

Final version

Trends in the Philosophy of Information

Luciano Floridi

Dipartimento di Scienze Filosofiche, Università degli Studi di Bari; Faculty of Philosophy and IEG, OUCL, University of Oxford.

Address for correspondence: St Cross College, OX1 3LZ, Oxford, UK;
luciano.floridi@philosophy.oxford.ac.uk

1. Introduction

“I love information upon all subjects that come in my way, and especially upon those that are most important.” Thus boldly declares Euphranor, one of the defenders of Christian faith in Berkley’s *Alciphron*.¹ Evidently, information has been an object of philosophical desire and puzzlement for some time, well before the computer revolution, Internet or the dot.com pandemonium. Yet what does Euphranor love, exactly? *What is information?*

As with many other field-questions (consider for example “what is being?”, “what is morally good?” or “what is knowledge?”), “what is information?” is to be taken not as a request for a dictionary definition, but as a means to demarcate a wide area of research. The latter has recently been defined as *the philosophy of information* (Floridi [2002], Floridi [2003b]). The task of this chapter is to review some interesting research trends in the philosophy of information (henceforth also PI). This will be achieved in three steps. We shall first look at a definition of PI. On this basis, we shall then consider a series of open problems in PI on which philosophers are currently working.² The conclusion will then highlight the innovative character of this new area of research.

2. Defining the Philosophy of Information

The philosophy of information may be defined as the philosophical field concerned with

- a) the critical investigation of the conceptual nature and basic principles of information, including its dynamics, utilisation and sciences, and
- b) the elaboration and application of information-theoretic and computational methodologies to philosophical problems.³

The first half of the definition concerns the philosophy of information as a new field. PI appropriates an explicit, clear and precise interpretation of the classic, Socratic question “*ti esti...?*” (“what is...?”), namely “What is the nature of information?”. This is the clearest hallmark of a new field. PI provides critical investigations that are not to be confused with a quantitative theory of data communication (information theory). On the whole, we shall see that its task is to develop an integrated family of theories that analyse, evaluate and explain the various principles and concepts of information, their dynamics and utilisation, with special attention to systemic issues arising from different contexts of application and the interconnections with other key concepts in philosophy, such as knowledge, truth, meaning and reality.

By “dynamics of information” the definition refers to:

- i) *the constitution and modelling of information environments*, including their systemic properties, forms of interaction, internal developments, applications etc.;
- ii) *information life cycles*, i.e. the series of various stages in form and functional activity through which information can pass, from its initial occurrence to its final utilisation and possible disappearance;⁴ and
- iii) *computation*, both in the Turing-machine sense of *algorithmic processing*, and in the wider sense of *information processing*. This is a crucial specification. Although a very old concept, information has finally acquired the nature of a primary phenomenon only thanks to the sciences and technologies of computation and ICT (Information and Communication Technologies). Computation has therefore attracted much philosophical attention in recent years. Nevertheless, PI privileges

“information” over “computation” as the pivotal topic of the new field because it analyses the latter as presupposing the former. PI treats “computation” as only one (although very important) of the processes in which information can be involved.

From an environmental perspective, PI is critical and normative about what may count as information, and how information should be adequately created, processed, managed and used. Methodological and theoretical choices in ICS are also profoundly influenced by the kind of PI a researcher adopts more or less consciously. It is therefore essential to stress that PI critically evaluates, shapes and sharpens the conceptual, methodological and theoretical basis of ICS, in short that it also provides a *philosophy of ICS*, as this has been plain since early work in the area of philosophy of AI (Colburn [2000]).

It is worth stressing here that an excessive concern with contemporary issues may lead one to miss the important fact that it is perfectly legitimate to speak of a philosophy of information even in authors who lived before the information revolution, and hence that it will be extremely fruitful to develop a historical approach and trace PI’s diachronic evolution, as long as the technical and conceptual frameworks of ICS are not anachronistically applied, but are used to provide the conceptual method and privileged perspective to evaluate in full reflections that were developed on the nature, dynamics and utilisation of information before the digital revolution. This is significantly comparable with the development undergone by other philosophical fields like the philosophy of language, the philosophy of biology, or the philosophy of mathematics.⁵

The second half of the definition indicates that PI is not only a new field, but provides an innovative methodology as well. Research into the conceptual nature of information, its dynamics and utilisation is carried on from the vantage point represented by the methodologies and theories offered by ICS and ICT (Grim et al. [1998] and Greco et al. [2005]). This perspective affects other philosophical topics as well. Information-theoretic and computational methods, concepts, tools and techniques have already been developed and applied in many philosophical areas,

- to extend our understanding of the cognitive and linguistic abilities of humans and animals and the possibility of artificial forms of intelligence (e.g. in the philosophy of AI; in information-theoretic semantics; in information-theoretic epistemology and in dynamic semantics);

- to analyse inferential and computational processes (e.g. in the philosophy of computing; in the philosophy of computer science; in information-flow logic; in situation logic; in dynamic logic and in various modal logics);
- to explain the organizational principles of life and agency (e.g. in the philosophy of artificial life; in cybernetics and in the philosophy of automata; in decision and game theory);
- to devise new approaches to modelling physical and conceptual systems (e.g. in formal ontology; in the theory of information systems; in the philosophy of virtual reality);
- to formulate the methodology of scientific knowledge (e.g. in model-based philosophy of science; in computational methodologies in philosophy of science);
- to investigate ethical problems (in computer and information ethics and in artificial ethics), aesthetic issues (in digital multimedia/hypermedia theory, in hypertext theory and in literary criticism) and psychological, anthropological and social phenomena characterising the information society and human behaviour in digital environments(cyberphilosophy).

Indeed, the previous examples and the various chapters in this volume show that PI, as a new field, provides a unified and cohesive, theoretical framework that allows further specialisation.

