Perception and Testimony as Data Providers

Luciano Floridi[1, 2, 3]

[1]Oxford Internet Institute, [2]Faculty of Philosophy, and [3]Department of Computer Science, University of Oxford.

Address for correspondence: Oxford Internet Institute, 1 St Giles, Oxford, OX1 3JS, United Kingdom; luciano.floridi@oii.ox.ac.uk

Abstract

The article addresses two questions. First, if knowledge is accounted information, how are we supposed (to apply this analysis in order) to understand perceptual knowledge and knowledge by testimony? In the first part of the article, I articulate an answer in terms of a re-interpretation of perception and testimony as data providers rather than full-blown cases of knowledge. Second, if perception and testimony are correctly understood as data providers, how are we supposed (to apply this analysis in order) to understand the semantic value of the data provided by such processes? In the second part of the article, I argue in favour of a constructionist hypothesis about how data may become meaningful for human cognitive agents through a process of *repurposing* of natural data/signals. The conclusion of the paper is that human agents are natural-born data hackers.

Keywords

Action-based Semantics; Data; Helmholtzian theory of perception; Natural Meaning; Neo-Kantian epistemology; Non-natural Meaning; Perception; Philosophy of Information; Semantic Information; Symbol Grounding Problem; Testimony.

1. Introduction

What is the relationship between information and knowledge? Recently (Floridi 2010, 2011, 2012) I argued that information—understood as well-formed, meaningful and truthful data—upgrades to knowledge if and only if it is correctly accounted for. The basic idea is rather simple. Each piece of factual, propositional information that *p* (e.g., "the dishwasher's yellow light is flashing"), can be analysed in terms of a Boolean question and answer ("Is the dishwasher's yellow light flashing?" + "Yes"), which, as a stand-alone item, does not yet constitute knowledge, but poses further questions about itself. Such further questions require the right sort of information flow in order to be answered correctly, through an appropriate network of relations with some informational source. This network of information flow is what accounts for the information, and upgrades it to knowledge. To see intuitively why information requires such an accounting network, consider that, if all Alice can do—when asked by Bob why she holds the information that the dishwasher's yellow light is flashing—is to repeat that this is what the yellow light is actually doing, the fact that the dishwasher's yellow light is indeed flashing only warrants at most the conclusion that Alice is informed about the state of the dishwasher's yellow light, but nothing else. For all Bob knows, Alice might have uttered "the dishwasher's yellow light is flashing" as the only English sentence she can master, or she might have dreamed or guessed correctly the state of that particular light. Indeed, the light that Alice reports to Bob to be flashing might have stopped flashing, but then, when Bob goes to the kitchen to check the status of the dishwasher, another dishwasher's light, also yellow but different from the one to which Alice was referring, might have started flashing, making Alice right, yet only accidentally so. This is all very well known. The proposal to resolve such difficulties is to analyse knowledge informationally (Floridi 2011). The result is a definition of knowledge according to which an epistemic agent *S* knows that *p* if and only if:

i) *p* qualifies as semantic information (it is well-formed, meaningful and truthful data);
ii) *q* accounts for *p*, that is, A(*q*, *p*);
iii) *S* is informed that *p*; and
iv) *S* is informed that A(*q*, *p*).[1]

[1] I take condition (iv) to imply that *S* is informed that *q*. The reader who finds this unclear may add a further condition: (iv*) *S* is informed that *q*.

The articulation of the informational analysis of knowledge, in terms of a network theory of account, and its defence, especially against a potential Gettierization (Floridi 2004), are explicit tasks with which I have dealt in (Floridi 2012), so I shall not rehearse the arguments here. Rather, in the following pages I intend to investigate two important consequences of such an analysis:

a) if knowledge is accounted information, how are we supposed (to apply this analysis in order) to understand perceptual knowledge and knowledge by testimony?

In the first part of this article (sections 2-6), I shall articulate an answer to (a) in terms of a re-interpretation of perception and testimony as data providers rather than full-blown cases of knowledge. This, however, leads to a further question:

b) if perception and testimony are data providers, how are we supposed (to apply this analysis in order) to understand the semantic value of such data?

In the second part of this article (sections 7-10), I shall argue in favour of a constructionist hypothesis about how data may become meaningful for cognitive agents like Alice. This will build upon the naturalistic, action-based semantics developed in (Floridi 2011; Taddeo and Floridi 2005, 2007), and provide the missing link between that kind of semantics and the richer, more convention-based semantics enjoyed by Alice and Bob when having an ordinary conversation, basically their uniquely human way of generating and managing non-natural meanings.

