

The following text is very much an outline, prepared by just one of the editors after minimal interaction. It has a 'logician's slant', while Pieter must still add an 'algorithmic slant'. So, we have broken our promise of the March Workshop, which was to provide you with a well thought-out definition of information that can unify the Handbook, agreed upon by the two editors – and while we are at it, we have not kept to the promised deadline either. Even so, Pieter and I hope that the following can serve as an 'umbrella' for the book, and of course, we will develop it a lot further. But for the moment, your comments and suggestions would be quite welcome!

DRAFT EDITORIAL Handbook of the Philosophy of Information

Information Is what it Does

Johan van Benthem

The title of this editorial is reminiscent of David Lewis' famous analysis of the concept of 'meaning'. Do not ask what meanings *are*, but what they are supposed to *do*. And if there are some Marxian/Feuerbachian undertones to the title as well, in that we *are* what we *eat*, informationwise, then that may also have a kernel of truth.

1 History

Some history of the term 'information' will be put here, using parts of what Keith, Mike, and Pieter have already written. It will probably be short, just a page or two.

2 Just a Metaphor?

'Information' may be a wide-spread term, and even a wide-spread phenomenon, and yet it might be just a *metaphor* leading to vague philosophy, just like that of a 'system' or, in its softer senses, a 'game'. But the real situation seems more interesting. As with terms like 'energy' or 'money', there is indeed a general usage where little can be said beyond great generalities. Energy is what drives inanimate processes and animate activities, and what allows us to relate the effort involved in different cases of both. Money is what makes transactions possible without always having to bring the house that we sell on our backs to hand it over. But in both

cases, the general usage is backed up by pockets of more precise use in expert circles, and moreover, the latter uses are often grounded in mathematical theory: thermodynamics, or economics. The resulting interplay of broader and narrower usage seems to cause no problems to anyone. We understand the general usage, and we specialize and precisify when needed. These lessons can be transferred to other notions. That 'money' leads the way here would not even be a bad thing, if we recall Karl Marx' famous dictum that 'Logic is the Currency of the Mind'. An undefined slogan perhaps, but rich and suggestive! Indeed, when Keith Devlin says jokingly to lay audiences that "information is the tennis ball of communication", he actually formulates a very similar role for information as for money, viz. as the abstract currency that gets transferred when people say or observe things. And he also gets the idea right that information usually arises in complex multi-agent settings, where interaction is of the essence.

3 A Family of Related Notions?

There seems to be a serious issue whether information can stand on its own in a more philosophical analysis of the concept. Compare the case of *knowledge*. Most standard philosophical analyses, whether mainstream like Plato's or more avant-garde like Dretske's or Nozick's make it part of a cluster of notions, involving also truth, belief, justification, and perhaps even counterfactuals. We are not isolating concepts, we are charting their most intimate connections and clans – in short, the state of their affairs.

Candidates for such a clan of related concepts – not identical but intertwined – in our case are: *information*, *probability*, *complexity*, *meaning*. Nevertheless, our Handbook sticks to information as our main theme, with the others as side-themes throughout.

4 One, or Many Concepts of Information?

One might adopt as a null-hypothesis that there exists just one notion of information. But scrutiny of the landscape also reveals a dissimilarity with the cases of 'energy' or 'money'. For, there are competing respectable mathematical views of information! This diversity is not an unknown phenomenon. Carnap proposed a famous dichotomy between *Probability-1* for objective frequency and *Probability-2* for subjective chance, and this still corresponds to a major 'duality' in

both the mathematical and philosophical foundations of probability and the practice of statistics. Indeed, Carnap also tried something similar with 'information' in the early 1950s, when confronted with Shannon's work, and trying to relate that to his own work on logical information spaces. But his Information-1 versus Information-2 never caught on the way his views on probability did. Still, there may be a lot to it. Indeed, in the case of information as we see it, even more varieties may be needed. Here are three streams in our Book – the first two correspond roughly to 'Carnap' versus 'Shannon':

<i>Information-A</i>	Knowledge, logic, what is conveyed in informative answers
<i>Information-B</i>	Probabilistic, information-theoretic, measured quantitatively
<i>Information-C</i>	Algorithmic, code compression

We do not mean to say that these represent opposing camps or competing 'paradigms', but the three types do correspond to natural clusters of chapters – and it makes sense to chart the varieties. Later on, we will analyze these categories of 'information' a bit further, but for now, let's just make some quick preliminary observations.