3. Open Problems in the Philosophy of Information

PI possesses one of the most powerful conceptual vocabularies ever devised in philosophy. This is because we can rely on informational concepts whenever a complete understanding of some series of events is unavailable or unnecessary for providing an explanation. In philosophy, this means that virtually any issue can be rephrased in informational terms. This semantic power is a great advantage of PI understood as a methodology (see the second half of the definition). It shows that we are dealing with an influential paradigm, describable in terms of an informational philosophy. But it may also be a problem, because a metaphorically pan-informational approach can lead to a dangerous equivocation, namely thinking that since any x can be described in (more or less metaphorically) informational terms,

then the nature of any x is genuinely informational. And the equivocation obscures PI's specificity as a philosophical field with its own subject. PI runs the risk of becoming synonymous with philosophy. The best way of avoiding this loss of identity is to concentrate on the first half of the definition. PI as a philosophical discipline is defined by what a problem is (or can be reduced to be) *about*, not by *how* the latter is formulated. Although many philosophical issues seem to benefit greatly from an informational analysis, in PI one presupposes that a problem or an explanation can be legitimately and genuinely reduced to an informational problem or explanation. So the criterion to test the soundness of the informational analysis of x is not to check whether x *can* be formulated in informational terms but to ask what would be like for x not to have an informational nature at all. With this criterion in mind, we shall now review some of the most interesting problems in PI.

For reasons of space, only some research trends and issues could be included and even those selected are only briefly outlined and not represented with adequate depth, sophistication and significance. This is not only because of space, but also because the interested reader will find a wealth of further material in the other chapters of this Handbook. The issues included have been privileged because they represent macroproblems, that is, they are the hardest to tackle but also the ones that have the greatest influence on clusters of microproblems to which they can be related as theorems to lemmas. Some microproblems are mentioned whenever they seem interesting enough, but especially in this case the list is far from exhaustive. Some problems are new, others are developments of old problems, and in some cases philosophers have already begun to address them, but the review does not concern old trends and problems that have already received their due philosophical attention. There is also no attempt at keeping a uniform level of scope. Some problems are very general, others more specific. All of them have been chosen because they well indicate how vital and useful the new paradigm is in a variety of philosophical areas. Finally, whenever possible I have indicated which chapters in the Handbook are relevant to the problem under discussion.

4. The nature of information

This is the hardest and most central question in PI. It has received many answers in different fields but, unsurprisingly, several surveys do not even converge on a

single, unified definition of information (see for example Braman [1989], Losee [1997], Machlup and Mansfield [1983], Debons and Cameron [1975], Larson and Debons [1983]). Information is notoriously a polymorphic phenomenon and a polysemantic concept so, as an explicandum, it can be associated with several explanations, depending on the level of abstraction adopted and the cluster of requirements and desiderata orientating a theory. Claude E. Shannon, for one, was very cautious: “The word ‘information’ has been given different meanings by various writers in the general field of information theory. It is likely that at least a number of these will prove sufficiently useful in certain applications to deserve further study and permanent recognition. *It is hardly to be expected that a single concept of information would satisfactorily account for the numerous possible applications of this general field.* (italics added)” (Shannon [1993], p. 180). Thus, following Shannon, Weaver [1949] supported a tripartite analysis of information in terms of (1) technical problems concerning the quantification of information and dealt with by Shannon’s theory; (2) semantic problems relating to meaning and truth; and (3) what he called “influential” problems concerning the impact and effectiveness of information on human behaviour, which he thought had to play an equally important role. And these are only two early examples of the problems raised by any analysis of information.

Indeed, the plethora of different analyses can be confusing. Complaints about misunderstandings and misuses of the very idea of information are frequently expressed, even if to no apparent avail. Sayre [1976], for example, already criticised the “laxity in use of the term ‘information’” in Armstrong [1968] (see now Armstrong [1993]) and in Dennett [1969] (see now Dennett [1986]), despite appreciating several other aspects of their work. More recently, Harms [1998] pointed out similar confusions in Chalmers [1996], who “seems to think that the information theoretic notion of information [see section 3, my addition] is a matter of what possible states there are, and how they are related or structured [...] rather than of how probabilities are distributed among them” (p. 480).

Information remains an elusive concept. This is a scandal not by itself, but because so much basic theoretical work, both in science and in philosophy, relies on a clear grasping of the nature of information and of its cognate concepts. We know that information ought to be quantifiable (at least in terms of partial ordering),

additive, storable and transmittable. But apart from this, we still do not seem to have a much clearer idea about its specific nature.

Information is often approached from three perspectives: information as reality (e.g. as patterns of physical signals, which are neither true nor false), also known as *ecological* information; information about reality (semantic information, which is alethically qualifiable and an ingredient in the constitution of knowledge); and information for reality (instruction, like genetic information, algorithms and recipes). Many extensionalist approaches to the definition of information as/about reality provide different starting points. The following list contains only some of the most philosophically interesting or influential, and I shall say a bit more about each of them presently. They are not to be taken as necessarily alternative, let alone incompatible:

1. the communication theory approach (mathematical theory of codification and communication of data/signals (Shannon and Weaver [1949 rep. 1998]; see also the chapter by Topsøe and Harremoës) defines information in terms of probability space distribution;
2. the algorithmic approach (also known as Kolmogorov complexity, Li and Vitányi [1997]; see also the chapters by Grunwald and Vitányi and by Adriaans) defines the information content of X as the size in bits of the smallest computer program for calculating X (Chaitin [2003]);
3. the probabilistic approach (Bar-Hillel and Carnap [1953], Bar-Hillel [1964], Dretske [1981]; see also the chapter by Dretske), is directly based on (1) above and defines semantic information in terms of probability space and the inverse relation between information in p and probability of p ;
4. the modal approach defines information in terms of modal space and in/consistency (the information conveyed by p is the set of possible worlds excluded by p);
5. the systemic approach (situation logic, Barwise and Perry [1983], Israel and Perry [1990], Devlin [1991]) defines information in terms of states space and consistency (information tracks possible transitions in the states space of a system);

6. the inferential approach defines information in terms of inferences space (information depends on valid inference relative to a person's theory or epistemic state);
7. the semantic approach (Floridi [2004c], Floridi [2005b]) defines information in terms of data space (semantic information is well-formed, meaningful and truthful data).

