

Preface

The ultimate aim of this volume is to further the philosophical reflection on technology within the context of Luciano Floridi's philosophy of technology. Philosophical reflection on technology is as old as philosophy itself, dating back to the Ancient Greek philosophers. The themes that have dominated the philosophical discourse on technology since then can be roughly categorized into three: (i) the social, cultural, and political impacts of technological developments; (ii) the epistemological status of technological knowledge, especially in relation to scientific knowledge; and (iii) the ontological status of the products of technology, i.e., technological artifacts. Luciano Floridi's philosophy of technology, which is based on his philosophy of information, has something to say about each of these themes. Not only that, his philosophical analysis of new technologies leads to a novel metaphysical framework in which our understanding of the ultimate nature of reality shifts from a materialist one to an informational one, in which all entities, be they natural or artificial, are analyzed as informational entities (Floridi 2010). This is the main rationale behind the choosing of his philosophy of technology as the topic of this volume.

There is no doubt that the information and communication technologies of the twentieth century have had a significant impact on our daily lives. They have brought new opportunities as well as new challenges for human development. According to Floridi, however, this is not the whole story. He claims that these new technologies have led to a revolutionary shift in our understanding of humanity's nature and its role in the universe. By referring to an earlier categorization, he calls this the "fourth revolution." The Copernican revolution was the first, leading to the understanding that we as humans are not at the center of the universe. The second revolution was the Darwinian realization that we are not unnaturally distinct or different from the rest of the animal world. The third was the Freudian revolution, which taught us that we are not as transparent to ourselves as we once thought. With the fourth revolution, says Floridi, "we are now slowly accepting the idea that we might be informational organisms among many agents ..., *inforgs* not so dramatically different from clever, engineered artefacts, but sharing with them a global environment that is ultimately made of information, the infosphere. The information revolution [the fourth revolution] is not about extending ourselves, but about re-interpreting who we are" (Floridi 2008a).

This radical claim forms the basis of Floridi's philosophy of technology. Given this basis, philosophical reflection on technology is not only valuable in and of itself, but also brings a completely new framework of analysis for philosophy. In other words, philosophical reflection on technology takes a central role in philosophical analysis. To give an example, Floridi's analysis of object-oriented programming methodology (Floridi 2002), which relies on a method borrowed from a branch of theoretical computer science called Formal Methods, paves the way for defining a new macroethical theory, i.e., Information Ethics. The method he borrows from Formal Methods is the method of levels of abstraction. By using this method, Floridi claims that the moral value of human actions is not different in kind than the moral evaluation of other informational objects. The idea behind the method of levels of abstraction is quite simple and straightforward: the reality can be viewed from different levels. The roots of this simple idea go back to Eddington's work in the early decades of the twentieth century (Eddington 1928). Let me give a brief example in Floridi's own words:

Suppose, for example, that we interpret p as Mary (p =Mary). Depending on the LoA and the corresponding set of observables, p =Mary can be analyzed as the unique individual person called Mary, as a woman, as a human being, as an animal, as a form of life, as a physical body, and so forth. The higher the LoA, the more impoverished is the set of observables, and the more extended is the scope of the analysis (Floridi 2002).

Perhaps the most crucial feature of the method of levels of abstraction is that the identification relation between two variables (or observables) is never absolute. Rather, the identification is always contextual and the context is a function of the level of abstraction chosen for the required analysis (Floridi and Sanders 2004a).

Floridi utilized his method not only in Information Ethics but also in several other subfields of philosophy. The following quote from his *Minds and Machines* article (2008), in which he responded to some objections raised against the method of levels of abstraction, provides a list of the areas in which the method has been used.

Jeff Sanders and I were forced to develop the method of abstraction when we encountered the problem of defining the nature of agents (natural, human, and artificial) in Floridi and Sanders (2004b). Since then, we have been applying it to some long-standing philosophical problems in different areas. I have used it in computer ethics, to argue in favour of the minimal intrinsic value of informational objects (Floridi 2003); in epistemology, to prove that the Gettier problem is not solvable (Floridi 2004c); in the philosophy of mind, to show how an agent provided with a mind may know that she has one and hence answer Dretske's question "how do you know you are not a zombie?" (Floridi 2005a); in the philosophy of science, to propose and defend an informational approach to structural realism that reconciles forms of ontological and epistemological structural realism (Floridi 2004b); and in the philosophy of AI, to provide a new model of telepresence (Floridi 2005b). In each case, the method of abstraction has been shown to provide a flexible and fruitful approach (Floridi 2008c).

