

Outline of a Theory of Truth as Correctness for Semantic Information

Luciano Floridi

Research Chair in Philosophy of Information and GPI, University of Hertfordshire
Faculty of Philosophy and IEG, University of Oxford

Abstract: The article develops a correctness theory of truth (CTT) for semantic information. After the introduction, in section two, semantic information is shown to be translatable into propositional semantic information (*i*). In section three, *i* is polarised into a query (*Q*) and a result (*R*), qualified by a specific context, a level of abstraction and a purpose. This polarization is normalised in section four, where $[Q + R]$ is transformed into a Boolean question and its relative yes/no answer $[Q + A]$. This completes the reduction of the truth of *i* to the correctness of *A*. In sections five and six, it is argued that (1) *A* is the correct answer to *Q* if and only if (2) *A* correctly saturates (in a Fregean sense) *Q* by verifying and validating it (in the computer science's sense of "verification" and "validation"); that (2) is the case if and only if (3) $[Q + A]$ generates an adequate model (*m*) of the relevant system (*s*) identified by *Q*; that (3) is the case if and only if (4) *m* is a proxy of *s* (in the computer science's sense of "proxy") and (5) proximal access to *m* commutes with the distal access to *s* (in the category theory's sense of "commutation"); and that (5) is the case if and only if (6) reading/writing (accessing, in the computer science's technical sense of the term) *m* enables one to read/write (access) *s*. The last section draws a general conclusion about the nature of CTT as a theory for systems designers not just systems users.

Keywords: Computer Science; Correctness; Information Theory; Philosophy of Information; Semantic information; Truth

1. Introduction

In recent years, philosophical interest in the nature of information has been increasing steadily.¹ In particular, one of the current debates concerns the veridical nature of semantic information. The debate is somewhat old,² but has been re-ignited by the proposal to analyse semantic information in terms of *well-formed*, *meaningful* and *veridical data* (Floridi [2004b]). Admittedly, the analysis – according to which semantic information encapsulates truth, exactly as knowledge does – has attracted some criticisms for being too restrictive.³ Such criticisms, however,

have been proved unjustified.⁴ As a result, there is now a growing consensus about the following approach.

Semantic information is primarily understood in terms of *content about a referent*. I shall discuss the formal nature of content in the following pages but, at the moment, suffice to say that it is analysable in terms of *well-formed* and *meaningful data*.⁵ Strings or patterns of data may constitute sentences in a natural language, but of course they can also generate formulae, maps, diagrams, videos and other semiotic constructs in a variety of physical codes, being further determined by their appropriate syntax (well-formedness) and semantics (meaningfulness). By "about a referent" one is to understand the ordinary and familiar way in which some well-formed and meaningful data, constituting semantic information, concern or address a topic. Following Dretske (1981) and

¹ For an updated overview and guide to the literature see Floridi (2004a).

² For example, Bar-Hillel and Carnap (1953) and Devlin (1991) argued against the veridical nature of semantic information, whereas Dretske (1981) and Grice (1989) argued in its favour.

³ See for example the discussion in Fetzer (2004), with a reply in Floridi (2005a); or the objections moved by Colburn (2000a), Colburn (2000b) and Dodig-Crnkovic (2005).

⁴ Floridi (2007) and Sequoia-Grayson (2007).

⁵ On the analysis of data see Floridi (2008a).

Dretske (1988), one may easily recognise this “aboutness” feature in propositional attitudes such as “Mary is informed that there is some beer in the fridge”, where “being informed” is used in the *statal*⁶ sense, i.e., in the sense that Mary holds that information. This is the condition into which *a* enters once *a* has *acquired* the information (*actional* state of being informed) that *p*. It is the sense in which a witness, for example, is informed (holds the information) that the suspect was with her at the time when the crime was committed. In the rest of this paper, we shall be concerned with only this standard, *statal* and epistemically oriented concept of semantic information.

In Floridi (2005a), I argued that a definition of semantic information in terms of alethically-neutral content – that is, strings of well-formed and meaningful data that can be additionally qualified as true or untrue (false, for the classicists among us), depending on supervening evaluations – provides only necessary but insufficient conditions: if some content is to qualify as semantic information, it must also be *true*. One speaks of false information in the same way as one qualifies someone as a false friend, i.e. not a friend at all. This leads to a refinement of the initial definition into:

[DEF] *p* qualifies as semantic information if and only if *p* is (constituted by) *well-formed, meaningful and veridical data*.

[DEF] captures the general consensus reached by the debate and mentioned at the outset of this section. According to it, semantic information is, strictly speaking, inherently *truth-constituted* and not a contingent *truth-bearer*, exactly like knowledge but unlike propositions or beliefs, for example, which are what they are independently of their truth values and then, because of their truth-aptness, may be further qualified alethically.

⁶ The distinction is standard in linguistics, where one speaks of passive verbal forms or states as “*statal*” (e.g. “the door was shut (state) when I last checked it”) or “*actional*” (e.g. “but I don’t know when the door was shut (act)”). In this paper, I deal only with the *statal* sense of “is informed”. This is related to cognitive issues and to the logical analysis of an agent’s “possession” of a belief or some knowledge.

[DEF] offers several advantages. For example, it plays a crucial role in the solution of the so-called Bar-Hillel-Carnap Paradox (Floridi [2004c]) and provides a necessary element for a subjectivist theory of epistemic relevance (Floridi (2008c)). Here, it is worth emphasising that it forges a robust and intuitive link between semantic information and knowledge. More specifically, the veridical thesis contained in [DEF] corresponds to the one characterising the definition of knowledge. Taking advantage of this parallelism, one may rely on the ordinary apparatus of modal logic (e.g. Chellas (1980)) to formalise “*a* is informed that *p*” as $I_a p$, and hence formulate the veridicality thesis of semantic information in terms of the so-called veridicality axiom $\Box \phi \rightarrow \phi$, also known as **T**, **M** or **K2**, thus:

[VT] $I_a p \rightarrow p$

The intended interpretation of [VT] is that *a* is informed that *p* only if *p* is true. In Floridi (2006), I have shown that information logic (*IL*) can then be satisfactorily modelled in terms of an interpretation of the relation “*a* is informed that *p*” based on the axioms of normal modal logic **B**. [VT] associates *IL* to epistemic logics (*EL*) based on normal modal logics **KT**, **S4** or **S5**. And it differentiates both *IL* and *EL* from doxastic logics (*DL*) based on **KD**, **KD4** and **KD45**, since, of course, no *DL* satisfies the veridicality axiom. It follows that *IL* allows truth-encapsulation (i.e., it satisfies [VT]) without facing either epistemic or doxastic collapse, i.e., merely morphing into another epistemic or doxastic logic. So knowledge encapsulates truth because it encapsulates semantic information, which, in turn, encapsulates truth, as in a three dolls matryoshka.

Despite its advantages, any approach endorsing [DEF] raises two major questions (Floridi (2004b)). One is upstream:

a) what does it mean for semantic information to be truthful?