Each extentionalist approach can be given an intentionalist reading by interpreting the relevant space as a doxastic (i.e. belief-related) space, in which information is seen as a reduction in the degree of uncertainty or level of surprise given a state of knowledge of the informee (see the chapters by Baltag, Moss and van Ditmarsch and by Rott).

Communication theory in (1) approaches information as a physical phenomenon, syntactically. It is not interested in the usefulness, relevance, meaning, interpretation or reference of data, but in the level of detail and frequency in the uninterpreted data (signals or messages). It provides a successful mathematical theory because its central question is whether and how much data, not what information is conveyed.

The algorithmic approach in (2) is equally quantitative and solidly based on probability theory. It interprets information and its quantities in terms of the computational resources needed to specify it.

The remaining approaches all address the question “what is *semantic* information?”. They seek to give an account of information as semantic content, usually adopting a propositional orientation (they analyse examples like “The earth has only one moon”). Do (1) or (2) provide the necessary conditions for any theory of semantic information in (3)-(7)? Are all the remaining semantic approaches mutually compatible? Is there a logical hierarchy? Do any of the previous approaches provide a clarification of the notion of data as well? Most of the problems in PI acquire a different meaning depending on how we answer this cluster of questions. Indeed, positions might be more compatible than they initially appear owing to different interpretations of the concept(s) of information involved.

Once the concept of information is clarified, each of the previous approaches needs to address the following question.

5. The dynamics of information

The question does not concern the nature of management processes (information seeking, data acquisition and mining, information harvesting and gathering, storage, retrieval, editing, formatting, aggregation, extrapolation, distribution, verification, quality control, evaluation, etc.) but, rather, information processes themselves, whatever goes on between the input and the output phase. Communication theory, as the mathematical theory of data transmission, provides the necessary conditions for any physical communication of information, but is otherwise of only marginal help. The information flow—understood as the carriage and transmission of information by some data about a referent, made possible by regularities in a distributed system—has been at the centre of logical studies for some time (Barwise and Seligman [1997], Van Benthem [2003]), but still needs to be fully explored. How is it possible for something to carry information about something else? The problem here is not yet represented by the “aboutness” relation, which needs to be discussed in terms of meaning, reference and truth. The problem here concerns the nature of data as vehicles of information. In this version, the problem plays a central role in semiotics, hermeneutics and situation logic. It is closely related to the problem of the naturalisation of information. Various other logics, from classic first order logic to epistemic, erotetic and dynamic logic, provide useful approaches with which to analyse the logic of information, but there is still much work to be done (Van Benthem and Van Rooy [2003], Allo [forthcoming], Allo and Floridi [forthcoming], Floridi [forthcoming]).

Information processing, in the general sense of information states transitions, includes at the moment effective computation (*computationalism*, Newell [1980]; Pylyshyn [1984]; Fodor [1975];[1987]; Dietrich [1990]), distributed processing (*connectionism*, Smolensky [1988]; Churchland and Sejnowski [1992]), and dynamical-system processing (*dynamism*, Van Gelder [1995]; Van Gelder and Port [1995]; Eliasmith [1996]). The relations between the current paradigms remain to be clarified (Minsky [1990], for example, argues in favour of a combination of computationalism and connectionism in AI, as does Harnad [1990] in cognitive science), as do the specific advantages and disadvantages of each, and the question

as to whether they provide complete coverage of all possible *internalist* information processing methods.

The two previous questions in §§ 4 and 5 and are closely related to a third, more general problem.

6. The challenge of unified theory of information

The reductionist approach holds that we can extract what is essential to understanding the concept of information and its dynamics from the wide variety of models, theories and explanations proposed. The non-reductionist argues that we are probably facing a network of logically interdependent but mutually irreducible concepts. The plausibility of each approach needs to be investigated in detail. Both approaches, as well as any other solution in between, are confronted by the difficulty of clarifying how the various meanings and phenomena of information are related, and whether some concepts of information are more central or fundamental than others and should be privileged. Waving a Wittgensteinian suggestion of family resemblance means only acknowledging the problem, not solving it. The reader interested in a positive answer the question may wish to read the essays collected in Hofkirchner [1998]. A defence of a more skeptical view, following Shannon, can be found in Floridi [2003a].

7. The data grounding problem: how data acquire their meaning

We have seen that most analyses of the nature of information tend to concentrate on its semantic features, quite naturally. So it is useful to carry on our review of problem areas in PI by addressing next the cluster of issues arising in informational semantics. Their discussion is bound to be deeply influential in several areas of philosophical research. But first, a warning. It is hard to formulate problems clearly and in some detail in a completely theory-neutral way. So in what follows, the semantic frame will be adopted (see above § 4, (7)), namely the view that semantic information can be satisfactorily analysed in terms of well-formed, meaningful and veridical data. This semantic approach is simple and powerful enough for the task at hand. If the problems selected are sufficiently robust, it is reasonable to expect that their general nature and significance are not relative to the theoretical vocabulary in which they are cast but will be exportable across conceptual platforms.

We have already encountered the issue of the nature of data. Suppose data are intuitively described as uninterpreted differences (symbols or signals). How do they become meaningful? This is the data grounding problem.

Searle [1990] refers to a specific version of the data grounding problem as the problem of intrinsic meaning or “intentionality”. Harnad [1990] defines it as the symbols grounding problem and unpacks it thus: “How can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads? How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols?” (p. 335).

Arguably, the frame problem (how a situated agent can represent, and interact with, a changing world satisfactorily) and its sub-problems are a consequence of the data grounding problem (Harnad [1993], Taddeo and Floridi [2005]). In more metaphysical terms, this is the problem of the semanticisation of being and it is further connected with the problem of whether information can be naturalised.