First, these lines are not disjoint: they share important features. E.g., A, B are both *relative* to situations, while C seems to be more absolute, but it may be relative to the source of the code? Also, in a way to be discussed later, classical Information Theory, a B-type enterprise, shares important concerns with the logical A-type approach.

Next, the three strands proposed here are *not* the familiar *qualitative/quantitative* distinction. The latter does not seem very illuminating – and in any case, all of A, B, and C can have more quantitative and more qualitative versions. Just think of the basic representation theorems in Information Theory which relate qualitative and numerical notions of information. Of course, there is still a substantial issue of what it means to have quantitative measures of information, and the way in which a 'one-dimensional' number compresses the content of a message. More on that below.

Finally, our three streams have interesting counterparts, even in the case of probability. After all, in addition to Carnap's two views of that notion, there was one he omitted, viz. von Mises' famous intuition of randomness, defining probability via random sequences. This may be viewed as a third algorithmic notion, say, Probability-3 – especially, in the mathematical form which this took through recursive and other place selections, in the work of Martin-Löf, van Lambalgen, and others. In fact, the analogy philosophy of information/philosophy of probability may be worth exploring further.

5 Bottom-up Approach: Compare Stable Uses

One can analyze the notion of information 'top-down', looking at the phenomenon in its entirety, trying to find some umbrella definition capturing some essential features. One can also start 'bottom-up' with existing uses in the technical literature that have stable features. And then, roughly the same three stances emerge as earlier on. The following discussion draws on my paper 'Two Logical Concepts of Information', to appear in a forthcoming commemorative volume for Jon Barwise. Jon was a logician who, together with John Perry, has done much to put the concept of information on the map in logic and linguistics, drawing upon Dretske's work. He was also one of the founders of Stanford's pioneering Center for the Study of Language and Information.

(i) *Information as range, and reduction of uncertainty*

The first notion in much of the literature is this. Information may be encoded in a range of possibilities: the different ways the real situation might be. For instance, at the start of a card game, the range consists of the different possible deals of the cards. More dynamically, new information is that which reduces my current range – that is: more information = smaller range. This is the standard logical sense of information in which a proposition P updates the current set of worlds W to $\{w \text{ in } W \mid w \text{ makes } P \text{ true}\}$. This notion is relative to what we might call a 'logical space' in Carnap's sense, a particular scheme for describing the options. It is also relative to agents, since the update happens to what they know about the world. This is the main notion of information used in our Handbook chapters by McCarthy, Baltag/van Ditmarsch/ Moss, Groenendijk/Kamp/Stokhof, and Rott. Once again,

It is about agents' logical spaces of alternative options for some yet unknown

actual situation (the latter is what the information is 'about'), and new information typically has to do with update actions changing the current state.

The next pervasive notion in the literature may be called

(ii) *Information as correlation*

On this view, one situation carries information about another if there is a suitable correlation connecting the two. This is the sense in which dots on my radar screen carry the information that there are airplanes out there. Note that the information may be there, even though there is no agent to pick it up. Unlike what Procol Harum sings, "signposts" do not "cease to sign" when there are no humans left on our planet. In philosophy, this sense of information is central to the work of Dretske, and Barwise & Perry who all claim to have it ultimately from Shannon's information theory. This is the view in our chapters by Bais & Doyne Farmer, Devlin, Dretske, Kelly, Seligman, Topsoe & Harremoes, and I guess quite a few others. Some of its key features are:

Information is crucially about something, and thereby a relative notion: a relation between a receiving situation and a described/sending situation, where the 'quality' of the information depends on the reliability of the correlation (as in 'smoke means fire', a favourite example through the ages)

