and dangerous and unsanitary environmental conditions, but he was also thinking of the need for access to good education and for energy production. Surely, the list would have to include also the need for peace in war-torn countries, for human rights, and much else besides.

But if these are today’s most pressing issues, what should the brightest *philosophical* minds be doing? Maybe they should stop philosophizing and start doing something about this messy world instead. We could shut down our philosophy departments and put an end to philosophers’ corruption of the youth. But this solution smacks of self-defeat. It would be like burning the wicker basket we are traveling in because our hot-air balloon is descending too quickly. Philosophy is the sort of thing you need to keep in a good world, not what you should get rid of in a bad one. Athens is a better place with Socrates. The fact is that philosophy can be extremely helpful, for it can forge and refine the intellectual tools we need for dealing with the most challenging problems that confront us. The brightest philosophical minds can contribute insights and visions, analyses and syntheses, theories and critiques, questions and answers that can help us to solve these problems. If this still sounds rather useless compared to

actually doing something, consider that philosophy takes care of the roots, so that the rest of the plant might grow more healthily.

But even if we accept this role for philosophy, which ideas, theories, and perspectives should philosophers be designing now and for the foreseeable future, so that their contributions will be timely and helpful? Which philosophical questions should they be addressing?

The answers will have to do with this particular time in human history. Philosophical “upgrading” moments are rare, and they are usually prompted by important transformations in the surrounding reality. Since the Nineties, I have been arguing that we have reached one of those moments—a turning point in our history. The epochal transition from an analogue to a digital world and the rapid development of information technologies are changing every aspect of our lives: education, work, and entertainment; communication, business, and commerce; love, hate, and anything in between; politics, conflicts, and peace; culture, health, and even how we remember the dead. All this and more is being relentlessly transformed by technologies that have the recording, transmission, and processing of information as their core functions.

As information technologies come to affect all areas of life, they are becoming implicated in our most important problems—their causes, effects, and solutions, the scientific investigations aimed at explaining them, the concepts created to understand them, the means of discussing them, and even, as in the case of Bill Gates, the wealth required to tackle them.

Furthermore, information technologies don’t just modify how we *act* in the world; they also profoundly affect how we *understand* the world, how we relate to it, how we see ourselves, how we interact with each other, and how our hopes for a better future are shaped. All these are old philosophical issues, of course, but we must now consider them anew, with the concept of information as a central concern.

This means that if philosophers are to help enable humanity to make sense of our world and to improve it responsibly, information needs to be a significant field of philosophical study. Among our mundane and technical concepts, information is currently not only one of the most important and widely used, but also one of the least understood. We need a philosophy of information.

How to Ask a Question

In the fall of 1999, NASA lost radio contact with its Mars Climate Orbiter, a \$125 million weather satellite that had been launched the year before. In

a maneuver to enter the spacecraft into orbit around Mars, the trajectory had put the spacecraft far closer to Mars than planned, so that it directly entered the planet's atmosphere, where it probably disintegrated. The reason for this unhappy event was that for a particular software file, the Lockheed Martin engineering team had used English (imperial) units of measurement instead of the metric units specified by the agency, whose trajectory modelers assumed the data they were looking at was provided in metric.

This incident illustrates a simple lesson: successful cooperation depends on an agreement between all parties that the information being exchanged is fixed at a specified level. Wrongly assuming that everyone will follow the rules that specify the level—for example, that impulse will be expressed not as pound-seconds (the English unit) but as newton-seconds (the metric unit)—can lead to costly mistakes. Even though this principle may seem obvious, it is one of the most valuable contributions that philosophy can offer to our understanding of information. This is because, as we will see, failing to specify a level at which we ask a given philosophical question can be the reason for deep confusions and useless answers. Another simple example will help to illustrate the problem.