The other is downstream:

b) how can semantic information turn into knowledge?

Both questions are prompted by [DEF] but neither is specifically about [DEF] only, so each fails to provide a starting point for a *reductio ad absurdum*. They are rather information-theoretical versions of classic conundrums: (a) is a request for a theory of truth and (b) is a request for a substantive analysis of knowledge. Since the goal of this paper is to seek to answer only (a), let me brush (b) away by adding a final clarification.

[DEF] nests semantic information into knowledge so tightly that one is naturally led to wonder whether anything else might be missing, in order to escalate from the weaker to the stronger phenomenon, and hence between their corresponding concepts. Indeed, the threshold can be so fine that one may often overlook it and thus fail to distinguish between the two propositional attitudes, treating “Mary is informed that there is some beer in the fridge” and “Mary knows that there is some beer in the fridge” as if they were always losslessly interchangeable. In everyday life, this might be the norm and the conflation is usually harmless: it can hardly matter whether the bus driver is informed or knows that the traffic light is red. Philosophically, however, the distinction captures an important difference, and hence it is important to be more accurate. For it takes only a moment of reflection to see that one may be informed (hold the information) that p without actually knowing that p . Not only because holding the information that p does not have to be a *reflective* state (although it is not necessarily the case that $I_a p \rightarrow II_a p$, one may also object that $K_a p \rightarrow KK_a p$ is notoriously controversial as well) but also because, even when it is, it might still arguably be *opaque* and certainly *unjustified*.

Consider opacity first. It is open to discussion whether a messenger carrying (whether in her memory or in her hand it does not matter) an encrypted message p that she does not understand – even if she is informed that she carries p – may be said to hold the information that p . On the one hand, one may argue that she is not genuinely informed that p . On the other hand, one may retort that, if she can deliver the information that p (and we are assuming that she can) then she can legitimately be said to hold that information.

The interesting point here is not to solve the dispute, but to note that the dispute itself is reasonable, whereas, if the same messenger knows that p , there can be no doubt that she must also comprehend the information carried by p . It might be open to debate whether holding the information that p is necessarily a non-opaque state, but such a dispute would be pointless in the case of knowing that p .

Next, consider the degree of justification. Epistemic luck does not affect informativeness negatively. To see why, one may use a classic Russellian example: if one checks a watch at time t and the watch is broken but stopped working exactly at $t - 12$ hours and therefore happens to indicate the right time $t - 12$ at t , one still holds the information that the time is t , although one can no longer be said to know the time. The same applies to a more Platonic example in which a student memorises, but fails to understand, the proof of a geometrical theorem: she is informed (holds the information) that the proof is so and so, but does not really know that the proof is so and so. Generalising, Russell- Plato- or Gettier-type counterexamples may succeed in degrading “knowing” to merely “being informed” (“holding the information that”), but then “being informed” is exactly what is left after the application of such counterexamples and what remains resilient to further subjunctive conditionalization. The additional difficulty is that the counterexamples show both that some justificatory variable might have a key role to play in full epistemic states, besides reflectivity and transparency, and that this variable too is still insufficient to guarantee the delivery of knowledge every time. Sometimes, one may be (reflectively and transparently) informed that p and fully justified in holding the information that p and yet still fail to know that p .

Rotten as all this may be, it is not all, for there is further bad news. One can also prove that Gettier-type problems are logically unsolvable by showing that they are a subclass of the more general “coordinated attack” problem, which is demonstrably insolvable in epistemic logic (Floridi [2004b]). This entails that the tripartite account is not merely inadequate as it stands, as proved by Gettier-type counterexamples, but demonstrably irreparable in principle, so that efforts to

improve it can never succeed. Although it is useful to know that we should stop trying to fix this approach and start looking for a different one, the disappointing conclusion is that, as far as question (b) above is concerned, we lack even a promising strategy to upgrade I_{ap} to K_{ap} .

So much regarding (b), which we can now leave on one side. Prospects are much brighter when it comes to question (a). In this case, the challenge is not a shortage, but rather an overabundance of viable answers, since we are spoiled for choice by a variety of theories of truth.⁷ Admittedly, in the literature on semantic information there appears to be at least an implicit predilection for some version of a Tarskian and/or correspondentist approach.⁸ And yet, at least in principle, nothing prevents each of the major theories of truth from answering (a). They simply would have been refuted a long time ago if they couldn't. It follows that some initial tolerance towards a pluralistic approach to (a) might be unavoidable, if not methodologically welcome. Of course, if this were all that one could sensibly recommend about (a), there would be little reason to pursue any further investigation. There is, however, another way of approaching (a), which opens up an interesting line of enquiry.

Consider the strategy sketched above. It consists in selecting the best available theory of truth and testing how well it might be applied and adapted in order to explain the truthfulness of semantic information. With some negligible adjustments, such a top-down approach is comparable to the so-called "design pattern" technique (Gamma et al. (1995)) in software engineering (Sommerville (2007)). This consists in identifying and specifying the abstract features of a design structure, which are then generally reusable solutions to commonly occurring problems in the construction of an artefact. In our case, we have several design patterns for the concept of truth. We know that they are robust,

because they have been tested and refined since Ramsey, if not Aristotle. We also know that they are reusable: although they have been developed to deal primarily with propositional or sentential truths, one may reasonably expect them to be effectively adaptable to truthful data (e.g. a truthful map) as well. So, when our artefact, i.e. semantic information, is proved to require the particular feature of being truthful, a sensible alternative is to consider such design patterns and try to identify the ones that best satisfy the constraints and requirements imposed by the development of the artefact itself. Oversimplifying, one may answer (a) by choosing whichever pre-packaged theory of truth turns out to be most suitable. This strategy may be classic, is certainly viable but it is hardly innovative and I shall not pursue it in the following pages.

The other approach is bottom-up and suggests the sort of strategy that will guide the rest of this investigation. It consists in assuming the artefact itself as given, and then trying to discover the principles governing its properties and workings by analysing its structure, function and operations. In software engineering, this technique is known as "reverse engineering". This is "the process of extracting the knowledge or design blueprints from anything man-made" (Eilam (2005), p. 3). It consists in examining an existing artefact in order to identify its components and their interrelationships and hence create representations of it in other forms or at a higher level of generalization. Following this strategy, one may answer question (a) by assuming the occurrence of some semantic information and then disassembling it in order to reveal what its components are and how they interact with each other to deliver information. We have the artefact and we seek to understand its mechanism by taking it apart, hopefully in the right way and places. Note that this second strategy is perfectly compatible with the first, once it is realised that there is a virtuous cycle of feedback between design patterns and reverse engineering results. Contrary to the first strategy, however, reverse engineering promises to deliver a more innovative analysis, as it avoids approaching the problem of truth from pre-established theories and explores it

⁷ In this paper, I have relied especially on Lynch (2001), Engel (2002) and Künne (2003), among the many introductions and anthologies available on the major theories of truth, as particularly helpful.

