

This is a preprint version of the following article:
Brey, P. (2008). 'Do We Have Moral Duties Towards Information Objects?' *Ethics and Information Technology*.

Do We Have Moral Duties Towards Information Objects?

Abstract

In this paper, a critique will be developed and an alternative proposed to Luciano Floridi's approach to Information Ethics. Information Ethics is a macroethical theory that is to both serve as a foundation for computer ethics and to guide our overall moral attitude towards the world. The central claims of Information Ethics are that everything that exists can be described as an information object, and that all information objects, qua information objects, have intrinsic value and are therefore deserving of moral respect. In my critique of Information Ethics, I will argue that Floridi has presented no convincing arguments that everything that exists has some minimal amount of intrinsic value. I will argue, however, that his theory could be salvaged in large part if it were modified from a value-based into a respect-based theory, according to which many (but not all) inanimate things in the world deserve moral respect, not because of intrinsic value, but because of their (potential) extrinsic, instrumental or emotional value for persons.

Keywords: Information Ethics; intrinsic value; respect; informational realism; Object-Oriented Programming; bioethics; environmental ethics; artifacts; anthropocentrism

In a series of articles over a number of years, Luciano Floridi and Jeff Sanders have developed an ethical theory they call *Information Ethics* (IE), which is to serve as a philosophical foundation for the field of computer ethics (Floridi 1999, 2002, 2003, 2007a, b; Floridi and Sanders, 2001, 2002). According to Floridi and Sanders, computer ethics, understood as a field of applied ethics concerned with moral issues in the design and use of computer systems and digital technologies, has suffered from a lack of a foundational ethical theory, a *macroethics*, that unifies problems and issues in the field and legitimizes computer ethics as a philosophically important field of ethics. The undesirable alternative, they claim, is that computer ethics is no more than a random collection of applied moral issues that happen to involve computers. Rejecting approaches to computer ethics that deny the need for such foundations, as well as previous attempts at such underpinnings, they present IE as a compelling and viable

macroethical foundation for computer ethics. Floridi has presented the most extensive elaboration and defense of IE, and in this paper I will focus on his conception of IE. I will analyze Floridi's arguments for IE and will assess its prospects as a foundational theory for computer ethics.

Information Ethics: A Brief Overview

Floridi presents IE as a novel ethical theory that has standing independently of its possible application towards computer ethics. Whereas most ethical theories are *agent-oriented*, IE is *patient-oriented* in that it is not directly concerned with the question of how agents should behave, but rather with the question which kinds of things qualify as *moral patients*, which are objects deserving of moral consideration or respect, and how different classes of moral patients should be treated. In defining the class of moral patients for IE, Floridi takes the radical standpoint that everything that exists in the world is a moral patient, meaning that everything that exists is deserving of at least some respect. This moves Floridi beyond both the classical anthropocentric position that the class of moral patients includes only humans, and beyond biocentric and ecocentric positions according to which the class of moral patients consists of living organisms or elements of ecosystems.

In IE, everything that exists, whether a table, a human being, or a speck of dust is described as an *information object*. An information object is an object defined at an informational Level of Abstraction, as consisting of data structures that specify its attributes and its state, and functions and procedures that define how it behaves or reacts to other objects. The notion of Level of Abstraction (LoA) is used by Floridi to capture the fact that reality can be described at different levels of analysis, from abstract to concrete, each with their own ontology. At one LoA, a table is a solid object, whereas at another LoA, it is a collection of atoms in space. Floridi's claim is that there is an informational LoA at which reality can be described in informational terms. A rigorous definition of an information object can be given within the context of Object-Oriented Programming, a programming approach in which information processing takes place

through the interactions of modeled objects which belong to classes, and which behave and interact according to procedures (“methods”) associated with the classes to which they belong. Information entities are, by Floridi’s definition, part of the *infosphere*, which is “the environment constituted by the totality of information entities – including all agents – processes, their proprieties and mutual relations.” (Floridi, 1999, p. 44).

The foundational moral claim of IE is that all information objects, due to their status as information objects, have an *intrinsic moral value*, meaning that they have an inalienable moral worth of their own and are therefore deserving of moral consideration and respect. Floridi adds that this moral worth may be quite minimal, and can be overridden by other moral considerations. Every information object, Floridi claims, has minimal rights to persist and to flourish, that is, to improve and enrich itself. Corresponding to these minimal rights are minimal duties for (human) agents, who ought to respect information objects as ends in themselves. Generalizing from their duties for individual information objects, agents have a responsibility of stewardship towards the infosphere as a whole, to contribute to its continuous growth and flourishing, that is, to decrease its entropy and not to increase it. Floridi proposes a structured set of duties towards the infosphere, including duties not to cause, to prevent and to remove entropy from the infosphere and to promote the flourishing of informational entities and the infosphere as a whole (Floridi, 1999, 2007b).