Suppose you want to buy a secondhand car. You can treat the information about its price as an answer to a yes-or-no question: “Is the price of this car 5,000?” “Yes.” You see immediately that there is a problem, not with the answer, but with the question: it contains no indication of the type of currency. The correct way of putting the question to which “Yes” is a meaningful answer is: “Is the price of this car \$5,000?” We have just introduced what computer scientists call the correct *level of abstraction*: the variable represented by the symbol for U.S. dollar, not, for example, by the symbol for euro, is the level at which you are considering the 5,000. Using a metric or an English system of measurement is another choice of level of abstraction. (The reason computer scientists speak of levels and of abstraction is that an information-based system, such as a computer, can be understood as a hierarchy in which higher levels build on lower levels, and the mode of interaction with the system, for example the computer language that a programmer uses, depends on the level.)

To understand what philosophy of information needs to learn from this method of computer science, we can look at an example from Alan Turing, the great pioneer of computer science. He showed how philosophical and conceptual questions too could be answered only by fixing the level at which it would make sense to receive an answer. This is one of the most lasting lessons of his famous imitation game, also known as the

Turing Test of artificial intelligence, which he described in a 1950 paper. Here is a simplified version. Turing was interested in understanding the difference between human and mechanical computation. But he rightly refused even to try to provide an answer to the question as it is often asked, “Can machines think?,” because he considered it a problem “too meaningless to deserve discussion.” If you recall, in our simple example this would be like asking the price of the secondhand car in absolute figures, insisting that no currency is used to express it. Nonsense. Likewise, Turing objected that the question “Can machines think?” involved vague concepts such as “machine” and “think.” Although we could try to find out what these terms mean in common discourse through a large survey, this would be an absurd way of trying to answer the question. In other words, the question lacked a clear level of abstraction.

So Turing suggested replacing the question with the imitation game, which fixes certain variables in a rules-based scenario that is easily implementable and controllable. Suppose that A and B are a human being and a computer, but you do not know which is which. You can ask A and B any question simultaneously, but they are in another place and you can only interact with them by e-mail (or, in Turing’s day, by teleprinter). If after a reasonable amount of time you cannot tell which is the human and which the computer, then the computer has passed the test—that is, the computer is at least as good as the human in providing answers to the questions you asked. Turing’s test is based on a weaker version of Leibniz’s law of the identity of indiscernibles: if, everything else being equal, significant differences between A’s and B’s answers are indiscernible, then A and B are interchangeable. Given the same input of questions, the output of answers a human and a computer can generate are such that the differences between the two are insufficient for the purposes of unmistakable recognition.

Turing thought that computers would pass his test by the year 2000. He was wrong. But while most discussions of the test focus on his mistaken prediction, or on what consequences one should draw if computers were able to pass the test, the method of the test was in fact far more important than the prediction. By suggesting the imitation game, Turing specified a level of abstraction for asking a complex question about the capability of computers: the “currency” he chose for the game was human intelligence, but it could have been something else, from animal intelligence to human creativity, as many other versions of the game have shown. By specifying the level—human intelligence as measured by the imitation game—Turing was able to replace his original and vague

question with a new and answerable one, which may be summed up thus: “May one conclude that a machine is thinking, at the level of abstraction represented by the imitation game?”

After more than half a century, we are still learning this crucial lesson of fixing questions at certain levels. Take, for example, pancomputationalism—the idea that every physical system, and the entire universe, is a computational system. A 2012 *New Scientist* article claimed that “the universe is a computer, and everything that goes on in it can be explained in terms of information processing.” The physicist and computer scientist Stephen Wolfram has argued for a “new kind of science” that would study the computational nature of “a vast range of systems, from simple programs to brains to our whole universe.” And the computer scientist Gregory Chaitin writes that the universe “is constantly computing its future state from its current state” and that “actual computers like your PC just hitch a ride on this universal computation!”