The jury is still out as to the truth value of the claim stated in the last sentence of this quote. One thing, however, is certain. Floridi's method borrowed from the Formal Methods branch of theoretical computer science and its applications have led to prolific and novel discussions in many different areas of philosophy. For the purposes of this volume, one of the most important applications of the method is in computer ethics. As mentioned above, Floridi claims that his Information

Ethics is a macroethical theory that provides a foundation for computer ethics. His Information Ethics consists of two main theses: (i) information objects *qua* information objects can be moral agents; and (ii) information objects *qua* information objects can have an intrinsic moral value, although possibly quite minimal, and hence they can be moral patients, subject to some equally minimal degree of moral respect (Floridi 2002).

The contributions in Part I of this volume are mainly centered on Floridi's Information Ethics and the method of levels of abstraction. These are Gordana Dodig-Crnkovic's "[Floridi's Information Ethics as Macro-Ethics and Info-Computational Agent-Based Models](#)," M.J. Wolf, F.S. Grodzinsky, and K.W. Miller's "[Artificial Agents, Cloud Computing, and Quantum Computing: Applying Floridi's Method of Levels of Abstraction](#)," Richard Lucas' "[Levels of Abstraction and Morality](#)," and Federica Russo's "[The *Homo Poieticus* and the Bridge Between *Physis* and *Techne*](#)."

Dodig-Crnkovic's ultimate aim in her chapter is to provide a general framework for the distribution of moral responsibility in multi-agent systems, which include humans as well as technological artifacts. In order to lay the groundwork for achieving this aim, she starts by providing her own interpretation of Floridi's Information Ethics, which she has been developing since 2006. Her interpretation, called the Info-Computationalist interpretation, is characterized by a recursive self-sustaining loop in which "the bottom-up construction of informational structures gives rise to top-down information re-structuring." In other words, the aggregate of the bottom-level elements forms a collective state that has emergent properties that are not reducible to the properties of the bottom-level informational structures. These emergent properties in turn influence the behavior of all bottom-level structures. Dodig-Crnkovic's interpretation is, to say the least, a novel one, because it allows a structured interaction between different levels of abstraction. In addition to her novel interpretation, she also states the similarities between Floridi's Information Ethics and the pragmatic approach to moral responsibility. The classical analysis of moral responsibility requires an agent with free will, and thus limits the domain of moral responsibility only to humans. In contrast, in the pragmatic approach, moral responsibility is not a result of an individual's duty; rather, it is a role defined by the externalist pragmatist norms of a group. Dodig-Crnkovic claims that Floridi's Information Ethics falls under the category of the pragmatic approach, and in that respect has the potential of providing the foundation for a moral framework in which technological artifacts can be assigned moral responsibility. Armed with these two preliminary explanations, i.e., the Info-Computationalist interpretation and the pragmatic character of Information Ethics, she uses Information Ethics to construct an artificial morality framework in which moral responsibility in intelligent systems is distributed across all agents, including technological artifacts. In her artificial morality framework, moral responsibility is handled as a regulatory mechanism that assures the desirable future behavior of intelligent systems.

Wolf et al.'s chapter, in a sense, is a continuation of an earlier article of theirs that appeared in *Ethics and Information Technology* (2009). In that article, they use two different levels of abstraction for analyzing the ethics of designing artificial agents. Their first level of abstraction, LoAI, is the user's view of an "autonomous system"

such as a software package. The second level is the designer's perception of the system. Their ultimate conclusion in that paper is that the ethical responsibilities of a software designer significantly increase with the development of artificial agents because of the more intricate relationship between LoA1 and LoA2. In their contribution to this volume, they extend their original analysis by introducing a third level of abstraction, LoAS, the level that refers to society's perspective. This is important because new artificial agents not only have effects on individuals but also on the whole society that comprises those individuals. With this new addition, they test the applicability of Floridi's Information Ethics and the method of levels of abstraction to two new computing paradigms: cloud computing and quantum computing. Their overall conclusion is a positive one. They claim that although there are new challenges for Information Ethics in these two computational paradigms, Information Ethics has the potential of successfully meeting those challenges. It should be noted that their chapter also provides a nice and brief overview of the fundamental concepts of quantum computing.