⁸ See for example Popper (1935), Dretske (1981), Fox (1983), Israel and Perry (1990), Barwise and Seligman (1997) and Bremer and Cohnitz (2004).

from a new perspective. After all, the first strategy merely retrofits some already existing theory of truth to semantic information, instead of trying to develop a customised solution which may then be generalisable. The cost to be paid for this innovation is that our bottom-up strategy will also be uphill, if I may be allowed to combine the two metaphors: it is much more economical to choose from a pre-established menu than to develop a new approach. I can only hope that the reader will find the effort rewarding and the result enlightening. And now it is time to start climbing.

2. First Step: Translation

A large variety of kinds of semantic information, from traffic lights to train timetables, from road signs to fire alarms, falls within the scope of [DEF]. This is how it should be but it is awfully inconvenient for our purposes. For in order to reverse-engineer semantic information in such a way that its components might easily be identified, disassembled and explained, it would be far easier and more fruitful to concentrate on just one kind, the propositional one, which lends itself to such a treatment straightforwardly. So, our first step will be to ensure that all kinds of semantic information covered by [DEF] are indeed translatable into propositional semantic information, thus guaranteeing that what will be concluded about the latter may be extendable to the former. At this point, the reader who finds such “translatability” uncontroversial, or indeed trivial, may wish to skip the rest of this section. The one who finds it impossible may concede the restriction of scope as a matter of convenient stipulation, although the rest of this section purports to show that the burden of proof is on her shoulders. As for the rest of us, what follows should be sufficiently convincing to make our second step unproblematic.

Syntactically (or in terms of information theory), the propositional translatability of any kind of semantic information is unquestionable and a matter of daily experience. After all, analogue information is reproducible digitally to any chosen degree of accuracy, its digital version is equivalent to finite lists of zeros and

ones, and these can be further encoded into as many answers to questions asked in a suitably chosen language, and hence ultimately translated into statements of that language. That doing any of this would be sheer madness is irrelevant here. For the question is not how difficult or costly this process would be, e.g. in terms of accuracy, time and memory resources, but that it might be possible at all. More to the point is whether some non-propositional, semantic information – the sort of information provided by the map of the London Underground, for example – may always be translatable *semantically* into propositional semantic information, at least in principle. Mind, not all of it at once, and not even part of it at every level and for every kind and degree of detail (henceforth Level of Abstraction or simply LoA, Floridi (2008b)), but any of it at the right LoA, depending on needs and requirements. Since the difference between a syntactic and a semantic translation may not be very familiar, let me first introduce it with an example.

Consider being able to reproduce the map of the London Underground on graph paper by being told, say over the phone, the position and colour of each square on the paper: the communication over the phone would provide a syntactic translation, with the end result (the coloured graph paper representing the map) constituting a test about whether the translation worked. Contrast it now to being able to travel from one station to another on the London Underground, by receiving verbal instructions from someone who is navigating using the visual indications provided by the map. This is a semantic translation, and your trip is a test of its accuracy.

Suppose now that a semantic translation from non-propositional into propositional information, of the kind just illustrated, were sometimes impossible, even in principle. Then there would be some residual semantic information, conveyed non-propositionally (e.g., by the map), that one would necessarily be unable to convey propositionally, independently of the resources available. We would then have reached the limits of the informational powers of any natural language, even natural languages formally expanded, e.g. mathematically. Allegedly, we should still be able to point to the information in question

(in the previous example, suppose we are both looking at the same map), but we would be unable to generate the right sort of propositional content that could adequately convey it. This is a *reductio ad absurdum*. For here we are not engaging with some Wittgensteinian limits of the “sayable”, with Kantian *noumena*, with some linguistically-ungraspable sensations, or some mystical experience enjoyed while looking at the map of the London Underground. We are talking about what the map of the London Underground can encode, in terms of information about travelling through the network, positions of the stations, interconnections, available routes etc., which, allegedly, would be at least partly beyond the expressive power of any natural language to convey. But since natural languages have been acknowledged to be “semantically omnipotent” at least since Leibniz (Formigari (2004), pp. 91-92), one can arguably assume that the translation is always possible, even if it is likely to be onerous at times and hence often unfeasible in terms of resources. So, in the rest of the paper, we shall treat semantic information as possibly semiotic-dependent (it may always require a code) but not as semiotically bounded (codes are translatable propositionally, if expensively resource-wise) or, more formally and briefly:

$$[\text{TR}] \quad \forall x (\text{DEF}(x) \wedge \text{Non-prop}(x)) \rightarrow \exists y (\text{Prop-t}(y, x) \wedge \text{DEF}(y))$$

The intended interpretation of [TR] is that, if any data (the domain on which the quantifiers range) satisfy [DEF] but are not propositional, then there is a propositional translation of those data which also satisfies [DEF]. Note that we do not need to assume the stronger principle of translational equivalence: pictures may be worth thousands of words, but there might be thousands of words that are priceless. All that [TR] needs to guarantee is that the conclusions reached about the alethic nature of propositional semantic information will be exportable to the truthful nature of non-propositional semantic information as well. In other words, that what can be concluded about the truth of “there is some beer in the fridge” is equally applicable to the truthfulness of a picture conveying the same information visually.

3. Step two: Polarization

Once some information i is formulated propositionally, the second step is to follow a standard approach, in information theory, to the quantification of information, and disassemble i into a combination of a query Q and a result R . A query is to be understood as a request for data sent (e.g., an illocutionary act performed) by a sender to a receiver, in the form of a message. Thus, it might have the format of a question (“where is the beer?”) as well as of an imperative (“tell me where the beer is”), or a string of symbols in a search engine. A result is also to be understood as a message, the requested data, sent by the receiver to the querying sender. In short, we have (the asterisk is a reminder that the formula is provisional and will have to be refined):

$$[\text{POL}^*] \quad i = Q + R$$

That [POL*] is always achievable is warranted by the fact that any propositional i is equivalent to a message, and that any message is a combination of querying and resulting data encoded in the same set of symbols of the chosen language (alternatively: every p can be transformed into a request of whether p plus a result, but more on this in the next section). The polarization of i into $Q + R$ offers several advantages. We shall exploit four of them.