8. The semantic problem: how meaningful data acquire their truth value

Once grounded, meaningful data can acquire different truth values, the question is how. The question then gains new dimensions when asked within epistemology and the philosophy of science. It also interacts with the way in which we approach both a theory of truth and a theory of meaning, especially a truth-functional one (see the chapter by Groenendijk, Kamp and Stokhof). Are truth and meaning understandable on the basis of an informational approach, or is it information that needs to be analysed in terms of non-informational theories of meaning and truth? To call attention to this important set of issues it is worth asking two more place-holder questions:

1) can information explain truth?

In this, as in the following question, we are not asking whether a specific theory could be couched, more or less metaphorically, in some informational vocabulary. This would be a pointless exercise. What is in question is not even the mere possibility of an informational approach. Rather, we are asking

1.a) could an informational theory explain truth more satisfactorily than other current approaches? And

1.b), should (1.a) be answered in the negative, could an informational approach at least help to clarify the theoretical constraints to be satisfied by other approaches?

The second major question mentioned above is:

2) can information explain meaning?

Several informational approaches to semantics have been investigated in epistemology (Dretske [1981]; Dretske [1988]), situation semantics (Seligman and Moss [1997]), discourse representation theory (Kamp [1984]) and dynamic semantics (Muskens and Al. [1997]). Is it possible to analyse meaning not truth-functionally but as the potential to change the informational context? Can semantic phenomena be explained as aspects of the empirical world? Since the problem is whether meaning can at least partly be grounded in an objective, mind- and language-independent notion of information (naturalisation of intentionality), it is strictly connected with the problem of the naturalisation of information.

9. Information processing and the study of cognition

Information and its dynamics are central to the foundations of AI and of cognitive science (see the chapters by McCarthy and Boden). Both disciplines study cognitive agents as informational systems that receive, store, retrieve, transform, generate and transmit information. This is the *information processing* view. Before the development of connectionist and dynamic-system models of information processing, it was also known as the computational view. The latter expression was acceptable when a Turing machine (Turing [1936]) and the machine involved in the Turing test (Turing [1950]) were inevitably the same. It has recently become misleading, however, because computation, when used as a technical term (effective computation), refers now to the specific class of algorithmic symbolic processes that can be performed by a Turing machine, that is recursive functions (Turing [1936], Minsky [1967]; Floridi [1999]; Boolos et al. [2002]).

The information processing view of cognition, intelligence and mind provides the oldest and best-known cluster of significant problems in PI.⁶ Some of their formulations, however, have long been regarded as uninteresting. Turing [1950] considered “can machines think?” a meaningless way of posing the otherwise interesting problem of the functional differences between AI and NI (natural intelligence). Searle [1990] has equally dismissed “is the brain a digital computer?”

as ill-defined. The same holds true of the unqualified question “are naturally intelligent systems information processing systems?”. Such questions are vacuous. Informational concepts are so powerful that, given the right level of abstraction (LoA; Floridi and Sanders [2004]; Floridi and Sanders [forthcoming]), anything can be presented as an information system, from a building to a volcano, from a forest to a dinner, from a brain to a company, and any process can be simulated informationally - heating, flying and knitting. So pancomputationalist views have the hard task of providing a credible answer to the question: what would it mean for a physical system *not* to be an informational system (that is, a computational system, if computation is used to mean information processing, see Chalmers [1996] and Chalmers [online])? The task is hard because pancomputationalism does not seem vulnerable to a refutation, in the form of a realistic token counterexample in a world nomically identical to the one to which pancomputationalism is applied.⁷ A good way of posing the problem is not: “is ‘x is y’ adequate?”, but rather “if ‘x is y’ at LoA z, is z adequate?”.

10. Science and information modelling

In many contexts (probability or modal states and inferential spaces), we often adopt a conditional, laboratory view. We analyse what happens in “*a*’s being (of type, or in state) *F* is correlated to *b* being (of type, or in state) *G*, thus carrying for the observer the information that *b* is *G*” (Barwise and Seligman [1997] provide a similar analysis based on Dretske [1981]) by assuming that $F(a)$ and $G(b)$. In other words, we assume a given model. The question asked here is: how do we build the original model? Many approaches seem to be ontologically over-committed. Instead of assuming a world of empirical affordances and constraints to be designed, they assume a world already well-modelled, ready to be discovered. The semantic approach to scientific theories (Suppes [1960]; Suppes [1962]; Van Fraassen [1980]; Giere [1988]; Suppe [1989]), on the other hand, argues that “scientific reasoning is to a large extent model-based reasoning. It is models almost all the way up and models almost all the way down.” (Giere [1999], 56).

Theories do not make contact with phenomena directly, but rather higher models are brought into contact with other, lower models. These are themselves theoretical conceptualisations of empirical systems, which constitute an object being

modelled as an object of scientific research. Giere [1988] takes most scientific models of interest to be non-linguistic abstract objects. Models, however, are the medium, not the message. Is information the (possibly non-linguistic) content of these models? How are informational models (semantically, cognitively and instrumentally) related to the conceptualisations that constitute their empirical references? What is their semiotic status, e.g. structurally homomorphic or isomorphic representations or data-driven and data-constrained informational constructs? What levels of abstraction are involved? Is science a social (multi-agents), information-designing activity? Is it possible to import, in (the philosophy of) science, modelling methodologies devised in information system theory? Can an informational view help to bridge the gap between science and cognition? Answers to these questions are closely connected with the discussion of the problem of an informational theory of truth see above. The reader interested in some specific applications will find them in the chapters by Devlin and Rosenberg, and by Collier.

The possibility of a more or less informationally constructionist philosophy of science leads to our next cluster of problems, concerning the relation between information and the natural world.

11. The ontological status of information

Barwise and Seligman [1997] have remarked that “If the world were a completely chaotic, unpredictable affair, there would be no information to process. Still, the place of information in the natural world of biological and physical systems is far from clear.” (p. xi). This lack of clarity prompts a whole family of problems.

