

On the proper definition of information

András Kornai

ETHICOMP 2008

Abstract

Whether viewed as a Rawlsian primary good or as the building material of an entire Infosphere, the notion of *information* plays a fundamental role in the development of CE. We analyze this notion from the monosemic perspective of lexical semantics (Ruhl 1989) and offer a definition that provides a suitable basis for CE work in that it is not restricted to (true) propositional content and derives, rather than asserts, the ethical value of information.

1 Introduction

The proper definition of information is a complex matter, yet we often encounter drastic oversimplification. For many, only Shannon information can truly be called information, for many others (e.g. Himma 2008) information must be true before it merits the name ‘information’. In the first part of this paper we examine the notion of information from the perspective of ordinary language philosophy with special emphasis on the monosemic nature of the concept, and in the second part we begin to draw some more technical conclusions from the abstract notion thus arrived at. We take as our raw material a number of ordinary language expressions ranging from the broadly agreed upon *zip files contain the information in compressed format* to the highly controversial *information wants to be free* and seek for a single, unified definition that covers both everyday and technical usage.

Monosemic analysis is best understood in contrast to the the classical polysemic method of lexical semantics (Katz and Fodor 1963) whereby separate meanings such as *bachelor*₁ ‘unmarried adult man’, *bachelor*₂ ‘fur seal without a mate’, *bachelor*₃ ‘knight serving under the banner of another knight’, and *bachelor*₄ ‘holder of a BA degree’ are assigned to what are viewed as disjoint units of meaning. In monosemic analysis (also called ‘Saussurean’ and ‘Columbia School’ approach in Kirsner 1989, who calls the polysemic approach ‘cognitive’) we search for a single, general, and more abstract meaning such as ‘unfulfilled in typical male role’. This is not the place to fully compare and contrast the two approaches (Kirsner’s work offers an excellent starting point), but we note here a significant advantage of the monosemic approach, namely that it makes interesting predictions about novel usage, while the predictions of the polysemic approach border on the trivial. To stay with the example, it is possible to envision novel usage of ‘bachelor’ to denote a contestant in a game who wins by default (because no opponent could be found in the same weight class or the opponent was a no-show). The polysemic theory would predict that not just seals but maybe also penguins without a mate may be termed ‘bachelor’ – true but not very revealing.

The choice between monosemic and polysemic analysis need not be made on a priori grounds: even the strictest adherent of the polysemic approach would grant that ‘bachelor’s degree’ refers, at least historically, to the same kind of apprenticeship as ‘bachelor knight’. Conversely, even the strictest adherent of the monosemic approach must admit that the relationship between ‘obtaining a BA degree’ and ‘being unfulfilled in a male role’ is no longer apparent to contemporary language learners. That said, we still give methodological priority to the monosemic approach because of the original Saussurean motivation: if a single form is used, the burden of proof is on those who wish to posit separate meanings. One of the claims articulated in this paper is that ordinary language use of ‘information’ need not be carefully distinguished from the technical (Shannon) use. Rather, the Shannon theory is quite faithful to the ordinary meaning, and captures many, but not all, significant aspects of it, the same way as the Euclidean geometrical notions of line, circle, or triangle offer a rational reconstruction of the corresponding everyday notions. We will take the first steps toward another rational reconstruction of ‘information’, one that takes into account those aspects that the Shannon definition does not cover. In particular, we are interested not just in messages transmitted over (noisy) channels, but also in information phenomena that were not (and could not have been) much in Shannon’s purview six decades ago: *genetic information* and *computer languages*.