Lucas' chapter is an extensive and detailed criticism of Information Ethics. He criticizes three notions that form the fundamentals of Floridi's theory, which are interactivity, autonomy, and adaptability. Lucas' ultimate conclusion is that Information Ethics, mainly because of being only formally defined, is too artificial and too simple for a natural characterization of morality. Although Floridi thinks that Lucas' understanding of Information Ethics is based on serious misunderstandings and that Lucas' chapter is beyond repair (please see Floridi's reply at the end of this volume), the chapter paves the way for a closer scrutiny of some of the arguments that Floridi has provided in defense of Information Ethics. An example might be helpful at this point. The essential motivation of Information Ethics is to be able to count artificial agents as moral agents. It should be noted that this essential motivation is somewhat different than the motivation behind the earlier characterizations of computer and information ethics. Moor (1985) is a good example of the classic treatment of the subject. In one of their earlier characterizations of Information Ethics, Floridi and Sanders consider a set of possible objections to their main claim about the moral value of artificial agents. These are the teleological objection, the intentional objection, the freedom objection, and the responsibility objection. They then provide counterarguments against those objections. Lucas thinks that none of these counterarguments sufficiently overcome the four possible objections that Floridi and Sanders consider. Of course, whether Lucas is right in his assessment or not is a matter of debate, but Lucas' reasoning urges us to reevaluate the fundamental arguments provided for the philosophical value of Information Ethics. In that respect, it is a valuable contribution to this volume.

Russo, in her chapter, focuses on one particular aspect of Floridi's Information Ethics, the reconciliation of *physis* and *techne* in a constructionist manner. According to Floridi, traditional macroethical theories take the situation which is bound to moral evaluation as given, but this traditional approach ignores the poietic nature of humans as ethical agents. Ignoring the poietic nature of humans is the ultimate basis of the dichotomy between *physis* and *techne* (Floridi and Sanders 2003). The demarcation line between these two has been disappearing because of digital

technologies. Russo agrees with Floridi's analysis and attempts to take the analysis one step further. For Russo, the gradual disappearance of the demarcation line between *physis* and *techne* is not just a result of the new digital technologies; rather, it is dominated by new technologies in general. These new technologies include biotechnology and nanotechnology, which allow us to be "creating altogether new environments that pose new challenges for the understanding of us in the world." Floridi's Information Ethics, according to Russo, successfully accounts for the ethical implications of these new technologies, but, she continues, the epistemological implications are also at least equally important and need to be analyzed. This is what she aims to achieve in her chapter. In that respect, it would not be wrong to say that Russo takes Floridi's original analysis of digital technologies and applies it to a wider domain.

The two chapters in Part II provide novel ways of categorizing scientific and technological advancements on the basis of metrics different than Floridi's metric, which is based on introverted effects of scientific changes on the way we understand human nature. These are Anthony F. Beavers' "[In the Beginning Was the Word and Then Four Revolutions in the History of Information](#)" and Valeria Giardino's "[I Mean It! \(And I Cannot Help It\): Cognition and \(Semantic\) Information.](#)"

