Afterword

LIS as Applied Philosophy of Information: A Reappraisal

Luciano Floridi

Luciano Floridi, Dipartimento di Scienze Filosofiche, Universitá di Bari, and Department of Philosophy, Oxford University, Wolfson College, OX2 6UD, Oxford, UK

Library information science (LIS) should develop its foundation in terms of a philosophy of information (PI). This seems a rather harmless suggestion. Where else could information science look for its conceptual foundations if not in PI? However, accepting this proposal means moving away from one of the few solid alternatives currently available in the field, namely providing LIS with a foundation in terms of social epistemology (SE). This is no trivial move, so some reasonable reluctance is to be expected. To overcome it, the proposal needs to be more than just acceptable; it must be convincing. In Floridi (2002a), I have articulated some of the reasons why I believe that PI can fulfill the foundationalist needs better than SE can. I won't rehearse them here. I find them compelling, but I am ready to change my mind if counterarguments become available. Rather, in this contribution, I wish to clarify some aspects of my proposal (Floridi, 2002a) in favor of the interpretation of LIS as applied PI. I won't try to show you that I am right in suggesting that PI may provide a foundation for LIS better than SE. My more modest goal is to remove some ambiguities and possible misunderstandings that might prevent the correct evaluation

of my position, so that disagreement can become more constructive.

We often hear about the differences between the ordinary librarian, busily involved in managing and delivering a public service, and the information scientist or the LIS expert, deep in theoretical speculations. The line of reasoning here seems that a foundation for LIS should satisfy both, and that this is something that PI cannot achieve, hence the objection that PI is not "social" enough. I accept the inference, but I disagree on the premise. For I think we should distinguish as clearly and neatly as possible between three main layers.

There is a first layer where we deal with libraries, their contents and services. Compare this with the accountant's calculations and financial procedures. One may wish to develop a theory of everyday mathematics and its social practices--surely this would be a worthy and interesting study--but it seems impossible to confuse it with the study of mathematics as a formal science. The latter is a second layer. It is what LIS amounts to, what one learns, with different degrees of complexity, through the university curriculum that educates a librarian or an information specialist. There is then a third layer, in which only a minority of people is interested. We call it foundational. For mathematics, it is the philosophy of mathematics. I suggested PI for LIS. My point here is that it is important to acknowledge and respect the distinction between these three layers; otherwise one may criticize x for not delivering y when x is not there to deliver y in the first place anyway. When checking whether the bank charged you too much for an overdraft, you are not expected to provide an analysis of the arithmetic involved in terms of Peano's axioms. Likewise, a scientist may be happy with a clear understanding of statistics without ever wishing to enter into the philosophical debate on the foundations of probability theory. So I do not see why LIS cannot be provided with an equally theoretical approach, capable of addressing

issues that the ordinary practitioner and the expert would deem too abstract to deserve attention in everyday practices (mind that I'm talking about layers not people; one can wear different hats in different contexts; this is not the issue here). In the end, I agree that PI should seek to explain a very wide range of phenomena and practices. I would add that this is precisely the challenge ahead. The scope of PI spans a whole variety of practices, precisely because the aim of PI is foundationalist.

If we assume for a moment that LIS is Applied PI, and that PI can provide LIS with a conceptual foundation, the next question concerns how, more specifically, PI and LIS may interact. This special issue provides plenty of evidence of the sort of fruitful investigations prompted by a PI approach to LIS. Three more examples may further illustrate the point and shed some more light on the SE vs. PI debate. A PI approach to the foundations of LIS may be expected to work on the ontology of its (i.e., LIS') "objects," on a substantial theory of information dynamics, and on an ethical approach to the domain of information. I shall say a bit more on each topic in the following pages, but let me stress here that if you find these areas of inquiry important, you also may want to concede that they fall beyond the scope of any SE approach. Let us now return to PI itself.

A simple way of introducing PI is by referring to it as the philosophical discipline that attempts to answer the question "What is information?". I understand that the question itself can be tiresome, at least because there is no simple way of answering it, despite its misleading simplicity. Ordinary moves, like checking the dictionary, consulting an encyclopaedia, compiling a survey, or piling up quotations, won't do, and not just because we are the ones who write those sources and entries anyway. Imagine trying to answer questions like "What is life?", "What is the mind?" or "What is meaning?" in the same way. Questions such as these are ways of opening a

dialogue and launching investigations that may keep generations of scholars and scientists busy. They are like road signs indicating the direction in which we should move. Complaining about the lack of precise answers is pointless. Philosophical questions are inevitably open-ended. We have to leave them behind to step ahead.

Information is a slippery topic. This explains its philosophical attractiveness but does not justify sloppy treatment. On the contrary, before we attempt our first steps beyond the "big-question sign" on this icy surface, we should make sure that our skates are razor-sharp. The slipperier the topic, the sharper the skates. Thus enters conceptual analysis, which can help us to understand the next point.