Our central tool in search of abstract meaning will be the following syllogism: if XY is an ordinary language expression that presupposed X being in class C , and x can be used in place of X , then x belongs in class C . For example, X *wants something* in ordinary discourse presupposes the subject X to be capable of having wants and desires, so *information wants to be free* presupposes that information can have wants and desires. This is in sharp contrast to X *should be Y* which presupposes only that X is currently not in state Y but it can attain this state. In this particular case, usage overwhelmingly supports the former implication: Google provides over 121,000 hits for *information wants to be free* and less than 7,000 for *information should be free*. This is not to say that we must wholeheartedly endorse the notion that information can have wants and desires, or even to make the methodological claim that issues of this sort are amenable to resolution by the “wisdom of crowds”. At the same time, it does not bode well for a substantive discussion if we must dismiss 95% of those who express an opinion as a lunatic fringe, or, what is the same, there is something deeply worrisome about any theory that forces us to ignore 95% of our data.

2 Information is a precious fluid

Everyday use of the term *information* seems to rely on a metaphor (Lakoff and Johnson 1980) of the *precious fluid*, related both to ordinary liquids like water and to less tangible but nevertheless appealing notions like chi (life force), energy, electricity, etc. Information can *flow* from a source to a recipient over a *channel*, *pipe*, or other conduit, can be *diluted*, *compressed*, or *stored* in vessels of specific capacity, and so on. All this usage rests on a fluid/gas metaphor, and we shall talk of fluids, rather than gases, because information sources *dry up* rather than *empty up*, the information will *trickle out* rather than *hiss out* etc. Classical information theory provides a rational reconstruction of the fluid metaphor, with the *bit* as the fundamental volumetric unit from which other units such as channel capacity (bitrate) are derived by standard dimensional analysis (in this case, as bits/sec). This reconstruction turns the metaphor into a reliable tool for engineering communications devices, storage area networks, and the like.

While the practical success of Shannon theory is in and of itself a strong justification for this theory, Himma (2008) attributes too much resolving power to “mathematical attempts to formalize the defini-

tion of information” by claiming (p 3) that these “presuppose a semantic concept of information”. In fact, as noted in Fig. 1 of Floridi 2005, the opposite is true: Shannon’s definition stands completely neutral as to whether a piece of information is true or false. The key notion, *entropy*, is simply a measure of compressibility according to a precisely defined coding scheme, and information theory gives the exact same results for, and remains equally applicable to, encoding false messages as true ones. Lewis Carroll’s *Alice*, a clear instance of fiction, is just as amenable to transmission over a channel as his *Symbolic Logic*, and it would be simply wrong to say that the former, being false, requires zero bits (has zero information content) while the latter, being true, has positive information content. The ordinary language expression *one can learn a great deal about human nature from Shakespeare*, with the standard implication that Shakespeare’s works contain information, makes eminent sense, even though strictly speaking all the propositional content in Shakespeare is false.

In fact, we must view with skepticism the idea of deriving a semantic (truth-based) theory of information from any mathematical formulation, since the entire body of mathematics, including mathematical logic, stands neutral on what is true and what is not. Mathematics is perfectly capable of formulating alternative theories that cannot all be true. The universe is either 4-dimensional or, as string theory suggests, perhaps 10-dimensional, but the mathematical theory of 17-dimensional spaces is just as good as that of 19-dimensional spaces even though one of them is guaranteed to be false.

For CE, the main shortcoming of the Shannon theory is that it leaves the “precious” aspect of the notion unexplored: from the engineering perspective any data stream is information, and it is quite unclear how one could ascribe any value, let alone an ethically central value, to bits. Yet animistic statements like *our higher being is pure information* or *information wants to be free* abound in the CE discourse, indeed, for many people these statements provide the main reason for engaging in the discourse in the first place. To capture these ideas one must ascribe to information attributes like sentience or consciousness that will make many serious students of information ethics cringe. However, the ability to discuss matters that are not uniformly believed by all participants is an essential prerequisite of any meaningful debate, so we must seek a formulation that makes such arguments expressible, even if we do not fully share the underlying assumptions.

To use a non-computational example, the pro-choice/pro-life debate is not about definitions: the issue is not whether the fetus fits the proper definition of ‘human’, the issue is whether we should accord it human rights. Insisting on one’s own definition is simply a form of claiming the rhetorical high ground, and succeeds only in cutting off the debate. Here our goal is not to defend the claim that information has its own consciousness or at least its own identity, sentience, and volition, but rather to articulate the underlying implications, since it is these implications that we really care about.