In the conclusion, I shall outline some of the consequences of the two answers and contextualise them within the wider context of some previous research.

2. A First Potential Difficulty

The first question has actually been asked explicitly, by phrasing it in terms of a "potential difficulty" by Tommaso Piazza. In a recent, insightful article discussing my proposal for an informational analysis of knowledge, Piazza wrote:

> No less clearly, however, the considerations above [about the nature of the informational account of knowledge] also face Floridi's account with a *potential difficulty*, as they seem to sustain a reasonable doubt about the very viability of this strategy: if one believes that *knowledge can be acquired through perception, or by testimony*, and one also believes that *in those cases there is no accounting or explaining information which could explain the epistemic status to which it is upgraded*, one could well be tempted to suggest that

> *knowledge could not, at least not in general, be analysed as accounted information*; for at least in the cases just envisaged, an explanation of it will have to proceed by taking into account the justificatory role which perception and testimony seem to perform ((Piazza 2010), p. 79, italics added).

I believe Piazza to be mostly right, but perhaps in a way that may not entirely satisfy him, for I shall argue that his premises can be accepted, indeed strengthened, without accepting his conclusion, but first, let me clarify the background against which the discussion is best understood.

3. Some Background

All the empirical information about the world that we enjoy flows, and keeps flowing, to us through our sensorimotor interactions with the world: *directly*, through our perception of the world, possibly mediated by some technologies; and *indirectly*, through our perception of other (possibly even artificial) epistemic agents' perception of the world. We either saw it or read it somewhere, to put it simply, if slightly incorrectly (for we might have heard it, or tasted it, and reading after all is also a case of seeing etc., but I am sure the point is clear). Thus, Aristotelians and Empiricists of various schools are *largely* correct in holding that *nihil est in intellectu quod non prius fuerit in sensu*. "Nothing is in the understanding that was not earlier in the senses", if and only if (the biconditional qualifies the "largely" above) what we are talking about is empirical information about the external world.[2]

If we distinguish the *direct* and the *indirect perception* of the world by referring to the former as sensorimotor perception or simply *perception*, the first-hand testimony of our senses, and to the latter as *testimony*, the second-hand perception by proxy, we see immediately that the potential difficulty, highlighted by Piazza in section two, concerns the only two sources of empirical information available to cognitive agents like us. Without any external perception, either direct or indirect, we could not even be brains in a vat, bio-batteries in a *Matrix*-like world, or dreamers in a Cartesian scenario, because, in each of these cases, we could not be fed any data through our senses. Obviously, it is quite important to check how far the potential difficulty affects the proposal to analyse knowledge as accounted information. Before doing so, however, let me first clear the ground of a potential misunderstanding.

---

[2] Of course, Leibniz's qualification "excipe, nisi ipse intellectus" ("but the intellect itself") remains correct, *New Essays on Human Understanding*, Book II, Ch. 1, § 2.

Knowledge and information states, as well as epistemic, cognitive, and informational processes, are sufficiently similar for our terminology to be interchangeable in most daily circumstances, without any significant loss either in communication or in pragmatic efficacy. This fact reminds us that some tolerance might be sensibly acceptable, even in our technical language. There is an imprecise but still very reasonable sense in which, if Alice sees that such and such is the case, then Alice holds the information that such and such is the case, and *ipso facto* Alice knows that such and such is the case. Thus, if Alice sees a yellow light flashing, then she may rightly claim to know that there is a yellow light flashing in front of her. The same holds true for testimony: if there was a yellow light flashing, and Alice is told by Bob (where Bob could even be a robot or a parrot), who perceived that yellow light flashing, that there was a yellow light flashing, then Alice knows that there was a yellow light flashing. All this I am very happy to concede as uncontroversial in everyday scenarios and parlance.

The value of such a mundane equation—perceiving, or being provided (through testimony) with well-formed, meaningful and truthful data amounting to *p* is equivalent to being informed that *p*, which is equivalent to knowing that *p*—is that, by adopting it, we gain much simplicity. The cost is that we lose the possibility of drawing some conceptual distinctions, which become essential once we wish to be precise in our epistemology and philosophy of information, especially in view of a constructionist approach coherent with Helmholtz's Neo-Kantian philosophy of perception and with the most recent lessons we can learn from neuroscience (more on this later). This is partly why some philosophers, including myself, resist the equation's deflationism. "Partly" because the reluctance is due not only to the cost to be paid (a decrease in our ability to draw finer distinctions), but also to the fact that such cost is philosophically unaffordable once we realise that knowledge that *p* is a specific kind of information that *p*, the kind enriched by the capacity for answering relevant questions about *p*, that about which one is informed, by reference to further information that *q*, which accounts for *p*. Perception and testimony may be analysed along the same lines because—in the best (i.e., non-Gettierised, scepticism-free, error-free (Floridi 2004)) circumstances—they *end up conveying* information about their specific references, but they do not yet represent cases of knowledge: in slightly different ways (to be specified soon) they are our data providers. Let us consider perception first.