Floridi holds that IE provides a general ethical framework, including a general set of moral principles, that can be applied towards specific issues and cases in computer ethics. As noted in Floridi (2007b), a variety of such applications has already been undertaken, ranging from the ethics of the digital divide to the ethics of computer games. IE thus presents itself as a radical, unified macroethical foundation for computer ethics and a challenging ethical theory in its own right.

Floridi's Defense of Information Ethics

Environmental philosopher Baird Callicott (1995) has claimed that if we are to start valuing things as intrinsically valuable that we do not already value as such, we need good reasons to do so. Since people do not normally seem to assign intrinsic value to information objects, Floridi needs to provide strong arguments for us to start valuing them as such. An extended argument for IE is given in Floridi (2003). In this essay, Floridi presents both negative arguments against rival theories of IE and positive arguments for IE itself. His negative arguments are aimed against restrictive theories of intrinsic value and moral patienthood that limit them to classes of patients narrower in scope than the class of information objects. His positive argument purports to show that after these alternative theories have been eliminated, there are good reasons for holding that everything, *qua* information object, has intrinsic value and is therefore a moral patient.

Floridi first presents an argument against *anthropocentric* conceptions of intrinsic value, according to which only human beings are moral patients and therefore deserving of respect. Anthropocentric conceptions are standard in Western ethics. Anthropocentrists attribute intrinsic value to human beings and not to animals or inanimate objects because they claim human beings to have certain unique properties, like rationality or free will, by which they and they alone are entitled to respect. Floridi starts by observing that anthropocentrists typically hold that human beings are intrinsically valuable even when they are lacking the relevant uniquely human properties, for instance through birth defects, injury or illness. But if that is the case, he argues, their attributions of intrinsic value must be based on some other property of human beings that is not unique to them but that they share with animals, for instance the capacity for suffering. This shows that, even if properties like free will and rationality provide extra reasons for respect, there are properties in humans not unique to them that justify attributions of intrinsic value, and hence that the class of moral patients must be expanded to include all entities that possess these generic properties.

Biocentrists argue, by a line of reasoning similar to Floridi's, that the class of moral patients with intrinsic value should therefore be expanded to living organisms. Yet, Floridi presents an argument which is intended to both undermine biocentrism and to function as a positive argument for IE. This argument starts with the observation that there are objects for which it has become customary, for at least some people, to attribute intrinsic moral worth to them and to treat them with respect. Such objects include natural objects like large rocks in natural environments and cultural artifacts like objects of cultural heritage. Floridi then proceeds to claim that the minimal condition shared by these and other objects to which we attribute intrinsic value is that they are information objects, and he concludes that therefore all information objects must have at least some intrinsic value.

Criticism

Before I will start critiquing Floridi's argument for IE, let me first indicate some major points of agreement and support. First, I agree with Floridi that it would be a good thing if we could have a macroethical foundation for computer ethics, I applaud their efforts to develop such a theory, and I find their orientation towards a patient-oriented approach centered on the notions of intrinsic value and moral respect intriguing. Second, I hold that Floridi's argument against anthropocentric conceptions of intrinsic value convincingly demonstrates their untenability and shows the necessity of expanding the class of moral patients beyond human beings. And third, I hold that Floridi presents a plausible case that this expansion should not just include biological organisms but at least some inanimate objects as well.

What goes wrong in Floridi's argument for IE is the final step, in which he infers from the fact that both biological organisms and some inanimate objects deserve respect that everything is deserving of respect and therefore intrinsically valuable. This inference is invalid in two respects. First, it wrongly equivocates between deserving respect and possessing intrinsic value, and second, it overgeneralizes from specific classes of entities to all of reality.

The first problem with the argument lies in Floridi's inference from the fact that some objects, like rocks and objects of cultural heritage deserve respect that they therefore have intrinsic value. This inference can only be validly made if it is the case that an entity is deserving of respect only if it has intrinsic value. Elsewhere in his paper, however, Floridi (2003) states, correctly I believe, that objects can also deserve respect for possessing a different kind of value, which he calls *extrinsic value*. Extrinsic value is value possessed by objects because of contingent properties: the role they play or the symbolic meaning they have. A piece of cloth, for example, has extrinsic value for being respected as a flag, but its symbolic meaning as a flag is utterly contingent and may be lost, together with its extrinsic value. According to Floridi, objects with intrinsic value command respect for their inherent, inalienable properties, whereas objects with extrinsic value command respect for contingent properties that they may lose.