2.1 Identity

It is one of the most striking aspects of fluids that they lack identity: pouring two glasses of wine in a pitcher results in a mixture from which the original varieties can no longer be recovered, and subparts (drops) of one fluid are indistinguishable from one another. From the perspective of everyday ontology, fluids are best viewed as mereological entities, infinitely divisible, the polar opposites of individuals (which are, by definition, indivisible). H₂O is H₂O no matter where we find it – information, on the other hand, can generally be traced to specific *sources* or *leaks* based on its composition. In fact, there is a whole field of human endeavor, philology, that is largely devoted to developing methods that enable tracing various pieces of information to their sources. Since the pieces are individuated by cutting the bitstream at carefully selected points, we conclude that *linear organization* is a defining characteristic of information. The strings AB and BA are not identical, even if their components are the same. In

addition to preservation of linear order, the everyday notion of identity preservation requires also the preservation of suprasegmental/hierarchical organization (intonation in particular) and annotation of the source: *duo cum dicunt idem, non est idem*. When we say *the minutes preserved all relevant information from the meeting* it is implied that they make clear who said what.

2.2 Sentience

To say that faster than light travel would make *Einstein turn in his grave*, or that *the principles of special relativity are violated* by it is to say that a specific theory (some kind of structured set of axioms/information entities) is capable of sensing something that goes against it. We do not have to take a naive animistic view of theories whereby theories (or by standard rhetorical exaggeration, even the dead bodies of their creators) can sense actual events, just as we do not have to believe in Poseidon to believe that the sea is a dangerous place, it is enough for theories to be sensitive to statements of events. In other words, a theory imposes, by means of deductive rules, a classification on statements, which can (i) follow from the axioms (ii) their negation can follow from the axioms or (iii) be undecidable relative to the axioms. Thus we conclude that information entities come, implicitly or explicitly, with rules of composition (deduction). This is certainly weaker than the traditional (Bentham) sense of sentience, which requires the capacity for suffering: no claim is made, at the ordinary language usage level, that individual pieces of information are capable of suffering, just that structured ensembles of information (theories) can. The axioms are just a component of these theories, and the suffering, such as it is, is evidently sensed by the other component, the deductive apparatus. The pain metaphor, with its implied sentience, is widely used when we dismiss theories *on pain of triviality* or *on pain of contradiction*.

2.3 Volition

Fermat's Last Theorem has *acquired* a proof – the Riemann Hypothesis is still *in need* of one. This assumption will *bite* you in the ass. Everyday expressions like the preceding make sense only if we take it for granted that information entities have desires and powers on their own, the way genies in bottles do. The whole meme trope is meaningless without ascribing not only sentience (deductive capabilities) but also powers to change their environment, to structured sets of information. The absolute minimum that we need to posit is the capability of some piece of information to *override* some other piece, and a rational reconstruction of this notion will of necessity invoke some version of default or defeasible logic. But there is at least one form of information, the genetic code, that is capable of much more: replication, self-repair, and, by means of lumbering protein machines, exercising far-reaching control over its environment. Again we do not need to take a stance on the whole selfish gene debate, whether humans are fully controlled by their genes, or where the limits of their free will are, all that is required for ordinary language to make sense is the assumption that human consciousness is not necessarily the *only* source of control, other, possibly overriding, volitions may coexist with it.

2.4 Reverence

The Axiom of Choice needs to be *approached with care*. One must *respect* the words of the sages. To the extent that information has life-like attributes it is natural that we may venerate it. It is evident that genetic information has life-like qualities – in fact, at our current level of understanding, it is quite hard to disentangle the genetic code from the very notion of life. (We note in passing that blood, also a precious fluid, stands in an equally strong metaphoric relationship with life and genes.) In fact, it seems

quite in keeping with the Aristotelian notion to ascribe entelecheia to genes, and one should consider whether other forms of information also have entelecheia. One form that clearly meets the definition is logia, and it is impossible to make sense of Western civilization from Heraclitus to John 1.1 without assuming widespread belief in, and significant adherence to, sacred propositions that the human will must be subordinated to.