First, [POL*] highlights the need to specify the *context* (C) in which, the *level of abstraction* (LoA) at which, and the *purpose* (P) for which the query is formulated and hence it is expected to be satisfied by the result. For the sake of simplicity, below I shall refer to the combination of these three parameters by means of the acronym CLP. The first two requirements were stressed by Austin (1950). “Where is the beer?” is asked by someone in some specific circumstance (the context), by relying on a specific granularity of discourse or detail, what I have called LoA. In our example, there might be no beer (if no beer has been purchased) or, if the sender of the query knows that some beer has been purchased, answering that “the beer is somewhere” would amount to a joke or a mistake in the choice of LoA, if the sender wishes to know the precise location of the

beer, e.g. left in the car or carried inside the house or placed in the fridge. The third requirement was stressed by Strawson (1964). LoAs are always teleological and queries are formulated (results are offered) for a purpose, even if the purpose might be implicit. In the example, one may wish to make sure that the beer has been placed in the fridge and not left in the car, for example. To recall a Fregean point, queries cannot acquire their specific meaning in isolation or independently of their CLP parameters. It is a bit of a pain, but we need to keep these variables in mind, lest the conceptual mess becomes unmanageable. So, as a memory aid, let me revise [POL*] by adding a combined index, thus:

$$[\text{POL}] \quad i^{\text{CLP}} = [Q + R]^{\text{CLP}}$$

A second advantage of the polarization of i into $Q + R$ is that it makes evident the role of R , which is to *saturate* Q , to adapt another Fregean idea lately borrowed by information theory.⁹ Although it is trivial to apply [POL] to any piece of information, p , like “the beer is in the fridge”, in order to obtain:

[Ex. 1] Query	Where is the beer?" +
Result	“In the fridge” =
<hr/>	
Information	“The beer is in the fridge”

it is important to keep in mind that the correct interpretation of $i^{\text{CLP}} = [Q + R]^{\text{CLP}}$ in [POL] is not as (i) a request for *confirmation* or (ii) a *test*, but as (iii) a genuine request to erase a data deficit through *saturation*. The difference is that, in (i) and (ii), the sender of the query already holds the information that p , but wishes to double-check it, or to check whether the receiver also holds that information; whereas in (iii), the sender lacks the information that p and wishes to acquire it from the receiver by obtaining the missing data. Having said this, let me hasten to clarify a point that might be a source of potential confusion. The polarization of i does *not* really involve two agents. I shall speak sometimes as if the querying sender and the saturating receiver were two different entities, but this is

⁹ In information theory *saturation* is the condition at which a communications system reaches its maximum capacity of traffic-handling.

only for heuristic purposes and ease of treatment. It is i that is being polarised, so sender and receiver are really the same entity. If you need an intuitive representation, imagine a language in which Mary can make statements not by uttering declarative sentences, but only by formulating questions followed by the appropriate answers. Her language does not enable her to say: “The beer is in the fridge” but only “Where is the beer? In the fridge”.

The third advantage is set-theoretic. Adopting a standard extensional theory of questions,¹⁰ it is easy to see that [POL] allows us to treat “is correctly saturated by” as a relation r from a countable set of queries $A = \{Q \mid Q \in A\}$ to a countable set of results $B = \{R \mid R \in B\}$. Note that r is not a function because two or more propositional i , e.g. “the beer is in the fridge” and “the beer is in the kitchen” are analysed as “where is the beer?” + “in the fridge” and “where is the beer?” + “in the kitchen”, thus mapping the same Q_1 both to R_1 and to R_2 (see Figure 1). In section six, we shall see that the real crux is to provide an analysis of correctness that does not beg the question.

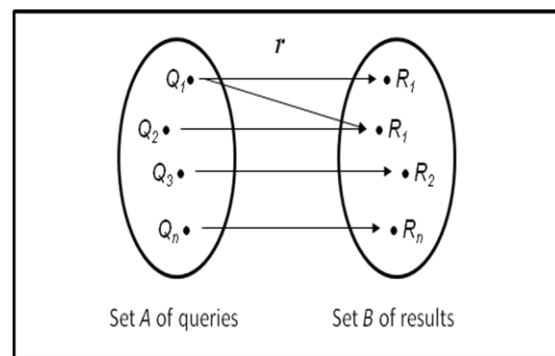


Figure 1: The relation “is correctly saturated by” assigns to each query Q in A at least one result R in B .

4. Third step: Normalization

In real life, queries and results share, in variable proportions, the amount of semantic content that is to be found in the corresponding semantic information. In [Ex. 1], the full semantic content to be found in “the beer is in the fridge” is allocated partly to Q ,

¹⁰ This is a rather standard approach, see Groenendijk and Stokhof (1994) and Szabolcsi (1997).

which contains a request for location and a reference to the object to be located, and partly to R , which contains a reference to the requested location of the object to be located. Although a step forward in the disassembling process, this is still unsatisfactory because it makes it very hard to quantify – precisely, consistently and uniformly across the whole class of $Qs + Rs$ – how much content is allocated to which side of the polarised information. In order to uncover what lies under the thick layer of content, it would be useful to shovel it all on one side, by shifting all the content, still embedded in R , to the left, until R is completely streamlined. At the same time, however, weakening R should not lead to an over-strengthening of Q into a rhetorical question, since a question that requires no answer would be a mere transliteration of i itself and would only defy the purpose. Luckily, a little trick from information theory comes to our rescue: we can reach the right balance, in shifting all the content on the side of the queries, by normalising them into yes/no questions, that is (again the asterisk reminds us that the formula is only a first approximation):

$$[NORM^*] \quad [Q + R]^{CLP \text{ norm}} \Rightarrow [Q_{0/1} + A_{0/1}]^{CLP}$$

The intended interpretation of $[NORM^*]$ is that a query Q and a result R , both CLP-parameterised, can be normalised into a Boolean Question Q and a Boolean Answer A (the $0/1$ subscripts are there to remind us of their Boolean nature), equally CLP-parameterised. This is very much easier done than said, so let us look at our example again. By applying $[NORM^*]$ to [Ex. 1], we obtain:

[Ex. 2] Question “Is the beer in the fridge?” +
 Answer “Yes”=
 Information “The beer is in the fridge”

Of course, this is not what happens in the real world, where one cannot expect a querying sender to be able always to maximise the content of her questions, for she often lacks much more than just a positive or negative saturation. However, recall that we are disassembling semantic information as a given artefact: all the content is already provided, and hence some idealization, typical of controlled experiments, is perfectly

reasonable. Recall also that $[NORM^*]$ does not really involve two agents. This time, imagine Mary being able to state that the beer is in the fridge only by uttering “is the beer in the fridge? Yes”.

Once again, $[NORM^*]$ offers several nice advantages for our analysis, four of which will be immediately useful for our next step.

The first advantage is syntactic: following standard programming languages,¹¹ we can now interpret “+” in $[POL]$ and $[NORM^*]$ more precisely as a *concatenation operator*, whereby a string Q and a string A are locked together to form a longer string i .

The second advantage is semantic: it is now easy to see that it is really Q and not A that sets the scope of the CLP parameters. A Boolean answer can only endorse the *context* (C) in which, the *level of abstraction* (LoA) at which, and the *purpose* (P) for which the Boolean question is formulated; it can neither change nor challenge them. So we can revise $[NORM^*]$ thus:

$$[NORM] \quad [Q + R]^{CLP \text{ norm}} \Rightarrow Q_{0/1}^{CLP} + A_{0/1}$$

The third advantage is set-theoretic: the normalization transforms the relation r “is correctly saturated by” into a function f from a still countable domain of Boolean questions A $\{Q \mid Q \in A\}$ to a codomain of only two possible Boolean answers $\{Yes, No\}$. Figure 2 provides a graphical illustration.