I believe that Floridi is right in arguing that objects with extrinsic value deserve respect. But, I want to argue, these are not the only types of objects that deserve respect. As Floridi claims, objects may also possess instrumental value (e.g. a hammer) or emotional value (e.g. a gift from one's lover). Such objects, I want to claim, can also be deserving of respect. This is the case when the value at issue is instrumental or emotional value *for others*. Respecting such objects is then a means for respecting others. We are taught, for instance, to have respect for the personal property of others because the objects in question have an instrumental or emotional value for others, which we want to respect. This, then, is not respect for the sake of these objects themselves but for the sake of their owners.

If objects can command respect for other reasons than their having intrinsic value, the question is what reasons we have to respect inanimate objects like rocks or objects of cultural heritage. Floridi suggests that we respect them because of their intrinsic value as information objects. But it seems more plausible that we respect them for either their extrinsic value or for their instrumental or emotional value to others. Rocks may for instance be respected for their extrinsic value as gifts of God or products of nature, or for

their instrumental or emotional value for present and future generations. Objects of cultural heritage may be respected for their extrinsic value as representatives of a particular school of art or as expressive symbols of national identity, or for their emotional value for humanity as a whole. It must be concluded, therefore, that Floridi has not argued convincingly that because objects like rocks and works of art deserve respect it follows that they have intrinsic value.

But let us suppose that Floridi could come up with a convincing argument that objects like rocks and objects of cultural heritage have intrinsic value. Floridi (2003) would then have us believe that in developing an account of the different things that have intrinsic value there must be a “most general possible common set of attributes which characterises something as intrinsically valuable” (p. 30) and, next, that the “least biased and most fundamental solution is to identify the minimal condition of possibility of an entity’s least intrinsic worth with its nature as an information object” (p. 31). These are bold statements, but no arguments are provided to support them. Most importantly, why should the correct account of intrinsic value be a general, minimalist, homogenous account, rather than a heterogenous and maximalist one, in which biological organisms, rocks, and objects of cultural heritage all have intrinsic value, but for different reasons unique to their own nature? Floridi fails to tell us why.

It can be concluded that Floridi’s main argument in favor of IE is lacking. In addition, a powerful argument can be made *against* IE. This is the argument that IE is committed to an untenable egalitarianism in the valuation of information objects. I will call this the *anti-egalitarian argument*. Within IE, it seems, no difference in value exists between different kinds of information objects: every information object, *qua* information object, is intrinsically valuable and therefore equally deserving of respect. This apparent egalitarianism has the undesirable consequence that, from the point of view of IE, a work of Shakespeare is as valuable as a piece of pulp fiction, and a human being as valuable as a vat of toxic waste. Floridi will no doubt want to reply that differentiation is possible because some objects have additional worth beyond their status as information objects. But note that any such sources of additional worth lie beyond the scope of IE, because IE

only assigns worth to things *qua* information objects. IE tells us that we should be equally protective of human beings and vats of toxic waste, or of any other information object, and that we have an (albeit overridable) duty to contribute to the improvement and flourishing of pieces of lint and human excrement. At best, this suggests that IE gives us very little guidance in making moral choices. At worst, it suggests that IE gives us the wrong kind of guidance.

A final problem with IE lies in its *ontology*. I will argue that to ground IE, Floridi requires an objectivist ontology for information objects, but fails to present a strong case for one. Information objects are either objectively real, meaning that they have mind-independent existence, or they are in some way mind- or observer-dependent. However, if information objects are to possess intrinsic value, they cannot be observer-dependent, because for an object to possess intrinsic value it must possess one or more properties that bestow intrinsic value upon it, such as the property of being rational, being capable of suffering, or being an information object. Such properties have to be objective and inalienable properties of the object in question, not subjective or contingent ones, because otherwise the assigned value is (at best) extrinsic, that is, resulting from the attribution of contingent roles or subjective meanings to objects.