Scientific theories, artistic expression, and other memes share the fundamental properties of genes and logia, and to the extent that information is viewed as the building material of the memes, the reverent attitude we see toward information, especially highly structured information, is explained by the same logic of behavior that dictates veneration of all forms of life. We do not make the normative statement that information, be it genetic or memetic, be it attributed to God, sages, geniuses, or ordinary mortals, *should* or *must* be venerated, we simply state the empirical fact that it *is*. Not only are there a sufficient number of people who regard various pieces of information with reverence, but in fact most people do so – those who hold nothing sacred or precious are often considered sociopaths, outside the realm of rational discussion. Therefore, any definition of information must include its power to inspire reverence in many, for it clearly has such power, even if some are immune to it.

3 Making the metaphor work

Where does this notion of information as a precious fluid leave us? As we said at the outset, the fluid part is well captured in Shannon's theory of information: what we need is an equally attractive rational reconstruction of being precious, sacred, a worthy object of reverence. But before turning to this issue, let us compare the pretheoretical notion of information that has emerged from the analysis of ordinary language, PF information, to the better accepted theoretical notion of semantic information (SI), as presented e.g. in Floridi (2005). In addition to the *maps* considered by Floridi, there are a host of information-bearing objects in everyday use: *books, theories, natural language utterances, computer code, genetic sequences, and databases* leap to mind – the list is not intended to be exclusive. Certainly the ordinary (or even technical) language utterance *this X contains (useful, good, misleading, ...) information* makes sense for all of the above, and there is no a priori reason to suppose that it makes different sense for different forms of information. Yet the SI definition, being restricted to propositional content, is not capable of treating the above construction in a homogeneous manner, since neither genetic nor computer code has propositional content.

In contrast, the PF criteria are met rather uniformly. The fact that our non-propositional examples, genetic and computer code, have identity and linear organization of the kind described in 1.1 is hardly debatable. In databases the linear structure is less than evident, but still present in the order db rows were created (an order that is heavily relied upon e.g. in unrolling transactions). Even in axiomatic theories, where the axioms are often considered to be in an unordered conjunction for logical purposes, their presentation order is generally rigidly maintained by tradition (as with Newton's Laws or the Laws of Thermodynamics). Maps are inherently two-dimensional, but the hierarchical (constituent) structure is still very much in evidence. Turning to sentience in the narrow sense of 'deductive apparatus' used in 1.2, it is clear that the genome, computer programs, and axiomatic theories must be surrounded by some kind of deductive or interpretative apparatus to get their full sense: genes are expressed by a complex biochemical mechanism, program code is executed on hardware, and the real meaning of the axioms is only revealed by deduction. For example, the Peano axioms do not *state* that $13+8=21$, though this is clearly implied by them. That natural language expressions (including logia) require an interpretative apparatus (grammar) of their own goes without saying, and maps have their own (vector or raster) semantics.

As for 1.3-4, whether these forms of information truly exhibit volition, or whether it is just our animistic manner of talking that ascribes them volition is hard to say (and the more formal definition introduced below stays neutral on this issue) but there does not seem to be any deeper reason to treat them non-uniformly in this regard, especially as various forms of information are largely interchangeable, for example a map can be generated by a program, or a gene sequence can become a database entry. Similarly, whether we treat any form of information as precious is an ethical decision we do not intend to dictate here, but to the extent we are prepared to venerate some we should be prepared to venerate all (though possibly to different degrees, just as we are reluctant to perform certain experiments on humans that we are willing to perform on mice).