It is often argued that there is no information without (data) representation. Following Landauer and Bennett [1985]; Landauer [1987]; Landauer [1991]; Landauer [1996], this principle is usually interpreted materialistically, as advocating the impossibility of physically disembodied information, through the equation “representation = physical implementation”. The view that there is no information without physical implementation is an inevitable assumption, when working on the physics of computation, since computer science must necessarily take into account the physical properties and limits of the carriers of information. It is also the ontological assumption behind the Physical Symbol System Hypothesis in AI and cognitive science (Newell and Simon [1976]). However, the fact that information

requires a representation does not entail that the latter ought to be physically implemented. Arguably, environments in which there are only noetic entities, properties and processes (e.g. Berkeley, Spinoza), or in which the material or extended universe has a noetic or non-extended matrix as its ontological foundation (e.g. Pythagoras, Plato, Leibniz, Hegel), seem perfectly capable of upholding the representationalist principle without also embracing a materialist interpretation (see Floridi [2004a] for a defence of this view). The relata giving rise to information could be monads, for example. So the problem here becomes: is the informational an independent ontological category, different from the physical/material and (assuming one could draw this Cartesian distinction) the mental? Wiener, for example, thought that “Information is information, not matter or energy. No materialism which does not admit this can survive at the present day” (Wiener [1948], 132).

If the informational is not an independent ontological category, to which category is it reducible? If it is an independent ontological category, how is it related to the physical/material and the mental? Answers to these questions determine the orientation a theory takes with respect to the following problem.

12. Naturalised information

The problem is connected with the semanticisation of data. It seems hard to deny that information is a natural phenomenon, so this is not what one should be asking here. Even elementary forms of life such as sunflowers survive because they are capable of some chemical data processing. The problem here is whether there is information in the world independently of forms of life capable to extract it and, if so, what kind of information is in question (an informational version of the teleological argument for the existence of God argues both that information is a natural phenomenon and that the occurrence of environmental information requires an intelligent source). If the world is sufficiently information-rich, perhaps an agent may interact successfully with it by using “environmental information” directly, without being forced to go through a representation stage in which the world is first analysed informationally. “Environmental information” still presupposes (or perhaps is identical with) some physical support but it does not require any higher-level cognitive representation or computational processing to be immediately usable. This

is argued, for example, by researchers in AI working on animats (artificial animals, either computer simulated or robotic). Animats are simple reactive agents, stimulus-driven. They are capable of elementary, “intelligent” behaviour, despite the fact that their design excludes the possibility of internal representations of the environment and any effective computation (Mandik [2002] for an overview, the case for non-representational intelligence is famously made by Brooks [1991]). So, are cognitive processes continuous with processes in the environment? Is semantic content (at least partly) external (Putnam)? Does “natural” or “environmental” information pivot on natural signs (Peirce) or nomic regularities? Consider the typical example provided by the concentric rings visible in the wood of a cut tree trunk, which may be used to estimate the age of the plant. The externalist/extensionalist, who favours a positive answer (e.g. Dretske and Barwise), is faced by the difficulty of explaining what kind of information and how much of it saturates the world, what kind of access to, or interaction with “information in the world” an informational agent can enjoy, and how information dynamics is possible. The internalist/intentionalist (e.g. Fodor and Searle), who privileges a negative answer, needs to explain in what specific sense information depends on intelligence and whether this leads to an anti-realist view.

The location of information is related to the question whether there can be information without an informee, or whether information, in at least some crucial sense of the word, is essentially parasitic on the meanings in the mind of the informee, and the most it can achieve, in terms of ontological independence, is systematic interpretability. Before the discovery of the Rosetta Stone, was it legitimate to regard Egyptian hieroglyphics as information, even if their semantics was beyond the comprehension of any interpreter? Admitting that computers perform some minimal level of proto-semantic activity works in favour of a “realist” position about “information in the world”.

Before moving to the next problem, it remains to be clarified whether the previous two ways of locating information might not be restrictive. Could information be neither here (intelligence) nor there (natural world) but on the threshold, as it were, as a special relation or interface between the world and its intelligent inhabitants (constructionism)? Or could it even be elsewhere, in a third world, intellectually accessible by intelligent beings but not ontologically dependent

on them (Platonism)? The reader interested in the physics of information is advised to read the chapter by Bais and Farmer.

13. The It from Bit hypothesis

Can nature be informationalised? The neologism “informationalised” is ugly but useful to point out that this is the converse of the previous problem. Here too, it is important to clarify what the problem is not. We are not asking whether the metaphorical interpretation of the universe as a computer is more useful than misleading. We are not even asking whether an informational description of the universe, as we know it, is possible, at least partly and piecemeal. This is a challenging task, but formal ontologies already provide a promising answer (Smith [2004]). We are asking whether the universe in itself could essentially be made of information, with natural processes, including causation, as special cases of information dynamics (e.g. information flow and algorithmic, distributed computation and forms of emergent computation). Depending on how one approaches the concept of information, it might be necessary to refine the problem in terms of digital data or other informational notions.

Answers to this problem deeply affect our understanding of the distinction between virtual and material reality, of the meaning of artificial life in the ALife sense (Bedau [2004]), and of the relation between the philosophy of information and the foundations of physics: if the universe is made of information, is quantum physics a theory of physical information? Moreover, does this explain some of its paradoxes? If nature can be informationalised, does this help to explain how life emerges from matter, and hence how intelligence emerges from life? “Can we build a gradualist bridge from simple amoeba-like automata to highly purposive intentional systems, with identifiable goals, beliefs, etc.?” (Dennett [1998], 262).