Figure 2: The function f (= is correctly saturated by) assigns to each Boolean question Q in A exactly one Boolean answer (either Yes or No) in B . Note that Q_3 , for example, corresponds to a negative truth, e.g. “the red

¹¹ For example in BASIC, C++, Java, Pascal, and Python.

wine is not in the fridge” in the case in which the fridge does not contain any red wine

Correctness is now a functional concept, but it is still premature to investigate it. At this stage, what matters is that the dramatic downsizing of the codomain of the function represents the extensional counterpart of a fourth, informational advantage: [NORM] shifts all the content in i to Q . We have seen that this re-location of content is what motivates the normalization in the first place. To understand how it works and why it is useful, we need to recall a few other elementary facts in information theory.

As is well-known, given a set of N equiprobable symbols, information theory quantifies the amount of information in a symbol thus:

$$\log_2(N) = \text{bits of information per symbol}$$

It follows that a coin ($N = 2$), by producing a head (h) or tail (t), delivers at most (if it is fair) 1 bit of information, whereas two coins ($N = 4$), deliver at most (again, if they are both fair) 2 bits of information (e.g. $\langle h, t \rangle$), and so forth.

Imagine now a biased coin, which makes obtaining h more likely. The more biased the coin is, the more likely h is, the less information is provided by the answer, the smaller the information deficit becomes, up to the point when, if both sides of the coin are heads, the bias is total, the probability of h is 1, the information conveyed by h is 0 bit and so is the receiver’s information deficit. All this means that, since [NORM] transforms queries into yes/no questions that can be answered by tossing a coin A with different degrees of bias, the worst scenario is one in which Q corresponds to an information deficit that requires at most 1 bit of information from A to be saturated. However, even a $A_{0/1}$ worth a full bit of information fails to add anything, in terms of semantic content, to what is already contained in Q . It follows that, whatever the specific semantic content in i is, [NORM] shifts it entirely to Q , exactly as we wished.

As a consequence, we now have an intuitive way of defining semantic content as unsaturated information or, more formally:

$$\begin{aligned} [\text{CONT}] \text{ Content in } i^{\text{CLP}} &= \text{Content in } Q_{0/1}^{\text{CLP}} \\ &= i^{\text{CLP}} - A_{0/1} \end{aligned}$$

$$= i^{\text{CLP}} - n \text{ bit of information, for } n = 0, 1$$

We have seen the case in which $n = 1$. For $n = 0$, the semantic information i^{CLP} , its content and the content in $Q_{0/1}^{\text{CLP}}$ overlap: this is the case with rhetorical questions (“are you joking?” when used to assert that you are joking), pseudo-questions (“could you close the door please?” asked in terms of a polite request instead of “I would like you to close the door”), self-answering questions (“were the four evangelists more than three?”) and tautological questions (“is $a = a$?” or “are bachelors unmarried?” where the noun and the qualification are both used not mentioned). Still following [CONT], it becomes easy to see how p and $\neg p$ may have exactly the same semantic content while counting as very different information.

[CONT] is not just interesting in itself but provides a reassuring test, since it is perfectly consistent with a theory of strongly semantic information (Floridi (2004c)). In particular, it shows that tautologies and contradictions are pure semantic contents, equally uninformative or, to phrase it differently, that they provide no semantic information about their referents, over and above their contents (in both cases the coin we are tossing has two identical sides, as it were). This is as it should be, so our reverse engineering seems to be proceeding in the right direction.

5. Fourth step: Verification and Validation

We have now disassembled semantic information into two components. By combining [POL] and [NORM], the result can be more succinctly formulated thus:

$$[\text{PN}] \quad i^{\text{CLP}} = Q_{0/1}^{\text{CLP}} + A_{0/1}$$

Let us now scrutinize each component separately.

On the one hand, we have seen that $Q_{0/1}$ sets the CLP parameters. Since it provides all the content in i , $Q_{0/1}$ also identifies its referent, that is, what i is about. We can express all this more precisely by saying that $Q_{0/1}^{\text{CLP}}$ identifies a *system* s (the referent of i) and provides all the *semantic content* (the content in i) for a

model of s (namely, $Q_{0/1}^{CLP} + A_{0/1}$) within a given context, at a particular LoA and for a purpose.

On the other hand, although $Q_{0/1}^{CLP}$ in [PN] is still neither a test nor a request for confirmation but a request for saturation, clearly the sort of saturation in question can no longer be a matter of content, as it was in [POL]. $A_{0/1}$ acts only as a Boolean key, that either fails to apply at all (see $\neg A_{0/1}$ in Figure 3) or that applies and then either locks or unlocks the content provided by $Q_{0/1}^{CLP}$, thus generating a partial model (henceforth just model) of the targeted system. Once again, a conceptual distinction and some technical vocabulary from software engineering (Fox (2007)) can help to clarify this crucial point.

Software Verification and Validation (V&V) is the overall process of checking the “fitness for purpose” of an artefact, by ensuring that the software being developed or modified:

- a) complies with some given *specifications*, regulations or pre-conditions imposed at the start of the development process; and
- b) accomplishes its intended purpose, meeting its *requirements*.

The two phases are complementary.

In phase (a), called *verification* (no relation at all with the philosophical concept), one checks whether one is constructing (or has constructed) what one has (or had) planned to construct, that is, whether the artefact is being developed in the right way. This means evaluating the consistency, completeness and correctness of the software during the stages of its development life cycle.

In phase (b), known as *validation* (again, no relation with the logical concept either), one checks whether one is constructing what is

required, that is, whether the right artefact is being developed. This means evaluating the correctness of the final software with respect to the user’s needs and requirements.

The V&V process applies to a variety of artefacts and products and helps to clarify the twofold role played by $A_{0/1}$ in [PN]. Let me first show how by relying on our example [Ex. 2]. Given the question “is the beer in the fridge?”, any Boolean answer – independently of whether it is “yes” or “no” – implicitly verifies (in the V&V sense) that the question complies with the pre-conditions (i.e., the specifications) regulating its proper formulation, including its context, LoA and purpose. A question like “Is the fridge in the beer?” fails to qualify as something that can receive either a “yes” or a “no” answer because it fails the verification check, since it blatantly fails to develop the semantic artefact in the right way. Once the question is verified – once it is shown to have been formulated properly – the specific answer, either “yes” or “no”, validates (gives a green or a red light to) its content. If this process seems to be prone to error recall that we started by assuming p in order to obtain Q and A , so the possibility of re-obtaining p by re-combining Q and A is a priori guaranteed by hypothesis and sceptical suggestions would merely be out of place here.

All this can be formulated more precisely by saying that $A_{0/1}$ saturates $Q_{0/1}^{CLP}$ by implicitly *verifying* its CLP parameters (roughly: both “yes” and “no” implicitly signal that the question is being asked in the right context, at the right LoA and for the right purpose) and explicitly *validating* its content, as a model of the system (roughly: “yes” and “no” provide a green or a red light for the question respectively). Figure 3 summarises how far we have progressed in reverse engineering semantic information.