It is vital to realize that there is no single concept or unified theory of information (UTI) as such. The same people who think otherwise would be happy to acknowledge that it makes no sense to ask for a satisfactory single definition for food in general, and yet, we have a much better idea of what counts as food than of what may count as information. What we need is analysis, analysis, and more analysis.

At this point we should be wary of the opposite mistake, namely thinking that if there is no UTI then there is no theory at all. Wrong. Abandoning the search for a UTI implies giving up the assumption that there might be a Ur-concept of information at the roots of a hierarchical reconstruction of the multifarious world of information phenomena. But, as I have argued in Floridi (2003a), the various meanings, uses, applications and types of information, including the related phenomena in the environment, still can be interpreted as a system gravitating around a core notion with theoretical priority. The core notion works as a hermeneutical device that influences, interrelates, and helps to access other notions. In PI, this central role has long been claimed by factual or epistemically-oriented semantic information. The basic idea is simple. To understand what information is, the best thing to do is to start by analyzing

it in terms of the knowledge it can yield about its reference. Factual information is the most important and influential sense in which information <u>qua</u> information "can be said," to use an Aristotelian phrase. However, I strongly doubt (perhaps I should be more honest and state that I do not believe at all) that any successful attempt can be made to reduce all other concepts to factual information. Factual information is like the capital of the informational archipelagos, crucially positioned to provide a clear grasp of what information is and a privileged gateway to other important concepts that are interconnected but not necessarily reducible to a single <u>Ur</u>-concept. The right model is not a hierarchy but a distributed network of connected concepts, linked by mutual and dynamic influences that are not necessarily genetic or genealogical.

From what I have just said, it should be clear that I do not think, let alone suggest, that there is only one concept of information; that I do not subscribe to a communication-based concept of information; that I am skeptical of any unified theory of information based on any privileged concept (including Shannon's theory of information); that it is far from me to argue that we need a philosophy of "digital" information because of the IT revolution and the computational turn (although the latter has certainly made the pressing need for PI more obvious and broadly felt); but that I consider PI deeply concerned with the historical and logical dynamics of information as well as with its conceptual analysis. In Floridi (2002c) I explicitly supported the importance of a reinterpretation of several episodes in the history of philosophy in the light of the new informational paradigm. And in Floridi (2002b) I have tried to apply the historical analysis to the transmission of a specific <u>corpus</u> of texts.

Another thesis to which I do not subscribe concerns the existence of information in the world. I'm neither a naive nor a critical realist, but I'm not an antirealist either

(one nice thing about PI is that it helps us to dislodge old dichotomies). The position I have being trying to develop and support is that information comprises data, which are (part of, or) in the world independently of the epistemic agent. Data are better seen as <u>constraining affordances</u>, that is, differences that invite or facilitate ("afford") certain interpretations in relation to intelligent data-processors like us, while impeding, or making more difficult, ("constrain") some others. So "where" is information? An analogy may help to introduce the right answer. Would you say that there is no food in the world unless there are food-consumers? Of course not. Even if there is no form of life on planet Z, there may still be nutrients, let's say some minerals or water, which could sustain some form of life on Z. On the other hand, grass is food only for a grass-eater, and to a cat it is as good as a piece of plastic. <u>Mutatis mutandis</u>, one may argue that a radiograph is a piece of information about my lungs only for someone who can read it, whereas someone else may object that reading the radiograph is only a way of acknowledging the presence of the relevant piece of information in the environment. So where should we place "information"? The debate about the <u>locus informationis</u> has seen a tension between internalists and externalists. Some people place information "inside" the mind (e.g., the radiologist's understanding of the radiograph of my lungs); some others insist that it is in the world (e.g., the state of my lungs represented by the radiograph). This is a pointless dispute. When we consider food, it is clear that it is neither in the world, as mere nutrients, nor is just a function of the consumers' digestive systems. Likewise, information, and semantics in general, is one of those "two-dimensional things" that are neither here nor there, but at the interface between us and the environment, like a threshold between the two spaces. They are relational phenomena. This "liminal" conception of information is not reducible to some form of externalism (naturalization of

information) or internalism (information is in the mind of the beholder), so I endorse neither. I'm a <u>liminalist,</u> if this can be a label. I much prefer it to <u>ontologically amphibious</u>.