This implies that information objects cannot be defined in terms of ordinary notions of information, because these tend to be receiver- and therefore observer-relative. In his own philosophical theory of information, for example, Floridi distinguishes between two kinds of information, semantic and environmental, both of which he describes as receiver-relative (Floridi, 2005a, b). Semantic information is man-made, predominantly linguistic, information, and clearly not the kind of information on which the notion of an information object can be based. Environmental information, which seems the more relevant notion here, is meaningful data in the environment. It is, for example, the information that there is a fire (because there is smoke) or that the car ran out of gas (because the gas gauge is on empty). Floridi states, however, that such information is observer-relative: it is defined relative to an information agent who receives the information (Floridi 2005b).

That environmental information is receiver-relative does not preclude one from arguing that it is objective in the sense of not being dependent on a single information agent: the rings on a tree trunk may be argued to contain potentially the same information for you, me, or a Martian, and in this sense, environmental information is not subjective but objective. This argument is invalid, however. For an agent to be informed by the rings that the tree is fifty years old, he must be in possession of particular concepts, like that of a tree, of a ring, and of time. These concepts are available meaningful structures that are used to process incoming data. They are human-made and depend on our cognitive and affective makeup and our particular mode of embodiment. There is no reason to believe that a Martian would possess the same or similar concepts, and that therefore to him, the tree trunk would contain the same information as it does for us.

In general, then, an object contains information only relative to the conceptual scheme of a receiver (or information agent). Receivers differ in their mentality and are embodied differently, and therefore employ different conceptual schemes that presuppose different ontologies of the world, resulting in different conceptions of the information that is contained in objects. To save the objectivity of information objects, Floridi would either have to argue that there is a single superior conceptual scheme which abstracts from the contingencies of agents and which can provide the basis for an objective notion of information, or that there is a more primordial notion of information, different from either semantic or environmental information, that has objective existence.

Floridi takes the latter approach by developing the position of *informational realism* (Floridi, 2008, 2004). This position is that the ultimate nature of reality is not material or substantial but informational, consisting of interacting information objects. Floridi argues that the most fundamental relation between two entities is the *relation of difference*, which is a basic, unqualified lack of uniformity between two items. This relation, Floridi claims, is binary and symmetrical. He and then goes on to define a *datum* as a concrete relation of difference. An information object, Floridi claims, is a cluster of such data. Floridi defines relations of difference, and hence data and

information objects, as mind-independent entities, that are moreover unknowable, like Kant's thing-in-itself, because agents can only know interpreted data. Strictly speaking, then, Floridi therefore defines information objects not in terms of *information* but in terms of a metaphysical conception of *data* on which the mind-dependent notions of semantic and environmental information depend.

I am willing to accept that this theory is not incoherent or inconsistent and that it could conceivably be true. But what are Floridi's positive arguments that it is, indeed, true? Floridi concedes that he needs good reasons for his theory, and argues that the acceptability of the theory depends on a demonstration of its applicability to the real world and its usefulness for dealing with the macroworld of everyday life and experience. He then presents some further arguments that his theory could have these virtues, and refers to previously published papers in which a promising start has been made in applying it to real-world topics (Floridi, 2008).

These arguments are all good and acceptable, but the fact is that Floridi has only just started with his demonstration that informational realism has these virtues of applicability and usefulness. A convincing demonstration to this effect would require much more work. Right now, a compelling argument for informational realism is lacking. In particular, Floridi has presented no argument that the notion of a mind-independent, primitive relation of difference even makes sense. We know what *mind-dependent* relations of difference are: they are relations that exist relative to the conceptual scheme of an agent who employs a set of criteria by which things can be measured to be different or identical. For instance, ethanol and methanol can be analyzed to be different according to some criteria but identical (i.e., both alcohol) according to others. But Floridi's objective relations of difference are noumenal and therefore unknown and, indeed, unknowable. How, then, can we make sense of such a concept? And why would we then go on to believe that such relations of difference bestow intrinsic value upon objects?

Unfortunately, the formalism of Object-Oriented Programming does little to make the notion of a noumenal relation of difference more understandable. Floridi does not claim that OOP provides a theory of model by which information objects can be defined, but he does claim that it can help to make the notion of an information object more precise. But if OOP does not specify adequate representational models of information objects, then it is not clear how it would do so. And there are, indeed, no reasons to believe that OOP can provide an ontological framework for the modeling of fundamental structures of reality. OOP is predominantly used to for task-oriented modeling of mundane systems, like processes in power plants or message flow in direct marketing. OOP is sometimes used for scientific modeling with the aim to provide accurate one-to-one mappings of some external reality, but such uses are controversial, as many do not hold OOP to be capable of providing reliable, objective mappings of reality in the way that successful sciences like chemistry and physics can. For example, Bertrand Meyer, one of the earliest and most visible proponents of OOP, has stated that in data modeling, “Reality is a cousin twice removed” (Meyer, 1997, p. 230). I conclude, therefore, that not only are there currently no compelling arguments for Floridi’s informational realism, but also that it is questionable that presently available formal methods allow for a more rigorous definition of an information object.