Beavers, in his chapter, gives us a different categorization of the technological revolutions that mankind has experienced in its entire history. As mentioned above, Floridi's categorization of the information revolution as the fourth revolution is based on the metric of the way scientific developments change our understanding of ourselves. Thus, according to this metric, scientific developments that have led to a reassessment of humanity's fundamental nature and role in the universe are counted as revolutionary. Of course, as Floridi himself states, other metrics are also possible. In his chapter, Beavers offers a different metric that is not supposed to be an alternative to Floridi's metric, but rather complementary. The suggested metric is the history of information flow itself. In other words, technological and scientific advancements are categorized according to "the kind of information that can be stored and transmitted, the speed of information transmission, its preservation, and its reach." This metric also gives us four revolutions: the Epigraphic Revolution, the Printing Revolution, the Multimedia Revolution, and the Digital Revolution. The last one, which corresponds to Floridi's fourth revolution, is characterized by the introduction of automated information processing. There are two interesting features of Beavers' categorization that I would like to mention in this short preface. The first is that in his categorization, the Digital Revolution is not considered as a discontinuity from the previous revolutions, because information transmission and coding were also present, albeit in different forms, in the previous revolutions. What the Digital Revolution has brought to the table is new and revolutionary technological affordances that are made possible by automated information processing. This interesting feature, perhaps, is what fundamentally differentiates Beavers' categorization from Floridi's categorization. The second point is that the trajectory of the history of information flow is not characterized merely by the evolution of particular technologies, but also by the evolution of the informational networks that those particular technologies enable. After establishing his new

categorization, Beavers situates the role of Philosophy, in particular the role of Philosophy of Information, in the historical context of the categorization by providing both valuable historical insights for the evolution of philosophical analysis and crucial questions that will help in the advancement of the Philosophy of Information as a new *philosophia prima*.

Giardino, in her chapter, also provides a different categorization of technological revolutions. Giardino argues that what Floridi calls the fourth revolution is in fact the second information revolution. The underlying reason for this difference is her analysis of information from a cognitive perspective. She thinks that we have been living in an informational environment all along, and that the infosphere includes all informational cognitive agents and cognitive tools. Given this understanding, any artifact that aids symbolic activities becomes an informational cognitive tool. Humans have been living in an informational environment since the time of the invention of the first tool that aided symbolic thinking. For Giardino, the correct characterization of information revolution(s) should be based on how information is transmitted across generations. Thus, the first information revolution is characterized by the transformation of the transmission from sequences of DNA to cultural transmission. The second information revolution, i.e., Floridi's fourth revolution, is characterized by the switch from cultural transmission to online transmission, according to Giardino. One of the valuable features of this chapter, among many others, is its interdisciplinary character. Giardino nicely brings together the literature on Philosophy of Information with the literature on Developmental Psychology and the literature on Cognitive Science.

The contributions in Part III take Floridi's Philosophy of Technology and Philosophy of Information as their basis and apply them to different domains: Elena Pasquinelli's "[What Happens to Infoteachers and Infostudents After the Information Turn?](#)" to education, Raphael Cohen-Almagor's "[Content Net Neutrality: A Critic](#)" to the regulation of freedom of speech on the Internet, and Armando Malheiro da Silva and Fernanda Ribeiro's "[Information Science and Philosophy of Information: Approaches and Differences](#)" to Information Science.

With the changes brought about by the information revolution, we humans have become inforgs that live in the infosphere, according to Floridi. The information revolution has led to a reontologization of our ordinary environment, where the divide between online and off-line has been disappearing. Our environment, the infosphere, "will become increasingly synchronized (time), delocalised (space) and correlated (interactions)," says Floridi. Pasquinelli, in her chapter, in light of Floridi's description of the infosphere, analyzes the past and possible future effects of the information revolution on educational institutions, practices, and actors. She starts her chapter with a diagnosis: the information revolution has not yet revolutionized education. For her, the main reason for this is the reluctance of educational institutions and actors in adopting the new tools, approaches, and paradigms that are made possible by information and computational technologies, especially in comparison to the institutions and actors of other domains. She then compares two different ways of changing the educational institutions and practices. The first is the top-down approach mostly adopted by policy makers. She cites the "One Laptop per Child"