This is an intriguing metaphor, but it should not be taken seriously. Apart from a large number of scientific difficulties, including the endorsement of a controversial form of determinism, the view fails to apply Turing’s method of levels of abstraction. To understand why, consider that informational concepts are so extensive that, given the right level, anything can be presented as a computational system, from a building to a volcano, from a forest to a dinner, from a brain to a company. Likewise, any process can be represented computationally: digesting, flying, and knitting all become forms of input-process-output according to some program.

So the point is to ask not whether there is a level of abstraction at which the whole universe can be viewed as computational—because it is trivially true that there is—but rather whether that is the *right* level at which to analyze the universe without blurring all differences, like a night in which all cows are black, to borrow a saying from Hegel. Also, since it is possible to think in some vague sense of all systems as computational, pancomputationalism lacks an important feature of any good theory: a counterexample must be possible in principle. But pancomputationalism is simply not vulnerable to a refutation. By failing to look at computation at a specified level, all things become computational in some way. It’s like asking if machines can think, or if the car costs 5,000.

What philosophy can offer to contemporary debates that involve the concept of information, whether we discuss the intelligence of computers or the makeup of the universe, is clarity about how to ask the right questions so that answers are possible and useful. Failing to ask the right questions can only lead to confusions and misunderstandings.

Know Thyself Anew

At the risk of oversimplifying, 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 in the modern era have had deep impacts of both kinds. In changing our understanding of the world around us and how we can interact with it, they also modified our conception of who we are and may expect to become—a story that is the subject of Friedel Weinert's book *Copernicus, Darwin, and Freud: Revolutions in the History and Philosophy of Science*.

First, after Copernicus, the heliocentric cosmology displaced the Earth and hence humanity from the center of the universe. Second, Darwin showed that all species of life have evolved over time from common ancestors through natural selection, thus displacing humanity from the pinnacle of rational design. And third, following Freud, we acknowledge that the mind is more than pure rationality, that it is influenced by hidden processes and subject to forces of which we are unconscious. If we are reluctant to consider psychoanalysis a strictly scientific enterprise like astronomy or evolutionary theory, we might yet be willing to concede that contemporary neuroscience plays a revolutionary role in overturning the idea that we are transparent to ourselves.

There may be reasons to question this picture of three revolutions. After all, Freud, who himself presented this exact historical timeline in "A Difficulty in the Path of Psycho-Analysis," interpreted the three moments as part of an ongoing process of reassessment of human nature. Each of the three, he wrote, is a blow against "the universal narcissism of men, their self-love." This way of looking at it was obviously self-serving, as it helped to support his suggestion that his own blow in the realm of psychology was the most severe of the three, and therefore met so much resistance. But while attributing such significance to his work may have been narcissistic in its own way, in retrospect it does not seem so unreasonable. Similarly, when we now perceive that something very significant and profound has happened to human life and self-perception after the computer revolution, our intuition may again be right, because we are experiencing what may be described as a fourth revolution in the process of dislocation and reassessment of humanity's nature and role in the universe.

Computer science and its technological applications have exercised both an extrovert and an introvert influence. They have provided unprecedented powers over natural and artificial realities; and by doing

so, they have 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 at the center of the growing “infosphere” that surrounds us because we are not the only smart agents able to carry out complex tasks. Our computers are often better than we are at dealing with information. So when we define artificial intelligence as the successful performance by digital technologies of tasks that, if left to humans, would require intelligence, we are really telling a new story about ourselves. Add to this the fact that we see ourselves increasingly as informationally embodied organisms (what I have called “inforgs”)—think for example of the idea of DNA as software—mutually connected and embedded in an informational environment that we share with both natural and artificial agents similar to us in many respects.

If we had to choose a representative figure of this latest revolution, it would undoubtedly be Turing. He did more than anyone to lay the conceptual foundations for computer science. And by helping us to think about computation, cognition, intelligence, and how we might see ourselves in relation to computers, Turing has changed our philosophical anthropology as much as Copernicus, Darwin, and Freud. This has had a significant impact on what it means to do philosophy after Turing, as philosophers now face the task of how best to understand ourselves in this new era.