4. Perception and the Phaedrus' Test

Epistemologically, our bodies are our cognitive interfaces with the world. Their sensory and kinetic apparatus implements hard-wired levels of abstraction (more technically, we are embodied, cognitive gradients of abstraction, (Floridi 2008)), which determine the range and type of data (observables) that can be negotiated, acquired and processed. Perception is then a general term that refers to the process of data input through which epistemic agents like us acquire first-hand data about their environments, at the levels of abstraction offered by their bodies. In the best scenarios, such data come *from* the world, and this guarantees their facticity. However, they are not *about* the world, and we shall see that this requires some explanation. Suppose Alice sees a yellow light flashing on the panel of her dishwasher at home. Such a process of data input is fallible, but it can be corrected at least through *redundancy* (e.g., Alice sees the yellow light flashing and hears the noise associated with it), *control and/or intervention* (e.g., Alice double-checks that the yellow light is actually flashing by turning on the light in the kitchen and moving closer to the dishwasher), *reasoning* (e.g., Alice infers that it is time for the yellow-light to be flashing), and *social interaction* (e.g., Alice notices that Bob too sees the yellow light flashing). The data input can also be *enhanced* (e.g., through a pair of glasses) and *augmented* (e.g. through a remote monitoring system). By itself, such a first-hand, data-gathering process may be considered a case of knowledge acquisition, but then any elementary signal-processing gadget, like Arduino,[3] would qualify as an epistemic agent, and this seems to be a bullet not worth biting. Let me explain.

In some circumstances, we are not much better off than the aforementioned gadget. Suppose that, when Alice sees a yellow light flashing on the panel of her dishwasher at home, she actually hasn't got a clue about what it might mean. At this stage, all she has acquired, through such perception, is at most the information (equivalent to the propositional content) that a yellow light on the dishwasher's panel is flashing. If Alice has further background information—e.g., about the covariance between the yellow light flashing and the dishwasher running out of salt (see footnote 8 below)—then, by perceiving the light flashing, she may also acquire that further information about the low level of salt. All this is uncontroversial.

[3] "Arduino is an open-source electronics prototyping platform based on flexible, easy-to-use hardware and software. It's intended for artists, designers, hobbyists, and anyone interested in creating interactive objects or environments. Arduino can sense the environment by receiving input from a variety of sensors and can affect its surroundings by controlling lights, motors, and other actuators", from the official website http://www.arduino.cc/

What is notoriously open to debate is whether this yellow-light-flashing-in-front-of-her kind of perception, by itself, may amount to more than just information-gathering at best. I hold that, if we wish to be epistemologically accurate, it does not. It is not enough for Alice to *perceive* a yellow light flashing to *know* that there is a yellow light flashing in front of her, for two sets of reasons. First, because several complex concepts and experiences must already be in place and at play: light; yellow; flashing; the fact that lights can flash; that flashing lights of any colour on the panel of a white good are normally not decorative features (like on a Christmas tree) but signals; that, as signals, lights being off are less perspicuous and hence conventionally less indicative than lights being on; that lights might not work properly but a flashing light on a dishwasher's panel is normally working well and it is meant to be intermittent; etc.. Second, but equally importantly, because the perceptual data input (to simplify: the stream of photons that turns into a there-is-a-yellow-light-flashing-there), plus the conceptual framework (the yellow light flashing there means …) required to formulate and make sense of it, further demand an account (explanation) in order to graduate from information to the higher status of knowledge. In other words, unless Alice understands and is able to answer (at least potentially, or implicitly) a whole series of "how come" questions—how come that the light is flashing? How come that it is the yellow light and not another light that it is flashing? How come that the light is yellow? etc.—her epistemic status is no better than Arduino's or indeed the dishwasher manual's, where one can read that "the yellow light flashing indicates that the dishwasher is running out of salt". Call this the Phaedrus' test:[4] Alice may claim to have more than mere information about the yellow light flashing if she can pass it.