Two views, or one? As hinted at in Section 4, the views (i) and (ii) are not necessarily antagonistic. In what follows we take *Shannon's Information Theory* (*IT*, for short) as a running thread for discussing this and other major issues. First, *IT* itself may be viewed as combining features of both options (i) and (ii). It is definitely about reducing uncertainty, be it in a quantitative manner asking for the *average* reduction of uncertainty, summarizing many possible update actions. But is also about correlation between a source and a receiver across a channel. In fact, this double aspect of *IT* suggests that notions (i) and (ii) may be analogous, and it encourages us to find out just how. For instance, we might think of the two 'situations' in notion (ii) as agents in some extended sense. And although notion (ii) and *IT* share a special emphasis on analyzing *channels* for information flow, this could really be part of the logical update apparatus in (i) as well: its absence is just a fluke of the current research agenda – and the first studies on channels in a logic-based framework topic are coming up. And finally, combinations of logical spaces

and *IT*-style correlation measures co-exist happily in some modern semantics for natural language: cf. the work by Robert van Rooij on questions and answers, or Rohit Parikh on general messaging.

Eventually, the better view may be that (i) and (ii) are aspects of one and the same notion of information – but this may still develop as this Editorial gets rewritten. Note that Dunn's chapter is already taking this broad line, since all notions in this Section make sense, and occur, in computer science.

Numerical versus qualitative But perhaps most famously, *IT* *quantifies* the notion of information. At first sight, this seems a sharp divide. Scientists and engineers love this idea, since we can now 'compute with information'. Philosophers and logicians often hate it, because all the many-dimensional content of an informational event is 'flattened' into a one-dimensional number. Messages with totally different content can become equivalent in this way. So far, the editors have slightly different takes on this. Both agree that both notions of information have value. But Johan would start from the logical notion, and ask *just when* can we get by with the impoverished *IT* account of information, foregoing deeper logical analysis of content? Pieter would point out that in most realistic scenarios involving informational events, logical micro-descriptions are either unavailable, or the cost of computing them becomes prohibitive. In that case, the statistical view *the only way we have* of finding essential features of the relevant process – and despite the one-dimensionality of the bit measure, it shows the same 'unreasonable effectiveness' that mathematics has in general.

Probability and recurrence Another major feature of *IT* is its probabilistic feature of defining an 'average reduction' of uncertainty. If we read these probability values as *frequencies* (Carnaps' Probability-1), then they refer to a repeated process of sending the same message over the same channel. This longer-term temporal aspect is lacking from (i) as stated above, which apparently describes one-shot information. Likewise, (ii) apparently describes not average correlation, but *specific* links between situations on single occasions. But eventually, the difference is not so great. When thinking of a genuine 'informational connection', recurrence of events over time is essential. A mere juxtaposition of two individual events, or of two true propositions, does not constitute real 'information flow'. The latter notion also involves robustness under recurrence, re-usability, and so on. So, again, this

temporal feature would be shared by both notions so far. Incidentally, an a-temporal reading of the probabilities would be as subjective chances (Carnap's Probability-2). Then we get closer to a-temporal logical views of information and knowledge, where adding subjective probability would allow for more fine-grained beliefs about the reliability of the incoming proposition.

So far, so good. But we also have chapters (Adriaans, Grunwald & Vitanyi) with yet another central notion of information, which we might call

(iii) *Information as structure*

Kolmogorov's algorithmic complexity does correspond to some further legitimate uses of 'information', having to do with the amount of compression in a string. E.g. there is a sense in which we find a mathematical paper 'more informative' than a sports report of the same length, since its content can hardly be compressed, whereas that of the sports match can. I feel much less confident about writing up the major features of this notion of information – and Pieter will take over from me later at this point. There are at least some obvious questions of comparison with our earlier features:

Algorithmic information computed for some string is information *about what?*

Is there a channel involved: maybe the computational coding/decoding device?

And eventually, we need to see how this fits with the earlier two (or perhaps just one...) notions. For instance, Grunwald & Vitanyi do develop strong analogies between algorithmic complexity and the *IT* view of information. And likewise, Pieter's work on learning uses algorithmic complexity – while on the other hand, learning is very close to core tasks in the logical framework, viz. information update and belief revision. Samson Abramsky's chapter may also be relevant here, once we have all of it.

To be continued.