(OLPC) program as an example of the top-down approach and shows the difficulties involved in changing educational practices in this way. According to Pasquinelli, change from the top is difficult, mainly because of the sheer size of educational institutions and the long tradition of educational paradigms and practices. A second reason, which is clearly seen in the OLPC case, is that top-down changes usually do not include students, who are the ultimate users of education, in the design of changing programs. Then she proceeds to give an example of a bottom-up approach that she claims to be more promising. Her fascinating example is the experience of Math on MXit from South Africa. With this example, she urges educational institutions and actors to implement the new technologies from the bottom up. The ultimate goal of such changes, for her, is to turn students into infostudents and teachers into infoteachers. During this transformation, which will be slow and gradual, she says, the old paradigms of education will be challenged because of the new tools and approaches of the information revolution. As the dominant example of the old educational paradigms, she gives the Victorian school, which was defined by the following three characteristics: (i) a dedicated and separated space for learning, (ii) a dedicated time for learning, and (iii) well-defined roles for the learner and the teacher. With the Internet, mobile phones, and digital media, she says, learning could occur anywhere and anytime. Moreover, the demarcation line between the student and the teacher will be blurred to the point of disappearance. In short, Pasquinelli's chapter is an informative and fascinating one in which she urges us to reontologize and reconceptualize our environment for education.

Cohen-Almagor, in his chapter, uses Floridi's Information Ethics in order to identify the democratic regulative principles of freedom of speech on the Internet and the responsibility of Internet Service Providers and Web Hosting Services. He starts his analysis by distinguishing three different senses of "net neutrality": (i) net neutrality as a nonexclusionary business practice; (ii) net neutrality as an engineering principle, allowing traffic on the Internet in a nondiscriminatory manner; and (iii) net neutrality as content nondiscrimination. He calls the third sense Content Net Neutrality. Although he accepts the first two senses as the fundamental principles that should underlie Internet regulation, he rejects Content Net Neutrality. Following Floridi's proactive approach to Ethics, which states that the ethical obligation in the information age is not limited to ethical behaviors in the infosphere but needs to extend to actively shaping the infosphere for the betterment of the humanity, Cohen-Almagor urges us to regulate the available content on the Internet. He argues that content that is morally repugnant and/or at odds with democratic ideals should not be made available on the Internet, and that the primary responsibility for this lies with Internet Service Providers and Web Hosting Services. Throughout his discussion, he uses several striking examples that seem to support his position.

As Silva and Ribeiro point out, Information Science as an autonomous field of study that appeared in the late 1950s. Since then, this new field of inquiry, which could be seen as a continuation of the library sciences, has seen an immense and rapid growth. Despite this rapid growth, however, its nature has not yet been precisely defined. This is perhaps due to the inherently interdisciplinary character of the field. Most interdisciplinary fields, for example Cognitive Science, have gone

through a similar stage of development. Silva and Ribeiro, in their chapter, provide an all-encompassing framework for the nature and identity of Information Science. In their framework, Information Science is “a unitary yet transdisciplinary field of knowledge, included in the overarching area of the human and social sciences, which gives theoretical support to some applied disciplines such as Librarianship, Archivistics, Documentation and some aspects of Technological Information Systems.” After providing their framework, they turn to Floridi’s Philosophy of Information with the aim of finding a firm philosophical grounding for Information Science. While doing that, they state their own definition of information, which implies the following properties: structured by an action, integrated dynamical, has potentiality, quantifiable, reproducible, and transmissible. Their definition of information has some differences from Floridi’s definition of semantic information. Perhaps one of the crucial differences is their distinction between informational data and noninformational data. The analysis of the differences and similarities between their definition of information and Floridi’s semantic information is by itself valuable. Moreover, along the way they also bring together different threads of discussions, ranging from the French philosopher Ruyer’s work on visual sensation to Søren Brier’s Cybersemiotics. Given their analysis of Information Science and the connections they identify between Information Science and Philosophy of Information, it is plausible to conclude that Information Science could be understood as applied Philosophy of Information.

The main focus in Part IV is the epistemic and ontic aspects of Floridi’s Philosophy of Information. The contributions here are Eric T. Kerr and Duncan Pritchard’s “Skepticism and Information,” Joseph E. Brenner’s “Levels of Abstraction; Levels of Reality,” and Steve T. McKinlay’s “The Floridian Notion of the Information Object.”