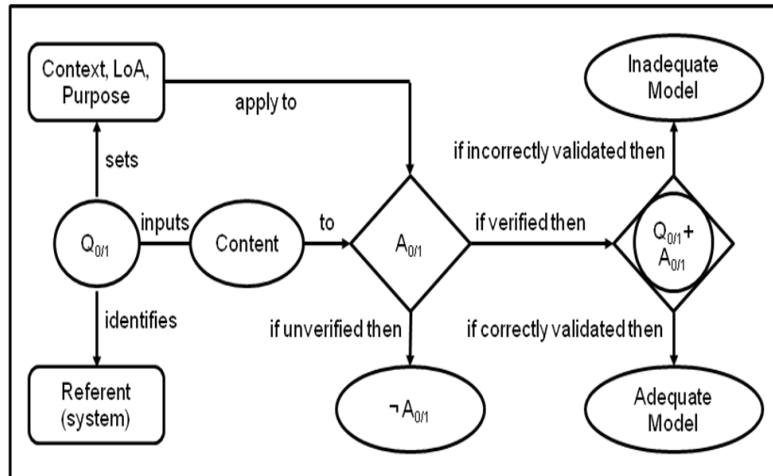


Figure 3: Summary of the first four steps in the Analysis of semantic information. The process starts with $Q_{0/1}$ on the left.

Clearly, a correct saturation consists in a *correct verification* and a *correct validation*. It has taken several clarifications and distinctions and quite a bit of technical vocabulary, but we have finally reached the heart of our problem.

6. Step five: Correctness

Let us quickly review our progress. Simplifying, we now know that p qualifies as semantic information about a system s if and only if p is true; that p is true if and only if A correctly saturates the Boolean question Q corresponding to p ; and that A correctly saturates Q if and only if it correctly verifies and validates it, thus generating an adequate model m of s . Having reduced truth (of semantic information) to adequacy (of the corresponding model m) via correctness (of A with respect to Q), our next challenge is the analysis of the correctness of A .

The challenge consists in negotiating two crossroads in a row. The first is represented by the twofold correctness of the saturation. Let me just highlight the fact that the correct verification of Q by A is a formal precondition for the development of an adequate model m of the targeted system s : it is necessary for, but does not contribute to, the truthfulness of i . In other words, the analysis of the correctness of the verification cannot help us in understanding what it means for semantic information to be truthful. At this crossroad, the really interesting path is represented by the correct validation of Q by A . By following

it, we encounter the second crossroad, represented by two further alternatives. For now we can either analyse correctness of the validation in terms of some concept of truth, thus showing consistency but also failing to provide a non-circular analysis of what it means for semantic information to be true. Or we can move forward, and check whether a further reduction of the correctness of the validation and hence of the adequacy of the issuing model in terms that are truth-poietic but not truth-dependent is possible. Let us quickly review the circular path first.

A useful way to test whether our reverse engineering process is still on the right track is by showing that we have not lost touch with our starting point. Statistics provides the standard analysis of what it means for a model to be adequate (Freedman et al. (2007)). A model is adequate with respect to its target system if it is *valid*. This is now the *statistical* (not the software engineering or the logical) concept of *validity*, which is to be understood as the result of a combination of *accuracy* and *precision*, two other technical concepts borrowed from statistics. Although one might have the impression that we are actually gaining some new ground, it is easy to see that this road only leads back to our starting point. For statistical *accuracy* is the degree of conformity of a measure or calculated parameter (belonging to the model) to its actual, that is, *true*, value (belonging to the system). And statistical *precision* is the degree to which further measurements or

calculations show the same or similar results (this is why it is also called *reproducibility* or *repeatability*). So it turns out that the statistical concepts of validity, accuracy and precision – even assuming that we could adapt them to our less quantitative needs and hence exploit them to clarify what we mean by an adequate model – ultimately presuppose a truth-dependent relation of *conformity* and hence cannot provide a foundational analysis of truth itself without begging the question. The silver lining in all this is that such internal coherence is reassuring: we have not got lost in some conceptual wilderness, while searching for the mechanism that generates semantic information. Encouraged by the knowledge that we could still go back to square one should we wish to do so, let us not press the panic button but push forward.

The second path should lead us away from semantics and epistemology, if we want to avoid ending up back where we started, and take us into the realm of pragmatics, that is, the realm of actual and hopefully successful *interactions* – between an agent a holding the information that p , the model m generated by p , and the system s modelled by m – that can provide some *exogenous* grounding for the evaluation of the quality of the model itself. In order to achieve this, I shall ask the reader to bear with me a bit longer, as I need to introduce two more technical concepts to make sense of such interactions.

One is that of *proxy*, and is borrowed from Information and Communication Technology (Luotonen (1998)). Technically, it refers to a computer agent (e.g., a network service)

authorized to act on behalf of another agent (the client), e.g., by allowing another computer to make indirect network connections to other network services (the server). In this sense, a proxy can be an interface for services that are remote, resource-intensive, or otherwise difficult to use directly. Note that the “proxying” system need not be a copy, an image, a representation or a reproduction of the “proxied” system (the client).

The other concept is that of *commutative diagram*, and is borrowed from category theory (Barr and Wells (1999)). Technically, it refers to a diagram of objects (vertices) and morphisms (arrows) such that, when selecting two vertices, one can follow any directed path through the diagram and obtain the same result by composition.

Adapting these two concepts to our needs, we can now reverse engineer the correctness of the validation, and hence the adequacy of the ensuing model, in terms of the commutativity of the accessibility relation, thus (see Figure 4 for a more intuitive presentation, all Greek letters in [COR] refer to paths in the diagram in Figure 4):

[COR] $A_{0/1}$ correctly validates $Q_{0/1}^{CLP}$ about a target system s identified by $Q_{0/1}^{CLP}$ if and only if $Q_{0/1}^{CLP} + A_{0/1}$ generates (β) an adequate model m of s ; and m is an adequate model of s if and only if m is a proxy (δ) of s such that, if a holds (α) $Q_{0/1}^{CLP} + A_{0/1}$, then a 's proximal access (γ) to m commutes with a 's distal access (ε) to s .




Figure 4: The meaning of [COR]. $Q+A$ is a simplification for $Q_{0/1}^{CLP} + A_{0/1}$.

[COR] offers two advantages and raises a problem. The first advantage is that it finally

introduces an explicit reference to an informee a . This is crucial, since semantic information

is an *objective* (i.e., not subjective) but also *liminal* (that is, neither internal nor external) and *relational* (that is, neither absolute nor relative) concept, like food. It makes little sense to talk about the presence and nature of food without any reference to the specific type of feeder. Likewise, something counts as semantic information only with respect to a specific type of informee.