The greatest advantage of insisting on propositional content is that conjunctions of propositions define states of affairs that we can assign probabilities to, and thus provide the hooks for computing information content via Shannon's definition. However, the genetic code does not express any state of affairs (let alone a true one), it is simply a collection of recipes for building proteins and regulating their interactions. The same can be said of computer programs, which also correspond to plans of actions rather than to states of affairs, and of many natural language expressions (in particular, imperatives) commonsensically considered informative. But not all is lost if propositional content is abandoned: the genetic code, computer code, and all forms of information are still subject to algorithmic compression, and a more general notion of compressibility/information content, Kolmogorov complexity (KC), is still available, filling the same role for the PF theory that Shannon entropy fills for the SI theory. The technical details of KC are beyond the scope of this paper (for a short introduction see Kornai 2008 Ch 7, and for an encyclopedic treatment see Li and Vitanyi 1997) but the main idea is simple: the KC of an object is the length of the shortest program that can generate it.

Let us briefly turn to the *information wants to be free* slogan to see how the PF approach would handle it. The empirical fact we wish to explain is that this slogan, practically meaningless under the SI approach, can nevertheless drive the behavior of multitudes of people who engage in file sharing, encryption cracking, and a host of similar activities at considerable risk. For better or worse, this slogan is seen by many as having greater moral force than all laws and regulations that prohibit file sharing or encryption cracking. How can this be? Free can mean two things, *gratis* and *libre*, and as most people in the community, we consider the second sense paramount (though historically, the original statement by Stewart Brand was clearly in the first sense). To be free thus means to be free of obligations, not just pecuniary ones. To the extent that information is viewed as having entelecheia, identity, sentience, and volition, the real question is whether we *want* such entities to exist in the world without obligations. The commonsense answer is that we do not, but this is an ethical choice rather than the consequence of some cleverly chosen definition. It is precisely because we can well imagine that information of the computational kind can acquire the characteristics of the genome (self-replicating, and gaining a measure of control over its environment) that we wish to saddle it with some obligations from the get-go.

Alternative viewpoints, such as the idea that we must respect the will of any living-like being, are certainly possible. But on the whole, humanity as a higher entity with its own entelecheia does not put the quest for liberty over everything else. We do not grant full self-government to nations (insisting that at the very least they show respect for the environment, refrain from mass destruction, and empower their own citizens), to smaller political or societal units, or even our own children. To the contrary, we expand considerable individual, familial, and societal effort on endowing our children with a sense of obligation, and on making sure that the overall structure is set up so as to minimize the harm those lacking moral scruples are capable of doing.

One key aspect of this undertaking is to create a deductive apparatus capable of distinguishing right from wrong. Elsewhere (Kornai 2008 Ch 6) we suggested using Ginsberg's (1986) System D as a foundation for lexical semantics, and here we put this system to use in describing PF information more formally. There are many notions that we will have to leave unanalyzed, in particular, we cannot within the confines of this paper fully formalize notions such as Agent, Action, Being, Event, Space, Time, or Thing. However, we use these terms in a meaning as close to their ordinary language meaning as possible. We assume discrete time and space only for expository convenience.

There are only concrete Things, each with two lists of attributes: those considered essential and those considered accidental. Essential attributes include a name, and possibly other unchangeable features, accidental features include a location, and possibly other changeable features. An Event is simply the change of some features of one or more Things, typically in a single time tick. We pass over the philosophical difficulties of defining which changes in the attribute list of which Things can be properly individuated as Events, not because the question is trivial, but because it does not impinge greatly on the PF concept of information. Some Things are capable of effecting change in the state of other Things (possibly including, or limited to, themselves), these are called Agents. Our notion of agency is weak, including the sea, wind, and other non-volitional forces that are nevertheless capable of effecting change. Things capable of observing and memorizing the state of other Things are called Beings. The state of their memory is called their (roll of drums) Information.

It is not assumed that Beings are necessarily capable of observing the Information state of other Beings. Their observation capabilities may vary considerably, e.g. being only partially able to observe the location of another Thing or Being (seeing the x coordinate but not the y coordinate), and their memory capacities, deductive powers etc. may also show variation. Some may be capable of reasoning with abstract models involving classes of things, others may be restricted to reasoning about the specific Things in their view. We may add to the model some completely passive indicators (Things tied to a location and displaying a single attribute such as temperature or free energy level), and declare that locations with a higher value of some indicator are more desirable. Beings which are Agents may be capable of moving themselves (changing their location attribute) or even of moving other Things or Beings. Moving another Being to a location with lower energy is considered bad – moving it to a location with higher energy is considered good. Moving oneself, or moving Things that are not Beings is considered morally neutral.