It is almost a truism to say that information should be “adequately created, processed, managed and used” (Floridi 2010). The bombardment of information that we all face in this day and age requires proper information management. As rightly pointed out by Kerr and Pritchard, proper information management requires paying attention to the connection between information and knowledge. After all, information is valuable as long as it paves the way for the acquiring of knowledge. In their chapter, Kerr and Pritchard focus on this important issue, i.e., the epistemic value of information. One of the milestones in the literature on the epistemic value of information is Dretske’s book *Knowledge and the Flow Information*, in which a comprehensive epistemology based on information is provided. One of the controversial features of Dretske’s framework is its denial of the principle of epistemic closure, which simply states that if an agent knows a proposition and knows that the proposition in question implies another one, then the agent also knows the implied proposition. Dretske’s main reason behind the denial of closure is that, for him, information about appearances can never completely rule out skeptical doubts. Kerr and Pritchard claim that Dretske is wrong and that there are ways in which information could address skeptical doubts. They examine two such ways in their chapter: Ram Neta’s contextual approach and John McDowell’s disjunctivism. Kerr and Pritchard’s chapter is valuable in and of itself. Moreover, it opens doors for a different

approach to the epistemic value of information. Dretske's epistemological analysis is done in a hybrid context of doxastic and informational concepts. Kerr and Pritchard's analysis of the closure principle may also be understood as showing a need for moving to a purely informational context of analysis for knowledge, and this is exactly what Floridi does in his *Philosophy of Information*.

In his chapter, Brenner provides an extensive comparison of his logico-ontological theory, which is called *Logic in Reality*, and Floridi's *Philosophy of Information*. According to Brenner, "the broad theory of information proposed by Floridi requires an understanding of the properties and role of information at all levels of reality, in all entities." In other words, a complete theory of information should clarify the relevant ontological properties of information. Given the Kantian spirit of his theory, however, Floridi is quite cautious in making any ontological commitment about reality and entities. The method of levels of abstraction is proposed as a more inter-subjective, socially constructible (hence possibly conventional), dynamic, and flexible way to further Kant's approach. This method, claims Floridi, needs to be seen as a step away from internal realism, but this does not imply that it is a step toward external realism (Floridi 2008b). Thus, according to Brenner, in its current status Floridi's *Philosophy of Information* seems to be incomplete. Brenner claims that his *Logic in Reality* remedies this problem and complements Floridi's theory, and he discusses this at length in his chapter. To put it briefly, *Logic in Reality* is an extension of logic to complex real processes, providing a framework for analyzing and making inferences about complex real world entities and processes at all levels of reality, including biological, cognitive, and social levels. It is obvious from this nutshell definition that the processes that *Logic in Reality* aims to address include information production and transfer, as well. Some of the philosophically interesting features of Brenner's *Logic in Reality* are as follows. First, the proposed logic is nonpropositional and non-truth-functional. Second, it is grounded in a fundamental dualism, dynamic opposition, that is claimed to be inherent in energy and present in all real phenomena. In other words, real complex phenomena are in a contradictory relation between themselves and with their opposites. Third, the dynamic opposition in energy is accompanied by the law of the included middle, and thus there is no room for the principle of noncontradiction. Fourth, *Logic in Reality* neither requires nor commits to abstract categorical structures that separate different aspects of reality. Thus, most of the absolute distinctions of the traditional philosophical analysis, such as the one between epistemology and ontology, disappear in the framework of *Logic in Reality*. Fifth, *Logic in Reality* is based on a process-ontological view of reality, which means that the ontological inventory of the world is composed of processes at different levels of complex real phenomena. A direct result of this fifth feature is that Brenner's *Logic in Reality* implies an ontological levelism. As clearly stated in his defense of *Informational Structural Realism* (Floridi 2008b), Floridi is committed to the epistemological levelism that his method of levels of abstraction implies, but, as a result of his Kantian general framework, he finds ontological levelism untenable (Floridi 2008c). Brenner states that the ontological levelism that Floridi finds untenable is a result of the misconception of reality as seen through the glasses of classical logic and the traditional object-based

ontological approach. In other words, according to Brenner, any ontological levelism that is based on an absolute distinction between epistemology and ontology is untenable, as Floridi rightly argues, but once the epistemology/ontology of Logic in Reality is adopted, then the ontological levelism becomes tenable and compatible with Floridi's Philosophy of Information.