5. Testimony and the Parrot's Test

Let us now turn to testimony. This is the process through which epistemic agents like us *transfer* information to each other. Note that testimony does not *generate* information: the GIGO (garbage in garbage out) rule applies. If Bob tells Alice that *p*—e.g., that the dishwasher's yellow light was flashing yesterday—then, at most, Alice now holds the

---

[4] Plato, *Phaedrus* [275d-e]: [Socrates]: Writing, Phaedrus, has this strange quality, and is very like painting; for the creatures of painting stand like living beings, but if one asks them a question, they preserve a solemn silence. And so it is with written words; you might think they spoke as if they had intelligence, but if you question them, wishing to know about their sayings, they always say only one and the same thing [they are unary devices, in our terminology]. And every word, when [275e] once it is written, is bandied about, alike among those who understand and those who have no interest in it, and it knows not to whom to speak or not to speak; when ill-treated or unjustly reviled it always needs its father to help it; for it has no power to protect or help itself.

information that *p*. Unless we quietly presuppose that Alice, the receiver of *p*, is actually doing *more* than just receiving and registering *p*—e.g., that Alice is not acting as a simple repository but she is also inferring something that Bob did not, namely that the dishwasher is running out of salt, or that she is evaluating the reliability of Bob as the source of *p*, but then all this "more" is where a theory of account is hiding—all we have, at the end of a testimony process, is at most the transfer of some information from the original source to the final target, through a network of senders and receivers. The best that can happen is that the informational baton is passed through the several nodes that are relaying it without being lost or decreased.

Luckily for us, testimony is not a Boolean process. The network is *resilient*—nodes can implement information correction procedures, as when a later epistemic agent recovers or reconstructs what was the original information and relays it in its corrected, restored format—and there is often plenty of *redundancy*—as when several epistemic agents act as independent sources, conveying the same information about the same event, or repeatedly sending the same information at different times and through different channels (Bob tells Alice that the yellow light was flashing, and so does Carol). Still, this is information *transfer*, not *yet* information (let alone knowledge) *generation*. The expression "knowledge by testimony" is really a shortcut not to be taken literally. Receiving *p* can hardly amount to knowing that *p*, for knowledge requires more than true content (or, which is equivalent, well-formed, meaningful and truthful data). If this were not the case, any database would be very knowledgeable indeed and all medieval scribes who copied Greek manuscripts without speaking much Greek at all would have been very learned. In other words, we would like Alice to pass the parrot test (Descartes' *Discourse on the Method,* Part Five): given that the yellow light was indeed flashing, being told, correctly, by a well-trained source like a parrot that the yellow light was flashing while Alice was not in the kitchen does not yet suffice to ensure that Alice *knows* that the yellow light was flashing. At best, she has now acquired that piece of information. If she cannot do anything else with it, then that is all the epistemic dividends she may enjoy.

6. Data Providers

Let us now put the two threads together. Perception is the process through which Alice acquires data about the world, which need to be made *meaningful* and properly *interpreted* (*semanticised*) in order to become information. Perception does not generate propositional semantic information in and of itself. Testimony is the process through which Bob transfers to Alice propositional information (also but not only) about the world, but does not yet generate propositional knowledge in and of itself. In both cases, what is missing, in order to gain empirical knowledge of the world in a precise epistemological sense, is the understanding (explanation in a different vocabulary, or account, as I would prefer) of the empirical information acquired. Such understanding (explanation) is obtained through the intelligent accounting of the available propositional semantic information. This is what I have argued in (Floridi 2011). Time to return to the potential difficulty.

It should now be clear that Piazza is right is stating that (first premise):

P.1) "knowledge can be acquired through perception, or by testimony",

as long as "acquired" in (P.1) is understood, as it should, as stating necessary but not yet sufficient conditions. Compare this to "*x* (a mortgage, a passport, a skill, etc.) can be acquired through *y* (a credit evaluation, a full application, the relevant training, etc.)". Indeed, in this sense, I have argued for a stronger thesis: empirical knowledge can be acquired *only* through perception or by testimony. If one day we are able to implant Wiki-microchips under our skin, it will still be a case of testimony.

Piazza is also right in stating that (second premise):

P.2) "in those cases [perception and testimony] there is no accounting or explaining information which could explain the epistemic status to which it is upgraded",

if we understand by (P.2) that unaccounted perception or testimony do not qualify *yet* as knowledge.

Where he seems to be mistaken is in drawing the following conclusion from the previous two premises:

C) "one could well be tempted to suggest that knowledge could not, at least not in general, be analysed as accounted information".