14. Conclusion

Our brief survey ends here. We have had a quick look to many questions of a wide variety of nature and scope. This should not be disheartening. On the contrary, we saw at the beginning of this chapter that Berkeley-Euphranor loved “information upon all subjects”. It has required several scientific, technological and social transformations, but philosophers have finally begun to address the new intellectual

challenges arising from the world of information and the information society. Michael Dummett recently acknowledged that ““Evans had the idea that there is a much cruder and more fundamental concept than that of knowledge on which philosophers have concentrated so much, namely the concept of information. Information is conveyed by perception, and retained by memory, though also transmitted by means of language. One needs to concentrate on that concept before one approaches that of knowledge in the proper sense. Information is acquired, for example, without one’s necessarily having a grasp of the proposition which embodies it; the flow of information operates at a much more basic level than the acquisition and transmission of knowledge. I think that this conception deserves to be explored. It’s not one that ever occurred to me before I read Evans, but it is probably fruitful. That also distinguishes this work very sharply from traditional epistemology.” (Dummett [1993], p. 186) Dummett is arguably correct. PI evolves out of the analytic movement, but does not seem to belong to it. It attempts to expand the frontier of philosophical research, not by putting together pre-existing topics, and thus reordering the philosophical scenario, but by enclosing new areas of philosophical inquiry—which have been struggling to be recognised and may not yet found room in the traditional philosophical syllabus—and by providing innovative methodologies to address traditional problems from new perspectives. Clearly, PI promises to be one of the most exciting and fruitful areas of philosophical research of our time. As this volume proves, it is already affecting the overall way in which new and old philosophical problems are being addressed, bringing about a substantial innovation of the philosophical system. This represents the information turn in philosophy.⁸

References

- Adams, F. 2003, "The Informational Turn in Philosophy", *Minds and Machines*, 13(4), 471-501.
- Allo, P. forthcoming, "Being Informative", *Lecture Notes in Artificial Intelligence*.
- Allo, P., and Floridi, L. forthcoming, "Logic and the Philosophy of Information", *Special Issue of Logique et Analyse*.
- Anderson, A. R. 1964, *Minds and Machines* (Englewood Cliffs: Prentice-Hall).
- Armstrong, D. M. 1968, *A Materialist Theory of the Mind* (London: Routledge & Kegan Paul).
- Armstrong, D. M. 1993, *A Materialist Theory of the Mind* 2nd edition (London: Routledge).
- Bar-Hillel, Y. 1964, *Language and Information : Selected Essays on Their Theory and Application* (Reading, Mass ; London: Addison-Wesley).
- Bar-Hillel, Y., and Carnap, R. 1953, "An Outline of a Theory of Semantic Information" repr. in Bar-Hillel [1964], pp. 221-74.
- Barwise, J., and Perry, J. 1983, *Situations and Attitudes* (Cambridge, Mass.: MIT Press).
- Barwise, J., and Seligman, J. 1997, *Information Flow: The Logic of Distributed Systems* (Cambridge: Cambridge University Press).
- Bedau, M. 2004, "Artificial Life" in *The Blackwell Guide to the Philosophy of Computing and Information*, edition, edited by L. Floridi (New York - Oxford: Blackwell), chap. 16.
- Berkeley, G. 1732, *Alciphron: Or the Minute Philosopher* (Edinburgh: 1948-57: Thomas Nelson).
- Boden, M. A. 1990, *The Philosophy of Artificial Intelligence* (Oxford: Oxford University Press).
- Boolos, G., Burgess, J. P., and Jeffrey, R. C. 2002, *Computability and Logic* 4th ed. (Cambridge: Cambridge University Press).
- Braman, S. 1989, "Defining Information", *Telecommunications Policy*, 13, 233-242.
- Brooks, R. 1991, "Intelligence without Representation", *Artificial Intelligence*, 47, 139-159.
- Chaitin, G. 2003, "Two Philosophical Applications of Algorithmic Information Theory", *Proceedings DMTCS'03 - Springer Lecture Notes in Computer Science*, 2731, 1-10.
- Chalmers, D. online, *A Computational Foundation for the Study of Cognition*,
- Chalmers, D. J. 1996, *The Conscious Mind : In Search of a Fundamental Theory* (New York: Oxford Univ. Press).
- Churchland, P. S., and Sejnowski, T. J. 1992, *The Computational Brain* (Cambridge MA: MIT/Bradford Press).
- Colburn, T. R. 2000, *Philosophy and Computer Science* (Armonk, N.Y.: M.E. Sharpe).
- Debons, A., and Cameron, W. J. (ed.) 1975, *Perspectives in Information Science : Proceedings of the Nato Advanced Study Institute on Perspectives in Information Science, Held in Aberystwyth, Wales, Uk, August 13-24, 1973* (Leiden: Noordhoff).
- Dennett, D. 1998, *Brainchildren* (Cambridge Ma: MIT Press).
- Dennett, D. C. 1969, *Content and Consciousness* (London: Routledge & Kegan Paul).