As the ontological basis of his Philosophy of Technology, Floridi defends a form of structural realism which he calls Informational Structural Realism. In this particular version of structural realism, objects are considered as structural entities which are nothing but a collection of data clusters. This gives rise to Floridi's notion of informational objects as the fundamental ontological entities. As a side remark, it should be noted that although Floridi's analysis of objects in informational terms is quite novel, the history of including information as a fundamental entity in the metaphysics of the world dates back to Wiener's work on *Cybernetics* (Wiener 1948). In order to establish his notion of informational object, Floridi heavily relies on the lessons that he draws from object-oriented programming (OOP), both in terms of the methodology and of the ontology of OOP in constructing his Philosophy of Information and therefore his Philosophy of Technology. McKinlay's chapter focuses on the similarities and differences between OOP and Floridi's Philosophy of Information with respect to their ontology. McKinlay claims that the objects of OOP cannot be the informational objects that Floridi needs in his ontology simply because the objects of OOP are referents, whereas Floridi's informational objects are supposed to be ontologically primitive. McKinlay's claim is almost a direct result of his nominalism about conceptual objects such as OOP classes. His defense of nominalism heavily draws upon Quine's ideas. In addition to its philosophical value in terms of calling our attention to the ontological issues surrounding information and artifacts, McKinlay's chapter also provides a nice introduction to object-oriented programming.

The last chapter of the volume, "[The Road to the Philosophy of Information](#)," is Floridi's reply chapter in which each of the contributions are critically evaluated.

This volume, in my humble opinion, is quite promising in terms of achieving its ultimate aim, which is to further the philosophical reflection on technology.

References

- Eddington, Arthur. 1928. *The nature of the physical world*. Cambridge: Cambridge University Press.
- Floridi, Luciano. 2002. On the intrinsic value of information objects and the infosphere. *Ethics and Information Technology* 4: 287–304.
- Floridi, Luciano. 2003. On the intrinsic value of information objects and the infosphere. *Ethics and Information Technology* 4(4): 287–304.
- Floridi, Luciano. 2004b. The informational approach to structural realism. final draft available as IEG – Research Report 22.11.04. <http://www.wolfson.ox.ac.uk/~floridi/pdf/latmoa.pdf>
- Floridi, Luciano. 2004c. On the logical unsolvability of the gettier problem. *Synthese* 142(1): 61–79.
- Floridi, Luciano, and Jeffrey W. Sanders. 2004a. On the morality of artificial agents. *Minds and Machines* 14(3): 349–379.

- Floridi, Luciano, and Jeffrey W. Sanders. 2004b. On the morality of artificial agents. *Minds and Machines* 14(3): 349–379.
- Floridi, Luciano. 2005a. Consciousness, agents and the knowledge game. *Minds and Machines*. 15(3–4): 415–444.
- Floridi, Luciano. 2005b. Presence: From epistemic failure to successful observability. *Presence: teleoperators and virtual environments* 14(6): 656–667.
- Floridi, Luciano. 2008a. Artificial intelligence's new frontier: Artificial companions and the fourth revolution. *Metaphilosophy* 39(4/5): 652–654.
- Floridi, Luciano. 2008b. A defence of informational structural realism. *Synthese* 161(2): 219–253.
- Floridi, Luciano. 2008c. The method of levels of abstraction. *Minds and Machines* 18: 303–329.
- Floridi, Luciano. 2010. *The philosophy of information*. Oxford: Oxford University Press.
- Floridi, Luciano, and Jeffrey W. Sanders. 2003. Internet ethics: The constructionist values of *Homo poieticus*. In *The impact of the internet on our moral lives*, ed. R. Cavalier, 195–214. New York: SUNY.
- Moor, James. 1985. What is computer ethics? *Metaphilosophy* 16(4): 266–275.
- Wiener, Norbert. 1948. *Cybernetics: or control and communication in the animal and the machine*. New York: Technology Press/Wiley.
- Wolf, M.J., K. Miller, and F.S. Grodzinsky. 2009. On the meaning of free software. *Ethics and Information Technology* 11(4): 279–286.