Nobody who understands the previous analysis and the two premises should be tempted to jump to such conclusion. He adds that

> "for at least in the cases just envisaged, an explanation of it will have to proceed by taking into account the justificatory role which perception and testimony seem to perform".

So perhaps the problem lies with the devilish concept of justification. Now, without entering into a lengthy discussion of the nature of justification and its role in epistemic processes (Floridi 1996), there are at least two ways in which perceiving that such and such is the case—e.g. seeing that the yellow light is flashing—justifies Alice in holding that such and such is the case. One is by interpreting the justification in terms of causal interactions. Reliabilist theories used to like this approach. It seems impossible to disagree with this interpretation: it is the perceptual (visual, in the example) process of data-input that causally makes possible the acquisition of the relevant bits of information about the yellow light flashing. Yet causality is not all that is being invoked here, since we are not looking for a mere descriptive account, but for a normative one. So the alternative is to use justification to mean exculpation. This, however, adds nothing to our, or Alice's, understanding of the case in question, even if it does add a note on her epistemic conduct in such circumstances. She did not dream it, nor imagine it, she did not project it out of fear, nor carelessly assumed it: she saw a yellow light flashing, eyes wide-open, double-checking, changing angle and perspective, perhaps asking Bob as well. Alice really did her best to make sure that what she actually saw was indeed a yellow light flashing. She did the right thing. The verdict is: causally sound and epistemologically not guilty. Yet all this is irrelevant to Alice's epistemic state. We still cannot tell whether she *knows* or is merely *informed* that *p*. As I have argued above and much more extensively and in detail in (Floridi 2011), being correct about *p* and having done everything reasonably possible to avoid being mistaken about *p* does not yet mean that one knows that *p*. An important part of the epistemic story is still missing. In order to see why, let us first turn to the second question.

7. A Second Potential Difficulty

Suppose the previous analysis is correct, or at least moves in the right direction. Perception and testimony are both to be understood as data providers. In the case of testimony, the suggestion seems less controversial, as long as one understands that testimony is a data providing process not in the sense that it transfers raw data about the world (think of the photons in the case of the yellow light flashing), but in the sense that it transmits well-formed and meaningful data from sender to receiver. Ultimately, perception deals with the world, testimony with information about the world: it is the difference between cooking with fresh ingredients (perception) and microwaving a pre-cooked meal (testimony). Testimony that conveys empirical information about the world ultimately depends on perception of the world based on data from the world. It is informative when the well-formed and meaningful data it transmits are also truthful, otherwise it is misinformative (unintentionally false) or disinformative (intentionally false). I shall return to such a crucial role at the end of section ten.

Testimony presupposes the occurrence of data already meaningful. This cannot be said of perception, and this raises a second, potential difficulty. For once perception is stripped off its high epistemological status—once perception no longer counts, philosophically, as a full-blown, genuine instance of knowledge, but rather as a necessary condition of possibility of empirical information and hence of knowledge—one may object that we have swung to the other extreme. For now it becomes difficult to explain how perception, so epistemologically impoverished, may progress to generate empirical knowledge at all. Recall the example of Arduino: artificial agents are very proficient at collecting, storing and manipulating data, and yet they do not go on to produce empirical knowledge, not in the sense in which Alice does. If the previous analysis reduces Alice's epistemic state to Arduino's, we have a new potential difficulty.

This second difficulty can be phrased in terms of a dilemma: either perception is *overinterpreted* informationally, but then this fails to explain how it differs from full-blown empirical knowledge (what is the difference between perceiving that such and such is the case and knowing that such and such is the case?) and why it does not require the ability to (explain, justify or) account for the information it provides; or perception is *underinterpreted* informationally, as the necessary source of the data that go on to constitute empirical knowledge, but then this fails to explain how such data can become full-blown empirical knowledge. We move from an inflated to a deflated view of perception, when what we need is just the right epistemological evaluation in between. As I argued above, working on the first horn of the dilemma looks unpromising. The alternative is to show that the data-based interpretation of perception is not stuck in the impasse of an underinterpretation. This is the task of the next three sections, for which we need more background.

8. More Background

Semantic information is a very slippery topic. If we know the relevant codes, we patently have no difficulty in understanding sentences, maps, formulae, road signs, or other similar instances of well-formed and meaningful data. And yet, scientists and philosophers have struggled to determine what exactly semantic information is and what it means for an agent to elaborate and understand it. One of the sources of the difficulty is known as the "symbol grounding problem" (SGP):

> How can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads? How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols?" (Harnad 1990), p. 335.