The second advantage is that [COR] explains the well-known fact that semantic information provides *distal* access to its target. If the agent in the bedroom upstairs asks whether the beer is located in the fridge, and the agent in the kitchen downstairs answers positively, then the agent upstairs, by having proximal access to this overall piece of information, gains distal access to the presence of the beer in the fridge, as long as the answer is *correct*. [COR] merely combines this into a single agent's informative state.

The problem concerns the interpretation of the relation of distal and proximal *accessibility*. If we were to interpret it alethically or epistemically this would obviously fail to take us off the semantic merry-go-round and, sooner rather than later, we would be sent back where we came from. The good news is that we do not need to go down that modal road. On the contrary, the sort of accessibility at stake here is a matter of pragmatic or factual interaction, which provides an exogenous grounding of correctness. It is the one that we find specified in computer science, where accessibility refers to the actual permission to *read* (technically, *sense* and *retrieve*) and/or *write* (again, technically *modify* and *record*) data as a *physical* process. The result is that *a*'s *proximal access* to *m* commutes with *a*'s *distal access* to *s* if and only if *a* can read/write *s* by reading/writing *m*.

The writing of *s* through the writing of *m* is admittedly rare, but it is useful to illustrate it in order to convey the sense of concrete interaction with the targeted system that is involved. Thus, we have left behind a magic culture that considered it an ordinary

phenomenon (cf. the practice of sticking pins in a doll as a method of cursing an individual). Nevertheless, self-fulfilling prophecies (Bill Gates confessing that "Microsoft's shares are overvalued"), performative sentences (the baptising priest declaring that "the name of this girl is Mary"), magic-placebo formulae (the guru concluding that "you are now healed"), authoritative-fictional descriptions ("Sherlock Holmes never visited the Bodleian Library" written by Conan Doyle), God's *intellectual intuition* that *p*, according to Kant, and other ways of "doing things with words" ("this train is not leaving the station" uttered by a dictator) are a good reminder that it is far from impossible to modify/record a system by accessing only its model. Of course, access to *m* is most commonly used in order to read (i.e., sense and retrieve) *s* by reading (ditto) *m*. One gains distal access to (part of) the actual, physical system represented by the fridge in the kitchen and its contents (one senses and retrieves the data in question at a distance) by gaining proximal access to its (partial) model represented by the semantic information "the beer is in the fridge". A way of conveying the same point is by relying on a subjunctive formulation: the proximal read/write access to *m* as a proxy of *s* commutes with the distal read/write access to *s* if and only by having read/write access to *m* one were having read/write access to *s*. This happens in space as well as time: imagine the question being "Will the train leave from platform one?" and the answer being "yes". Semantic information may be seen as a way of being telepresent (Floridi (2005b)).

We needed actual interaction with the system being modelled in order to ground exogenously the correctness of the (validation provided by the) answer to the question pertaining to it, and we have now obtained it. Our toiling is almost over. Putting together this last piece of our jigsaw puzzle, we obtain Figure 5 (the reader may check that this is simply the result of merging Figure 3 and Figure 4, even if this may not be immediately obvious visually):

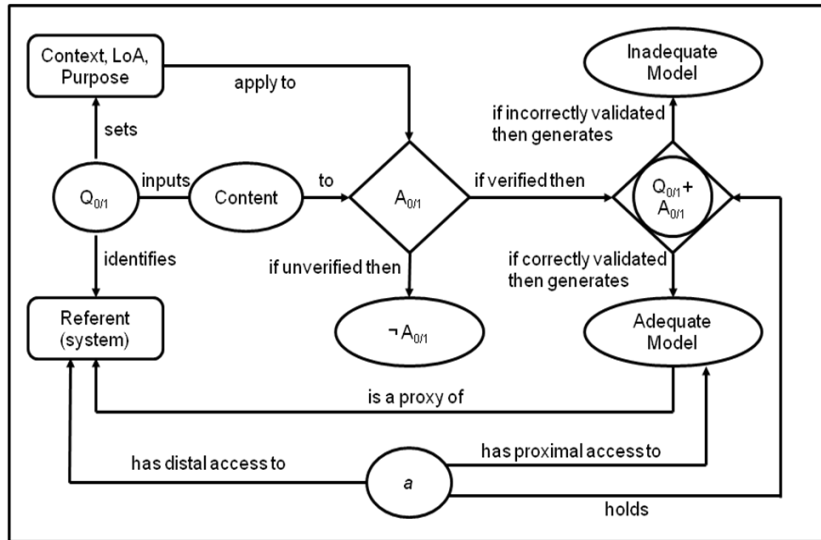


Figure 5: The Correctness Theory of Truth.

Figure 5 represents the blueprint of the mechanism that underlies the truthful nature of semantic information. If we apply it to our example, we obtain:

1. "the beer is in the fridge" qualifies as semantic information if and only if
2. "the beer is in the fridge" is true; this is the case if and only if
3. "yes" is the correct answer to (i.e., correctly saturates by correctly verifying and validating) the question "is the beer in the fridge?"; this is the case if and only if
4. "is the beer in the fridge?" + "yes" generate an adequate model m of the relevant system s ; this is the case if and only if
5. m is a proxy of s and proximal access to m provides distal access to s ; and finally this is the case if and only if
6. reading/writing m enables one to read/write s .

That is, if "the beer is in the fridge" qualifies as semantic information, then holding that semantic information is tantamount to accessing the particular feature of the system addressed by the model which, in our example, is the location of the beer inside the fridge.

7. Conclusion

We have come to the end of a rather long journey. The hope is that the effort might have been rewarding both in itself, if it has been clear enough, and in terms of the final result, if it has been sufficiently convincing. At this point, the reader will probably wish me to keep this conclusion as short as possible. I shall oblige, by adding only a final comment.

Theories of truth often seem to be developed with passive viewers of an outside world in mind, detached observers, whether inside or outside Plato's cave, TV watchers, radio listeners, movie goers, in short, *systems users*, according to the computer science terminology favoured in this article. The correctness theory of truth, proposed in the previous pages, should rather be seen as an attempt to cater for a different sort of customer, namely embodied and embedded, creative agents, who interact with reality, shape and build it, Plato's artisans, writers not just readers, players not audience, in short *systems designers*. To these customers, truth is about constructing and handling artefacts and interacting with them successfully, not merely experiencing them passively. Unfortunately, this is not very Greek, but it is still a very respectable tradition to which both Russell and Tarski belong, insofar as their

groundwork in model theory concerned the design of systems.¹²

¹² A first version of the paper was then presented at the *I International Meeting of Experts in Information Theories* (Leon, Spain, 6-7 November, 2008) and I am very grateful to Francisco Salto for his kind invitation and to the audience for the lively and valuable discussion. I would also like to acknowledge the help, useful comments and criticisms by Patrick Allo, Mark Jago, Sebastian

Sequoiah-Grayson, and Matteo Turilli. I remain deeply indebted to Michael Dummett and Susan Haack for their clarifications and feedback, which date to almost twenty years ago. I am afraid this paper has been a work in progress for quite some time. Joanna Gillies kindly copyedited the last version. All the aforementioned people helped me to improve the paper substantially but they are not responsible for any remaining mistakes.