- Dennett, D. C. 1986, *Content and Consciousness* 2nd edition (London: Routledge & Kegan Paul).
- Devlin, K. J. 1991, *Logic and Information* (Cambridge: Cambridge University Press).
- Dietrich, E. 1990, "Computationalism", *Social Epistemology*, 4, 135-154.
- Dretske, F. 1988, *Explaining Behavior* (Cambridge, Ma: MIT Press).
- Dretske, F. I. 1981, *Knowledge and the Flow of Information* (Oxford: Blackwell).
Reprinted in 1999 (Stanford, CA: CSLI Publications).
- Dummett, M. 1993, *Origins of Analytical Philosophy* (London: Duckworth).
- Eliasmith, C. 1996, "The Third Contender: A Critical Examination of the Dynamicist Theory of Cognition", *Journal of Philosophical Psychology*, 9(4), 441-463.
- Floridi, L. 1999, *Philosophy and Computing: An Introduction* (London; New York: Routledge).
- Floridi, L. 2002, "What Is the Philosophy of Information?" *Metaphilosophy*, 33(1-2), 123-145.
- Floridi, L. 2003a, "Information" in *The Blackwell Guide to the Philosophy of Computing and Information*, edition, edited by L. Floridi (Oxford - New York: Blackwell), 40-61.
- Floridi, L. 2003b, "Two Approaches to the Philosophy of Information", *Minds and Machines*, 13(4), 459-469.
- Floridi, L. 2004a, "Informational Realism", *ACS - Conferences in Research and Practice in Information Technology (Computers and Philosophy 2003 - Selected Papers from the Computer and Philosophy conference CAP 2003)*, 37, 7-12.
- Floridi, L. 2004b, "Open Problems in the Philosophy of Information", *Metaphilosophy*, 35(4), 554-582.
- Floridi, L. 2004c, "Outline of a Theory of Strongly Semantic Information", *Minds and Machines*, 14(2), 197-222.
- Floridi, L. 2005a, "Information, Semantic Conceptions Of ", *Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/information-semantic/>.
- Floridi, L. 2005b, "Is Information Meaningful Data?" *Philosophy and Phenomenological Research*, 70(2).
- Floridi, L. forthcoming, "The Logic of Being Informed", *Logique et Analyse*.
- Floridi, L., and Sanders, J. W. 2004, "The Method of Abstraction" in *Yearbook of the Artificial. Nature, Culture and Technology. Models in Contemporary Sciences*, edition, edited by M. Negrotti (Bern: Peter Lang),
- Floridi, L., and Sanders, J. W. forthcoming, "Levellism and the Method of Abstraction".
- Fodor, J. A. 1975, *The Language of Thought* (New York: Thomas Y. Crowell).
- Fodor, J. A. 1987, *Psychosemantics* (Cambridge, Ma: MIT/Bradford).
- Giere, R. 1988, *Explaining Science: A Cognitive Approach* (Chicago: University of Chicago Press).
- Giere, R. N. 1999, "Using Models to Represent Reality" in *Model-Based Reasoning in Scientific Discovery*, edition, edited by L. Magnani, N. J. Nersessian, and P. Thagard (Dordrecht: Kluwer),

- Greco, G. M., Paronitti, G., Turilli, M., and Floridi, L. 2005, "How to Do Philosophy Informationally", *Lecture Notes in Computer Science*, 3782, 623-634.
- Grim, P., Mar, G., and St. Denis, P. 1998, *The Philosophical Computer* (Cambridge Mass.: MIT Press).
- Harms, W. F. 1998, "The Use of Information Theory in Epistemology", *Philosophy of Science*, 65(3), 472-501.
- Harnad, S. 1990, "The Symbol Grounding Problem", *Physica Scripta*, D(42), 335-346.
- Harnad, S. 1993, "Problems, Problems: The Frame Problem as a Symptom of the Symbol Grounding Problem", *Psycoloquy*, 4(34).
- Haugeland, J. 1981, *Mind Design : Philosophy, Psychology, Artificial Intelligence* (Cambridge, Mass ; London: MIT Press).
- Haugeland, J. 1997, *Mind Design II : Philosophy, Psychology, Artificial Intelligence* Rev. and enl. (Cambridge, Mass. ; London: MIT Press).
- Hofkirchner, W. (ed.) 1998, *The Quest for a Unified Theory of Information : Proceedings of the Second International Conference on the Foundations of Information Science* (Amsterdam: Gordon & Breach).
- Israel, D., and Perry, J. 1990, "What Is Information?" in Hanson [1990], pp. 1-28.
- Kamp, H. 1984, "A Theory of Truth and Semantic Interpretation" in *Truth, Interpretation and Information*, edition, edited by J. Groenendijk, T. M. V. Janssen, and M. Stokhof (Dordrecht: Foris),
- Landauer, R. 1987, "Computation: A Fundamental Physical View", *Physica Scripta*, 35, 88-95.
- Landauer, R. 1991, "Information Is Physical", *Physics Today*, 44, 23-29.
- Landauer, R. 1996, "The Physical Nature of Information", *Physics Letter*, A 217, 188.
- Landauer, R., and Bennett, C. H. 1985, "The Fundamental Physical Limits of Computation", *Scientific American*, July, 48-56.
- Larson, A. G., and Debons, A. (ed.) 1983, *Information Science in Action : System Design. Proceedings of the Nato Advanced Study Institute on Information Science, Crete, Greece, August 1-11, 1978* (The Hague: M. Nijhoff).
- Li, M., and Vitányi, P. M. B. 1997, *An Introduction to Kolmogorov Complexity and Its Applications* 2nd (New York: Springer).
- Losee, R. M. 1997, "A Discipline Independent Definition of Information", *Journal of the American Society for Information Science*, 48(3), 254-269.
- Machlup, F., and Mansfield, U. (ed.) 1983, *The Study of Information : Interdisciplinary Messages* (New York: Wiley).
- Mandik, P. 2002, "Synthetic Neuroethology" in *Cyberphilosophy: The Intersection of Philosophy and Computing*, edition, edited by T. W. Bynum and J. H. Moor (New York - Oxford: Blackwell), 11-29.
- Minsky, M. L. 1967, *Computation: Finite and Infinite Machines* (Englewood Cliffs, NJ: Prentice Hall).
- Minsky, M. L. 1990, "Logical Vs. Analogical or Symbolic Vs. Connectionist or Neat Vs. Scruffy" in *Artificial Intelligence at Mit, Expanding Frontiers*, edition, edited by P. H. Winston (Cambridge, Ma: MIT Press),
- Muskens, R., and al. 1997, "Dynamics" in *Handbook of Logic and Language*, edition, edited by J. Van Benthem and A. Ter Meulen (Amsterdam: Elsevier), chap. 10.