6 The Process Stance: Information in Action

Next, we turn to a perspective shared by both editors, which is crucial to the title of this editorial. We borrow an idea from computer science – and philosophy, for that matter. In a computational perspective, it makes little sense to talk about data structures in isolation from the processes that manipulate them, and the tasks which these are supposed to perform. The same point was made in philosophy, e.g. by David Lewis, when saying that 'Meaning Is what Meaning Does'. We can only give good representations of meanings for linguistic expressions when we state how they are going to be *used*: in communication, disambiguation, inference, and so on. In a slogan:

Structure should always come in tandem with a process!

Now, which processes are relevant to the notion of information? In the Handbook, we find that authors almost always have some key process in mind for which their notion of information is supposed to work. In fact, there seems to be a great variety of paradigmatic informational processes, and this should certainly be highlighted:

communication, learning, computation, inference,
interaction (as in games), update, revision, ...

Some of these processes concern activities of single agents, whereas others are intrinsically multi-agent 'social' phenomena. Here are some questions about this list, which need to be made more systematic in the eventual editorial:

What does the notion of information *do* for each process?
Is the basic informational process a multi-agent one, with
single-agent activities a 'one-dimensional projection'?

E.g., is 'learning' as in formal learning theories just a one-agent projection of a shared activity of a two-agent system {Learner, Teacher}? Likewise, is the usual logician's notion of formal 'proof' as a formal string of symbols just the zero-agent projection of a multi-agent interactive activity of argumentation?

Again, our central concern of Section 5 will return here.

Does the same notion of information serve all
these informational processes? And should it?

One option would be to compare the three different approaches to information found in Section 5 on the same tasks in the current list. How would a logician go about them, and an information-theorist, and an algorithmician? This was the red thread in a recent issue of the Journal of *Logic, Language and Information* on 'Logical Theories of Information' – and it worked rather well. Such tasks might include:

question answering,
communication in general,
learning,
games,
interaction in general, etc.

Finally, the tasks in this Section should not be just a mere list. Again, some more general patterns can be discerned in our material. For instance, key tasks range from *single-agent* update or learning to intrinsically multi-agent tasks like communication, and eventually to the longer-term *interactive* setting of games, and players with inter-locking goals and strategies. In this connection, it would be useful to have some essentials of *game theory* covered, or at least mentioned, in the Handbook, since interaction is a basic phenomenon that comes up naturally in several chapters (update logic, linguistic semantics and pragmatics, human interaction). A more systematic perspective on the essential informational processes and activities, corresponding to major cognitive tasks, still needs to be added to this editorial.

To be continued.

7 Conclusion: Grand Unification or Co-Existence?

The purpose of this Handbook is not necessarily reductionist. The book can be a success even when no Grand Unification is reached. Here are the goals that we do have, in ascending order of ambition:

- (a) Let different bona fide traditions *meet*,
- (b) Find a common language: e.g., by making proposals for standardizing terminology,

- (c) Suggest new research topics by comparing concerns.

It seems clear to us that the book will be a success along these lines. Examples of this potential for communication and integration abound, when reading our chapter drafts. Here are some items, out of many:

- introduce a *multi-agent* perspective in correlationist and *IT* views,
- introduce explicit *channels* in the logical view,
- introduce enriched notions of *update* from the logical side (cf. the *BMD* and Rott chapters) into the *IT* perspective,
- move from interactive computation to *games*.

and explore broader consequences, such as a possible Universal Game for interactive information, just as the Universal Turing Machine for older notions of information.

The most ambitious goal of all might be developing one grand unified view of information in this Handbook. The editors do not expect that to happen, and they feel strongly that the book will *not* be a failure if we don't. Understanding and appreciating diversity, and creating channels for meaningful communication will be just as fine.

PS Remaining issues

What the Handbook does not cover We will briefly mention some potentially relevant topics that have been left out, such as information economics, game theory, and so on.

Main thrust of the chapters We will explain, perhaps in short separate introductions to the Parts, what our chapters do and do not claim to achieve. E.g., the final Part is not about exhaustive coverage of all relevant aspects of information in the natural and social sciences. It rather highlights some major themes – and for the rest, we hope that our authors will provide sufficient references to other themes and viewpoints.

Themes increasing coherence throughout the Handbook At our March Workshop, we also discussed a kind of 'grid structure', identifying major themes connecting different chapters, such as learning, interaction, or games. These themes will

correlate with the topics of Sections 6, 7 above – but things still have to be written up that way.