Admittedly, it would be too much of a stretch to attribute to Turing the foundation or even the beginning of a new philosophy of information. After all, he did not focus on the concept of information itself, or on problems about communication understood as information flow or transmission. The index of a large 2004 collection of his writings, *The Essential Turing*, does not even contain an entry for “information,” and David Luenberger’s 2006 book *Information Science* mentions Turing only in relation to his work as cryptographer during the Second World War and his short career thereafter.

And yet, without Turing, his groundbreaking work on information processing, its scientific and technological consequences, and his efforts to devise useful concepts for thinking about them, contemporary interest in the philosophy of information would be very hard to explain. The fact that we are today more likely to treat computers as communication machines rather than powerful calculators, and smart phones as mini-computers rather than telephones indicates how deep the influence of his work has been on our world.

The Time Is Right

Information is, in a way, the Cinderella in the history of philosophy. Any philosophy of knowledge, no matter whether ordinary (epistemology) or scientific (philosophy of science) requires an understanding of information—for instance in discussions of sensory perception and knowledge acquisition. There is no ethics without choices, responsibilities, and moral evaluations, all of which need a lot of relevant and reliable information and quite a good management of it. Logic was first a matter of the study of arguments, and then of mathematical proofs, but today it is also if not mainly a question of information extraction, transmission, and dynamics, and some branches of logic are really branches of information theory. Ontology, the study of the nature of being, would be meaningless without informational patterns—real, virtual, necessary, possible, or even impossible. The philosophy of mind needs informational mental states, and the philosophy of language without communication of information is pointless. Any philosophy of the *logos* is a philosophy of information, and Christian philosophy of religion is inconceivable without the informational concept of revelation. The list could be extended and refined to aesthetics, hermeneutics, philosophy of biology, philosophy of physics, and so forth, but the point is clear. To paraphrase Molière, Western philosophy has been speaking informationally without knowing it for twenty-five centuries. We have always relied on Cinderella working hard in the house of philosophy. It is time to acknowledge her great services, by designing the philosophy of our time to be properly conceptualized *for* our time.

While information has been a concept in the background for so long in the history of philosophy, it now also fits neatly in its foreground. Seventeenth-century philosophers redirected their attention from the nature of the knowable object (metaphysics) to the relation between the object and the knowing subject (epistemology). Problems surrounding how we come to know the world then consumed much of modern philosophy. Then, in the twentieth century, philosophers came to reflect primarily on how knowledge is organized—how it is stored and linguistically structured—thus moving from epistemology to philosophy of language and logic. And with the growth of the information society in which billions of people now spend their lives, some philosophers have increasingly focused on the very fabric of knowledge—information and its dynamics, including communication, flows, and processing. As a result, information has arisen as a concept as fundamental and important as being, knowledge,

life, intelligence, meaning, and good and evil—all concepts with which it is interdependent, and so equally worthy of autonomous investigation. It is also a more impoverished concept, in terms of which the others can be expressed and to which they can be related. This is why the philosophy of information may explain and guide the purposeful construction of our intellectual environment, and provide the systematic treatment of the conceptual foundations of contemporary society.

The future of philosophy of information depends on how well we engage with Turing's intellectual legacy, with today's most important problems, and with classic philosophical issues. I am optimistic. Philosophy of information can help us to expand the frontiers of our understanding by providing innovative methodologies for addressing problems from a contemporary perspective. Relying on Turing's intuition of the crucial importance of the method of abstraction, it ensures that such problems are addressed in the right way. Thanks also to Turing, the Baconian-Galilean project of reading and manipulating the alphabet of the universe has begun to be fulfilled in the computational and informational revolution, which has produced the tools that can describe and modify our environment and ourselves. From this perspective, the philosophy of information can be presented as the study of the informational activities that make possible the construction, conceptualization, and also the moral stewardship of reality, both natural and artificial, both physical and anthropological. It promises to be one of the most exciting and beneficial areas of philosophical research in our time.