- Newell, A. 1980, "Physical Symbol Systems", *Cognitive Science*, 4, 135 - 183.
- Newell, A., and Simon, H. A. 1976, "Computer Science as Empirical Inquiry: Symbols and Search", *Communications of the ACM*, 19 113-126.
- Pylyshyn, Z. W. 1984, *Computation and Cognition* (Cambridge, Ma: MIT/Bradford).
- Ringle, M. 1979, *Philosophical Perspectives in Artificial Intelligence* (Atlantic Highlands N.J.: Humanities Press).
- Sayre, K. M. 1976, *Cybernetics and the Philosophy of Mind* (London: Routledge & Kegan Paul).
- Searle, J. R. 1990, "Is the Brain a Digital Computer?" *Proceedings and Addresses of the American Philosophical Association*, 64, 21-37.
- Seligman, J., and Moss, L. S. 1997, "Situation Theory" in *Handbook of Logic and Language*, edition, edited by Van Benthem J. and Ter Meulen A. (Amsterdam: Elsevier), chap. 4.
- Shannon, C. E. 1993, *Collected Papers.*, edited by N. J. A. Sloane and A. D. Wyner (New York: IEEE Press).
- Shannon, C. E., and Weaver, W. 1949 rep. 1998, *The Mathematical Theory of Communication* (Urbana: University of Illinois Press). Foreword by Richard E. Blahut and Bruce Hajek.
- Smith, B. 2004, "Ontology" in *The Blackwell Guide to the Philosophy of Computing and Information*, edition, edited by L. Floridi (New York - Oxford: Blackwell), chap. 12
- Smolensky, P. 1988, "On the Proper Treatment of Connectionism", *Behavioral and Brain Sciences*, 11(1), 1-23.
- Suppe, F. 1989, *The Semantic Conception of Theories and Scientific Realism* (Urbana, Ill.: University of Illinois Press).
- Suppes, P. 1960, "A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences", *Synthese*, 12, 287-301.
- Suppes, P. 1962, "Models of Data" in *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, edition, edited by E. Nagel, P. Suppes, and A. Tarski (Stanford: Stanford University Press), 252-261.
- Taddeo, M., and Floridi, L. 2005, "Solving the Symbol Grounding Problem: A Critical Review of Fifteen Years of Research", *Journal of Experimental and Theoretical Artificial Intelligence*, 17(4), 419 - 445.
- Turing, A. M. 1936, "On Computable Numbers, with an Application to the Entscheidungsproblem", *Proceedings of the London Mathematics Society, 2nd series*, 42, 230-265. Correction published in Vol. 43 (1936), pp. 544-546.
- Turing, A. M. 1950, "Computing Machinery and Intelligence", *Minds and Machines*, 59, 433-460.
- van Benthem, J. 2003, "Logic and the Dynamics of Information", *Minds & Machines*, 13(4), 503-519.
- van Benthem, J., and van Rooy, R. 2003, "Connecting the Different Faces of Information", *Journal of Logic, Language and Information*, 12(4), 375-379.
- Van Fraassen, B. 1980, *The Scientific Image* (Oxford: Clarendon Press).
- van Gelder, T. 1995, "What Might Cognition Be, If Not Computation?" *Journal of Philosophy*, 92, 345-381.

- van Gelder, T., and Port, R. (ed.) 1995, *Mind as Motion: Explorations in the Dynamics of Cognition* (Cambridge, Ma: MIT Press).
- Weaver, W. 1949, "The Mathematics of Communication", *Scientific American*, 181(1), 11-15.
- Wiener, N. 1948, *Cybernetics or Control and Communication in the Animal and the Machine* 2nd ed. (Cambridge, Ma: MIT Press).
- Young, R. A. 2004, "Wittgenstein's Tractatus Project as Philosophy of Information", *Minds and Machines*, 14(1), 119-132.

Notes

¹ Berkeley [1732], Dialogue 1, Section 5, Paragraph 6/10.

² For a longer and more detailed discussion see Floridi [2004b].

³ The definition is first introduced in Floridi [2002]. The nature and scope of PI are further discussed in Floridi [2003b] and Floridi et al. [2005]. Floridi [2003c] provides an undergraduate level introduction to PI.

⁴ A typical life cycle includes the following phases: occurring (discovering, designing, authoring, etc.), processing and managing (collecting, validating, modifying, organising, indexing, classifying, filtering, updating, sorting, storing, networking, distributing, accessing, retrieving, transmitting etc.) and using (monitoring, modelling, analysing, explaining, planning, forecasting, decision-making, instructing, educating, learning, etc.).

⁵ See Adams [2003] for a reconstruction of the informational turn in philosophy and Young [2004] for an analysis of Wittgenstein's philosophy of information.

⁶ In 1964, introducing his influential anthology, Anderson wrote that the field of philosophy of AI had already produced more than a thousand articles (Anderson [1964], p. 1). No wonder that (sometimes overlapping) editorial projects have flourished. Among the available titles, the reader may wish to keep in mind Ringle [1979] and Boden [1990], which provide two further good collections of essays, and Haugeland [1981], which was expressly meant to be a sequel to Anderson [1964] and was further revised in Haugeland [1997].

⁷ Chalmers [online] seems to believe that pancomputationalism is empirically falsifiable, but what he offers is not (a) a specification of what would count as an instance of x that would show how x is not to be qualified computationally (or information-theoretically, in the language of this paper) given the nomic characterisation N of the universe, but rather (b) just a re-wording of the idea that pancomputationalism might be false, i.e. a negation of the nomic characterisation N of the universe in question: "To be sure, there are some ways that empirical science might prove it to be false: if it turns out that the fundamental laws of physics are noncomputable and if this noncomputability reflects itself in cognitive functioning, for instance, or if it turns out that our cognitive capacities depend essentially on infinite precision in certain analog quantities, or indeed if it turns out that cognition is mediated by some non-physical substance whose workings are not computable." To put it simply, we would like to be told something along the lines that a white raven would falsify the statement that all ravens are black, but instead we are told that the absence of blackness or of ravens altogether would, which it does not.

⁸ This chapter is based on Floridi [2002], Floridi [2004b] and Floridi [2005a]. I wish to acknowledge the kind permission by Blackwell and by the Stanford Encyclopedia of Philosophy to reproduce parts of the texts from those publications.