Conclusion

I have argued that Floridi has not yet succeeded in presenting a strong case for IE, and that several objections can be made against IE as currently defined. If Floridi wants to continue to defend IE in its current form, he will have to present both a compelling argument why we should attribute intrinsic value to information objects and a stronger case for his metaphysical theory of information objects, and account for the anti-egalitarian argument.

However, I believe that Floridi could solve many of the problems with IE that were discussed if he would no longer center IE around the notion of intrinsic value but around the notion of moral respect. As I have argued, inanimate objects can be deserving of

respect because of either their intrinsic, extrinsic, instrumental or emotional value. Moreover, being deserving of respect is a sufficient condition for being a moral patient. Floridi could hence argue that inanimate objects, although not possessive of intrinsic value, deserve respect because of either their extrinsic value or their (actual or potential) instrumental or emotional value for particular human beings (or animals) or for humanity as a whole. This move might have the implication that large classes of (informational) objects deserve no respect because they possess no non-intrinsic value and are therefore not moral patients, but perhaps a case could still be made that anything and everything that exists has at least *some* non-intrinsic value by which it is deserving of minimal respect.

The proposed transmutation of IE would likely also make IE easier to apply to issues in computer ethics. Digital information, while not possessive of intrinsic value, could be claimed to possess various sorts of non-intrinsic value by which it is deserving of respect and which bestows upon us duties to treat it with various kinds of respect, either for the sake of particular human beings (or animals) or for humanity as a whole. Similarly, our obligations towards both living beings and certain classes of inanimate objects would require us to refrain from creating, owning or distributing information that is disrespectful towards such entities and keeps them from persisting or flourishing. IE could hence have a bright future ahead, but only if moves from a value-centered ethics to a respect-centered ethics in which information objects must be respected due to their possession of various sorts of value.

References

- Callicott, J. B. 1995. "Intrinsic Value in Nature: a Metaethical Analysis," *Electronic Journal of Analytic Philosophy* 3 (Spring 1995).
- Floridi, L. 1999. "Information Ethics: On the Theoretical Foundations of Computer Ethics", *Ethics and Information Technology*, 1(1), 37-56.
- Floridi, L. 2002. "Information Ethics: An Environmental Approach to the Digital Divide", *Philosophy in the Contemporary World*, 9(1), 39-45.
- Floridi, L. 2003, "On the Intrinsic Value of Information Objects and the Infosphere", *Ethics and Information Technology*, 4(4), 287-304.

- Floridi, L. 2004. "Informational Realism," *Computers and Philosophy 2003 - Selected Papers from the Computers and Philosophy conference (CAP 2003)*, Conferences in Research and Practice in Information Technology, 37, ed. J. Weckert. and Y. Al-Saggaf (ACS), 7-12.
- Floridi, L. 2005a. Is Semantic Information Meaningful Data? *Philosophy and Phenomenological Research* 70(2), 351-70.
- Floridi, L. 2005b. Semantic Conceptions of Information. *Stanford Encyclopedia of Philosophy.*, ed. J. Zalta. (Online at <http://plato.stanford.edu/>.)
- Floridi, L. 2007a. "Global Information Ethics: The Importance of Being Environmentally Earnest," *International Journal of Technology and Human Interaction* 3(3), 1-11.
- Floridi, L. 2007b. "Information Ethics: Its Nature and Scope," *Moral Philosophy and Information Technology*, ed. Jeroen van den Hoven and John Weckert (Cambridge: Cambridge University Press), 40-65.
- Floridi, L. 2008. "A Defence of Informational Structural Realism," *Synthese* 161 (2), 219-253.
- Floridi, L., and Sanders, J. W. 2001. "Artificial Evil and the Foundation of Computer Ethics", *Ethics and Information Technology*, 3(1), 55-66.
- Floridi, L., and Sanders, J. W. 2002. "Computer Ethics: Mapping the Foundationalist Debate", *Ethics and Information Technology*, 4(1), 1-9.
- Meyer, B. 1997. *Object-Oriented Software Construction*, 2nd ed. Prentice Hall.