Is "liminalism" then just another form of antirealism? The answer depends on what we mean by the latter, and it requires some further explanations. When data acquire their meaning, we have what philosophers of language call "content." In Floridi (in press-a) and Floridi (in press-b), I have defended the view that to have factual information we need meaningful data (contents) that are also truthful. Two points need to be clarified here. First, content is a necessary condition for knowledge, but it is not the only one. Knowledge means something very precise in epistemology: it is content that is at least true (epistemologists speak of true beliefs) and possibly justified, that is, supported by some good reasons (guessing correctly that p is not yet knowing that p). The analysis of knowledge as true, justified belief is not satisfactory (see Floridi, in press-d), but this is no reason to think that we can do without the true condition. Speaking of "false knowledge" is nonsense, exactly like speaking of "married bachelors." Nobody can know that the earth is flat for the simple fact that it is not. Yet libraries are full of "false knowledge." So speaking of LIS as a discipline in which we are concerned with knowledge instead of content is at least imprecise and at worst mistaken. LIS deals with contents understood as meaningful data. This has nothing to do with data handling in the sense of a mechanical and brainless crunching and management of bytes. It is, rather, connected with the activity of stewardship of a semantic environment.

The second point concerns a possible commitment to a standard correspondence theory of truth. I said above that I do not think we should privilege a communication-based concept of information. The usual alternative--another tempting

dichotomy again--is to analyze information in terms of the representational

contribution it makes to our understanding of the world. So do I believe that

information represents the world, at least in some cases? Not quite, for I consider it

the wrong question to ask. I take the view, neither uncommon nor very popular among

philosophers, that the semanticization of data is a modeling process at some level of

abstraction (LoA). There is no space here to explain the methodology of LoAs

(Floridi & Sanders, in press-b), but one may grasp its gist by considering that,

according to the methodology, any access to data (and hence, any access to whatever

aspect of the world is under scrutiny) is mediated by an ontological commitment to a

level of abstraction that can be roughly understood as an interface. For example, we

are epistemic agents inevitably committed to a perception of space as Euclidean.

Now, the primary function of factual information seems to me to be the design, by the

agents inhabiting it, of an environment as meaningful as possible to the agents

themselves. Only part of this semanticization is adaptation-oriented. Most of it is

"superfluous." Once again, don't get me wrong: I'm not supporting the thesis that

"adaptation is beautiful." Quite the opposite, the secret of our special place in nature

seems to me to be hidden exactly in our "superfluous" detachment from the world

caused by our giving far more sense to it than it actually needs to make us prosper like

any other species. We are the animals that oversemanticize, and for no survival

purpose. But let's go back to the realist issue. We build our understanding of the

world by taking full advantage of the constraining affordances (data) offered by our

external sources at different LoA. Certainly, data only underdetermine the choice of a

particular LoA and the nature of its outcome, but underdetermination itself is not

Boolean and is inversely proportional to the degree of coherence among our LoAs.

Consider the following crossword analogy. Normally, crossword is a two-dimensional

puzzle, but one can easily imagine a three-dimensional version (a cube) in which coherence among the string of letters is even more difficult to achieve. Now consider an n-dimensional version, with n as large as one may wish. For example, a four-dimensional crossword would be one in which strings of letters have to satisfy the constraints that also regulate the diagonals of the cube, and so forth. This is what I mean by multidimensional intercoherence among LoAs.

Coherence among LoAs, however, may still guarantee at most some kind of internal "realism," if one forgets that the nature of the observables is also determined, partly, by the data being modeled. Whether empirical or conceptual, data afford only a certain range of models, and not all models are afforded equally easily. Another analogy may help here. Suppose you have to build a shelter. The design and complexity of the shelter may vary, but there is a limited range of "realistic" possibilities, determined by the nature of the resources available, the goals, etc. (size, building materials, location, weather, physical and biological environment, working force, technical skills, purposes, security, etc.). Not any shelter can be built. And the type of shelter that will be built more often will be the one that is more likely to take close-to-optimal advantage of the resources available. The same applies to data. Data are resources that make possible the construction of certain models, and the best models are those better tuned to their constraining affordances. This is what I mean by adequacy. Coherence and adequacy do not entail nor support naive or direct realism, or a correspondence theory of truth as this is ordinarily presented. Ultimately, LoAs construct models of data systems; they do not represent or photograph or portray or photocopy, or keyhole-spy or map or show or uncover or fax or monitor or ... the intrinsic nature of the systems they analyze, no more than an igloo describes the intrinsic nature of snow or the Parthenon indicates the real properties of stones. We

neither discover nor invent the world; we design it. So we understand it derivatively, only insofar as we understand its models. The world as we experience it scientifically or every day is the outcome of our modeling its data with a degree of intra-LoAs coherence as great as one may wish. This is neither a realist nor an antirealist but a constructionist view of information.

