The Philosophy of Information, some broader background

This is a working document by Pieter Adriaans, meant to illustrate the editors' ideas about the broader thrust of this volume, and make the entreprise a bit more grounded in philosophy. It is not intended, however, as a required 'ideology' for our authors.

The prefix 'the philosophy of ...' is a great tool to create conversational items at parties. Even the most whimsical topics (motorcycle maintenance, surrealism, Winnie the Pooh) look more interesting from a philosophical perspective. But it is not this kind of intellectual enhancement we are looking for with our phrase 'Philosophy of Information'. Some issues have been around in philosophy from the start.

The idea that the notion of 'information' might have philosophical relevance seems to be rather new. 'Information' does not occur as a separate entry in the index of the eight-volume Encyclopedia of Philosophy edited by Paul Edwards in 1967. The same holds for the well-known History of Logic written by Kneale and Kneale that first appeared in 1962. Actually the history of the term is rather complicated. In antiquity and the middle ages the concept of 'informatio' occurred as a technical notion in the work of such diverse authors as Cicero and Aquinas. In the 15th century the term 'information' appears to have found its way via influence of the French language to colloquial speech in various European countries. In the 20th century three crucial and interrelated developments shaped the modern scientific concept of 'information': 1) The mathematical definition of information in terms of probability of a message, 2) The definition of the bit as the fundamental unit of information, 3) The association of proof with computation as a sequence of well-defined events in the physical world.

When we read pre-Socratic philosophers like Zeno or Parmenides with this modern mind, we feel uneasy about the undifferentiated mix of formal, epistemological, ontological, ethical and esthetical questions. It is all there, but without the distinctions. The same holds for the notion of information. With hindsight one could say it has played a role in philosophy from the beginning, without being recognized as such. If we look at Kant's famous three questions: 'What can I know?' 'What must I do?' and 'What can I hope for?', then the first question at least partially involves a reflection on information. It is also immediately clear that a philosophy of information never can replace the whole of philosophy, since the concept of information gives us at best very limited assistance when pondering the second and third question.

According to Kant philosophy should not reach for knowledge of the transcendent but it should be transcendental: i.e., it should study necessary a priori conditions for the possibility of knowledge. Modern information science offers fascinating perspectives on such a program for a methodological foundation of the sciences. Even if somewhat removed from Kant's original ambitions, the modern view is transcendental in the sense that it can formulate a priori conditions for the possibility of the growth of scientific knowledge with unprecedented mathematical precision and clarity.

To be sure, substantial reflections on information fed by modern science made their way into mainstream philosophy in the seminal work of Carnap and others in the 1950s, and in a second wave, with celebrated publications by Dretske, Perry, and others in the 1980s and 1990s. So far this process has mostly affected particular themes in epistemology and the philosophy of language. We think it can be taken much further still, when a broader set of scientific results and philosophical issues is brought together. If one adopts 'the computational view' a number of new and powerful solutions to age old philosophical problems present themselves.

An example: if one lives in a world in which events are generated by computational processes, then the Solomonoff-Levin distribution (or universal distribution) assigns an a priori probability to events. This probability is related to their computational complexity. This gives us an entirely new perspective on Hume's induction problem, the study of heuristic search and the analysis of human creativity and learning. Another example: the concept of 'cognition as computation' allows us to formulate a partial a priori answer to the question of what can be known. In order for something to be knowable it must be computable. We have deep results on what can be computed and what not. We also know a good deal about the complexity issues involved. A third example: modern logic studies epistemic logics, dynamic logics, non-well-founded set theory, update logics, belief revision systems and a myriad of related formal systems that give us an entirely new perspective on the questions of what can be known and what we could believe and how these questions are interrelated.

Of course the question whether the 'computational paradigm' is correct is a matter of philosophical debate. This debate however can never be conducted properly without a thorough analysis of the philosophical issues involved. It would also be superficial to address these problems without a deeper understanding of the role of information and

computation in various sciences like physics, mathematics, biology, linguistics and cognitive science. In each of these disciplines 'information' and 'computation' play a vital role, but prima facie in very different ways. There seems to be no direct route from the study of the abstract notion of information to practical results in these disciplines. The study of information has a definite empirical component as it takes a different guise in various parts of reality: physical, biological, social or psychological. At the same time progress in modern science depends critically on the use of computers to manipulate large quantities of data, to facilitate co-operative work between groups of scientists and to calculate the consequences of complex models of substructures of the world around us. This last observation gives a very practical motivation for a foundational study of 'the computational paradigm'. In an even broader perspective it is clear that the use of computers has a profound influence on our cultures and our societies. Computers change the way we communicate and the way we work. They affect our art and our science and ultimately the way we think about ourself. In this context one could define the philosophy of information as:

1) philosophical analysis of the concepts of 'information' and 'computation',
2) philosophical analysis of the role of information in nature, science and culture,
3) analysis of the notion information in the traditional philosophical disciplines like methodology of science, epistemology, ontology, ethics and aesthetics.

A complete survey of these issues would be an ambition well beyond our capabilities. This volume wants to be a first contribution to a proper rethinking of some philosophical problems in the light of recent developments in the study of the concept of 'information'. It's aim is to be rather comprehensive on the first item on the list and to present elements of an analysis of the role of the computational paradigm in various sciences. Cultural issues and a deeper technical analysis of the concept of 'information' in relation to traditional philosophical disciplines will be casually discussed, but a proper treatment would require several other volumes.

The central motivation for the philosophical study of information can be summarized as follows. There is a philosophy of information because philosophy is in a sense *about* information. When we reason about the structure of reality, or about the beauty of music or the rightness of an action the information dimension is always present at the background. In this sense information is an essential phenomenon that is always found in every philosophical reflection. It was there when Zeno and Parmenides wrote

their first tentative analyses. It was there when Kant formulated his transcendental program. It was even around when people invented writing or when our ancestors learned to speak. I am pretty certain all those people never thought about motorcycle maintenance or Winnie the Pooh, although I am not sure about surrealism ;-).

Finally, this broad analysis also leads to some more concrete questions that we hope to get clearer on as the volume takes shape. We list only a few for now:

Probability Information is defined in terms of probability. There are various theories of probability (extrapolation of observed frequeny, probability of a theory given data, etc). How universal is our view of probability? What follows for the notion of information? Does the latter give rise to specific new views on probability?

Computation How universal are current models of computation as manipulation of macroscopic objects localized in space and time? What is the link between theories of computation and physics? Does the rise of quantum computing affect this view?

Epistemology How does theory of information translate into traditional epistemology, methodology of science, heuristics?

Learning Why is the world learnable: i.e., why do we live in a universe in which things can be learnt? This is really the question of complexity. We have a reasonable theory of (un)computability but no good theory that explains the efficiency of computational of agents in the real world. Why do we apparently live in regions that are dominated by benign probability distributions. (This leads to relations with problems in AI, and work of Dreyfuss, Husserl, etc.) Is there one unified model of learning or is there a whole taxonomy of relatively unrelated learning systems? Given all this, what are adequate formal models of learning?

Logic What are the consequences of putting information and multi-agent processes of communication, update and belief change at centre stage, rather than meaning and deductive inference? How does this change the agenda of the philosophy of logic, which is still largely dominated by inherited set issues from the foundational era? Can we get a better grasp of notions like expressiveness of a language, a taxonomy of sorts of natural communication arranged by function and complexity, and perhaps Gödel-like deeper (im-)possibility theorems for complete communication of information?

We will update this list of themes in our eventual editorial as the Book takes shape.