References

- Austin, J. L. (1950). Truth. In J. O. Urmson and G. J. Wamock (Eds.), *Philosophical Papers*, (pp. 117-133), Oxford: Oxford University Press.
- Bar-Hillel, Y., and Carnap, R. (1953). An Outline of a Theory of Semantic Information. In J. Bar-Hillel, *Language and Information* (pp. 221-274), London: Addison-Wesley
- Barr, M., and Wells, C. (1999). *Category Theory for Computing Science* 3rd Edition, Montreal: CRM Press.
- Barwise, J., and Seligman, J. (1997). *Information Flow: The Logic of Distributed Systems*, Cambridge: Cambridge University Press.
- Bremer, M., and Cohnitz, D. (2004). *Information and Information Flow - an Introduction*, Frankfurt/ Lancaster: Ontos Verlag.
- Chellas, B. F. (1980). *Modal Logic: An Introduction*, Cambridge: Cambridge University Press.
- Colburn, T. R. (2000a). Information, Thought, and Knowledge, In *Proceedings of the World Multiconference on Systemics, Cybernetics and Informatics*, (pp. 467-471).
- Colburn, T. R. (2000b). *Philosophy and Computer Science*, Armonk, N.Y.: M.E. Sharpe.
- Devlin, K. J. (1991). *Logic and Information*, Cambridge: Cambridge University Press.
- Dodig-Crnkovic, G. (2005). System Modeling and Information Semantics, In J. Bubenko, O. Eriksson, H. Fernlund, and M. Lind (Eds.) *Proceedings of the Fifth Promote IT Conference, Borlänge, Sweden*, Studentlitteratur: Lund.
- Dretske, F. I. (1981). *Knowledge and the Flow of Information*, Oxford: Blackwell.
- Dretske, F. I. (1988). *Explaining Behavior: Reasons in a World of Causes*, Cambridge, Mass.: MIT Press.
- Eilam, E. (2005). *Reversing: Secrets of Reverse Engineering*, Indianapolis, Ind.: Wiley.
- Engel, P. (2002), *Truth*, Chesham: Acumen.
- Fetzer, J. H. (2004). Information, Misinformation, and Disinformation, *Minds and Machines*, 14(2), 223-229.
- Floridi, L. (2004a). Information. In L. Floridi (ed.) *The Blackwell Guide to the Philosophy of Computing and Information*, (pp. 40-61), Oxford/New York: Blackwell).
- 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). Is Information Meaningful Data?, *Philosophy and Phenomenological Research*, 70(2), 351-370.
- Floridi, L. (2005b). The Philosophy of Presence: From Epistemic Failure to Successful Observability, *Presence: Teleoperators and Virtual Environments*, 14(6), 656-667.
- Floridi, L. (2006). The Logic of Being Informed, *Logique et Analyse*, 49(196), 433-460.
- Floridi, L. (2007). In Defence of the Veridical Nature of Semantic Information, *The European Journal of Analytic Philosophy*, 3(1), 1-18.
- Floridi, L. (2008a). Data. In: W. A. Darity (ed.) *International Encyclopedia of the Social Sciences*, Detroit: Macmillan.
- Floridi, L. (2008b). The Method of Levels of Abstraction, *Minds and Machines*, 18(3), 303-329.
- Floridi, L. (2008c). Understanding Epistemic Relevance, *Erkenntnis*, 69(1), 69-92.
- Formigari, L. (2004). *A History of Language Philosophies*, Amsterdam/Philadelphia: John Benjamins Pub.
- Fox, C. J. (1983). *Information and Misinformation: An Investigation of the Notions of Information, Misinformation, Informing, and Misinforming*, Westport, Conn: Greenwood Press.
- Fox, C. J. (2007). *Introduction to Software Engineering Design*, Boston, Mass; London: Pearson/Addison Wesley.
- Freedman, D., Pisani, R., and Purves, R. (2007). *Statistics* 4th ed., New York; London: W.W. Norton.
- Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1995). *Design Patterns: Elements of Reusable Object-Oriented Software*, Reading, Mass.; Wokingham: Addison-Wesley.
- Grice, H. P. (1989). *Studies in the Way of Words*, Cambridge, Mass.: Harvard University Press.

- Groenendijk, J., and Stokhof, M. (1994). Questions. In J. Van Benthem and Ter Meulen (eds.), *Handbook of Logic and Language*, North-Holland: Elsevier Science.
- Israel, D., and Perry, J. (1990). What Is Information?. In P.P. Hanson (ed.) *Information, Language, and Cognition*, (pp. 1-28), Vancouver: University of British Columbia Press.
- Künne, W. (2003). *Conceptions of Truth*, Oxford: Clarendon Press.
- Luotonen, A. (1998). *Web Proxy Servers*, Upper Saddle River, N.J.: Prentice Hall PTR.
- Lynch, M. P. (2001). *The Nature of Truth: Classic and Contemporary Perspectives*, Cambridge, Mass.; London: MIT Press.
- Popper, K. R. (1935). *Logik der Forschung : Zur Erkenntnistheorie der Modernen Naturwissenschaft*, Wien: J. Springer.
- Sequoiah-Grayson, S. (2007). The Metaphilosophy of Information, *Minds and Machines*, 17(3), 331-344.
- Sommerville, I. (2007). *Software Engineering* 8th, Harlow: Addison-Wesley.
- Strawson, P. (1964). Identifying Reference and Truth-Value, *Theoria*, 30, 96-118.
- Szabolcsi, A. (1997). *Ways of Scope Taking*, Dordrecht/ London: Kluwer Academic.

About the Author

Luciano Floridi

(Laurea, Rome University "La Sapienza", Mphil and Ph.D. Warwick, M.A. Oxford) is professor of Philosophy at the University of Hertfordshire, where he holds the Research Chair in Philosophy of Information, and Fellow of St Cross College, University of Oxford. He is the founder and director of the IEG, Oxford University Information Ethics research Group, and of the GPI, the University of Hertfordshire research Group in Philosophy of Information. In 2006, he was elected President of IACAP (International Association for Computing And Philosophy). He is the first philosopher to have been elected *Gauss Professor* by the Göttingen Academy of Sciences. In 2009, the American Philosophical Association awarded him the *Barwise Prize*, in recognition of his research on the philosophy of information. Floridi is best known as the founder of two major areas of research: *Information Ethics* and the *Philosophy of Information*. His research interests include the Philosophy of Information, Information and Computer Ethics, Epistemology and Philosophy of Logic and the History and Philosophy of Scepticism. His forthcoming books are the *Handbook of Information and Computer Ethics* for Cambridge University Press, *The Philosophy of Information*, for Oxford University Press, and *Information* for Oxford University Press - A Very Short Introduction series.