Approaching PI from a liminalist and constructionist perspective means adopting a metaphysical stance (Floridi, in press-c). And it is from this stance that information ethics (IE) should be evaluated. I have explained in other contexts why I believe in the importance of developing an ethics of stewardship toward the infosphere (see Floridi, 1999; Floridi and Sanders, 2001; Floridi and Sanders, 2002; Floridi and Sanders, in press-a). Here, I only wish to clarify an apparent misunderstanding. When I defend the minimal and overridable, intrinsic moral worth of informational objects (see Floridi, 2003b), I do not mean to refer to the moral value of an e-mail, or of Newton's Principia, or of any other piece of well-formed and meaningful data. Honestly, this would be rather silly. What I am suggesting is to approach the analysis of Being informationally, by adopting a minimal common ontology whereby human beings as well as animals, plants, artifacts, and so forth are interpreted as informational entities.

I hope the following analogy may be helpful, even if it is not really fair to the philosophical thesis at stake. Imagine looking at the whole universe from a chemical level of abstraction: you are 70% water and 30% something else. Now consider an informational level of abstraction. You are 100% a cluster of data. More precisely, you (as any other entity) are a discrete, self-contained, encapsulated package containing (i) the appropriate data structures, which constitute the nature of the entity in question: state of the object, its unique identity, and attributes and (ii) a collection

of operations, functions, or procedures, which are activated by various interactions or stimuli, namely messages received from other objects or changes within itself, and correspondingly define how the object behaves or reacts to them. At this level of abstraction, informational objects as such, rather than just living systems in general, are raised to the role of patients of any action. IE is then just an evolution of environmental ethics. Its fundamental tenet is that something is more elemental than life, namely Being understood informationally, and hence, something more fundamental than pleasure and pain, namely "entropy" (this is not the physicists' concept of entropy; entropy here means destruction of informational objects, that is, nothingness, in the vocabulary of the old substantialist metaphysics of Being). According to IE, one also should evaluate the duty of any moral agent in terms of contribution to the growth of the infosphere and any process, action, or event that negatively affects the whole infosphere--not just an informational object--as an increase in its level of entropy and hence an instance of evil. The ethical question asked by IE is: "What is good for an informational entity and the infosphere in general?" The answer is provided by a minimalist theory of deserts: any informational entity is recognized to be the center of some basic ethical claims, which deserve recognition and should help to regulate the implementation of any information process involving it, if possible. Approval or disapproval of any information process is then based on how the latter affects the essence of the informational entities it involves and, more generally, the whole infosphere, i.e., on how successful or unsuccessful it is in respecting the ethical claims attributable to the informational entities involved, and hence in improving or impoverishing the infosphere. IE brings to ultimate completion the process of enlarging the concept of what may count as a center of minimal moral concern, which now includes every informational entity.

Clearly, the relation between IE and LIS will be worth investigating.

References

[PI is a thriving area of research. Here I have limited myself to provide a list of works where I introduced it as a new philosophical field and I discussed its foundations. References to current literature less Floridi-centric can be found in the cited works. All the following papers are available from http://www.wolfson.ox.ac.uk/~floridi/ under "publications and papers online."]

Floridi, L. (1999). Information ethics: On the theoretical foundations of computer ethics. Ethics and Information Technology, 1(1), 37-56.

Floridi, L. (2002a). On defining library and information science as applied philosophy of information. Social Epistemology, 16(1), 37-49.

Floridi, L. (2002b). Sextus Empiricus: The transmission and recovery of Pyrrhonism. New York: Oxford University Press.

Floridi, L. (2002c). What is the philosophy of information? Metaphilosophy, 33(1-2), 123-145.

Floridi, L. (2003a). Information. In L. Floridi (Ed.), The Blackwell guide to the

philosophy of computing and information, (pp. 40-61). New York: Blackwell.

Floridi, L. (2003b). On the intrinsic value of information objects and the infosphere.

Ethics and Information Technology, 4(4), 287-304.

Floridi, L. (in press-a). Is information meaningful data? Philosophy and

Phenomenological Research.

Floridi, L. (in press-b). Outline of a theory of strongly semantic information. Minds

and Machines.

Floridi, L. (in press-c). Two approaches to the philosophy of information. Minds and

Machines.

Floridi, L. (in press-d). On the logical insolvability of the Gettier problem. Synthese.

Floridi, L., & Sanders, J. W. (2001). On the morality of artificial agents. In

Proceedings of CEPE 2001, Computer Ethics: Philosophical Enquiry, Lancaster

University, December 14-16.

Floridi, L., & Sanders, J. W. (2002). Computer ethics: Mapping the foundationalist

debate. Ethics and Information Technology, 4(1), 1-9.

Floridi, L., & Sanders, J. W. (in press-a). Internet ethics: The constructionist values of

Homo Poieticus. In R. Cavalier (Ed.), The impact of the Internet on our moral lives.

New York: SUNY.

Floridi, L., and Sanders, J. W. (in press-b). The method of abstraction. In M. Negrotti

(Ed.), Yearbook of the artificial. Nature, culture and technology. Models in

contemporary sciences. Bern: Peter Lang.