Going one step further, we may eliminate energy, temperature, and other external resources from consideration altogether, saying simply that destructively changing the Information of another Being is bad, adding to it in a way the recipient considers not bad is good. Notice that there is no provision here for the Information of a Being to be in any sense true, nor is there a provision for it being propositional: Beings, being von Neumann machines, can mix program instructions and data in their memory any way they wish. However, we do begin to see how ethical value is attributed to bits: changing some other Being's Information is grievous bodily harm, since the very identity of Beings (aside from some superficial parameters like name and location) is concentrated in the Information they hold.

How much a model of this sort is capable of making a uniform moral code (such as Gewirth's PGC) emerge remains to be seen, but it seems clear that a first step in making information free (within limits) is to experiment with setups such as the one described here to see where the difficulties in endowing Beings with Action capabilities are.

4 Conclusions

While ordinary language philosophy is considered outmoded, this paper makes the case that in situations where the underlying ontological issues are hard to disentangle it can still make a valuable contribution. In particular, the ordinary language view of information is more friendly to information stored in logia, computer programs, genomes, and other active (life-like) forms than the received (SI) view. By abandoning the restriction to propositional content (or, in programming terms, abandoning the idea that memory must store data to the exclusion of instructions) we arrived at a more nuanced (PF) view of information, defined simply as the memory content of some being capable of some actions (changes in its environment and itself), having some sentience, and capable of some memorization and execution of earlier memorized instructions.

The idea that information is a Rawlsian primary good can now be tied to the ethical justification for protections against GBH, rather than being postulated anew. Highly controversial statements such as *our higher being is pure information, the Infosphere is built entirely from information, or information wants to be free* are now analyzable in a model that does not prejudge the issues. Shannon information is not negated, but seen as a special case arising from specific assumptions (namely, prefix-freeness) concerning the class of codes allowed. There are many subtle difficulties both with Kolmogorov complexity (which is not Turing-computable) and multi-valued logic (necessary for the modeling of beings that can store contradictory pieces of information, see Belnap 1977), but the PF view, in spite of the mysticism at the roots of ordinary language, supports both formal analysis and large-scale simulation.

References

- Belnap, N. 1977: How a computer should think. In Ryle (ed) *Contemporary Aspects of Philosophy*. Newcastle upon Tyne: Oriel Press 30–56
- Floridi, L. 2005: Semantic Conceptions of Information. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/information-semantic> accessed April 2008
- Ginsberg, M.L. 1986: Multi-valued logics. In Ginsberg (ed) *Readings in Non-monotonic Reasoning*. Morgan Kaufman, San Mateo CA pp 251–255
- Himma, K. 2008: Information and Intellectual Property Protection: Evaluating the Claim that Information Should be Free. *APA Newsletters in Philosophy and Computers*, forthcoming.
- Katz, J. and Fodor, J.A. 1963: The structure of a semantic theory. *Language* 39 170–210
- Kirsner, R.S. 1993: From meaning to message in two theories. *Cognitive and Saussurean views of the Modern Dutch demonstratives*. In Geiger and Rudzka-Ostyn (eds) *Conceptualizations and mental processing in language*. Berlin: Mouton de Gruyter. 81–114
- Kornai, A. 2008: *Mathematical Linguistics*. Berlin: Springer Verlag
- Lakoff, G. and Johnson, M. 1980: *Metaphors we live by*. Chicago, IL: University of Chicago Press
- Li, M. and Vitanyi, P. 1997: *An introduction to Kolmogorov complexity and its applications*. Berlin: Springer Verlag.
- Ruhl, C. 1989: *On Monosemy: A Study in Linguistic Semantics*. SUNY Press, Albany