The difficulty in solving the SGP consists in specifying how agents can *begin* to elaborate *autonomously* their own semantics for the data (symbols, signals) that they manipulate, by interacting with their environments and other agents, without begging the question, that is, without relying on more or less hidden forms of *innatism* or *externalism*: semantic resources should be neither presupposed, as already "pre-installed" in the agents in question, nor merely "uploaded" from the outside by some other agents already semantically-proficient. If they are, we are really addressing a different kind of question.

In chapters six and seven of (Floridi 2011), I argued that all the main strategies proposed so far in order to solve the SGP fail to satisfy the previous conditions (clustered under the expression *zero semantic commitment condition* or Z condition), but they provide several important lessons to be followed by any new alternative. In light of such critical analysis, I elaborated a constructive proposal for a *praxical* solution to the SGP. There is neither space nor need to outline it here. Suffice it to say that the praxical solution is based on two main components: a theory of meaning—called *Action-based Semantics* (AbS)—and an architecture of agents—which models them as being constituted by at least two modules, M1 and M2. M1 operates at an *object level* (OL), interacting directly with the external environment, e.g., by navigating, detecting obstacles, avoiding them etc., thus outputting and inputting actions. M2 operates at a *metalevel* (ML), taking as data the actions at the OL and as target of its elaborations the internal states of M1. Any action that M1 outputs to, or inputs from, the environment defines a particular internal state (S_n) of M1. Hence actions and internal states are causally coupled: for any different action in M1 there is a different internal state S_n, and for all similar actions in M1 there is the same S_n. Simplifying, M2 uses such states to provide a semantics for the symbols developed by the system. Thanks to their architecture, agents can implement AbS, and this allows them to ground their symbols semantically as well as to develop some fairly advanced semantic abilities, including forms of semantically grounded communication and of elaboration of semantic information about the environment, and hence to overcome the SGP without violating the Z condition. The reader interested in the details (and viability) of the proposal is invited to read the two chapters. Here, it is important to stress that such a praxical solution points towards a more ambitious and challenging possibility: the elaboration of a theory of meaning that can enable us not to underinterpret perception as a data provider but rather account for its role in the elaboration of empirical knowledge. The hypothesis is that the praxical solution of the SGP provides the seeds for an explanation of how advanced semantic and linguistic skills may develop among higher biological agents in more complex environments when perception and later testimony are in question. This is what we shall see in the next two sections.

9. The Vice Analogy

In trying to show how the analysis of perception as data provider may lead to the elaboration of a meaningful experience of the world and hence to empirical knowledge of it, two converging strategies may be adopted. They may be quickly introduced as the two jaws of a vice. On the one hand, one may rely on *quantitative* analyses, especially, but not only, through information theory and Shannon information, the algorithmic theory of information and Kolmogorov information, and signalling theory. On the other hand, one may rely on *qualitative* analysis, especially, but not only, through truth-theoretic semantics, inferential role semantics, game-theoretic semantics, and meaning as use.

The limits of such strategies are well-known. Quantitative analyses are not meant to deal with semantics, while qualitative analyses offer at most semantic *criteria* (how can one tell whether Alice understands the meaning of *p*), or presuppose meaningful contents (how Alice successfully handles meanings by becoming proficient in a particular perceptual or linguistic game), but are not meant to explain how semantics (including meaningful data) arises in the first place.

When misapplied, both kinds of analyses are a way of cheating. Quantitative analyses do less than they are said (but not meant) to do. When misused, this leads to a *semantics-from-syntax fallacy* so flagrant in the failures of classic AI. Qualitative analyses presuppose (correctly) what they are said (mistakenly) to deliver. When misused, this leads to a *semantics-from-semantics fallacy* and the failures of current solutions to the SGP (see Z condition above).

The question is whether there is a way to bridge the physical/syntactic side, addressed by quantitative analyses, and the mental/semantic side, addressed by qualitative ones, in order to explain how perception, and later on testimony, as a data proving processes, may lead to the generation of meaning. The answer is that pragmatics might help. Here is a quick list of some lessons we have learnt from the two kinds of strategies recalled above:

a) there is an active component dealing with meanings, this is the semantic engine (*agent*);
b) *interactions* between the environment (*system*) and the agent elicit the data used by the agent as constraining affordances to create semantic information (a *model*) of the system;
c) *semanticisation* (the generation and attribution of meaning to data/signals) is a functional relation between meaningless input and meaningful output;
d) evolutionarily, models of the system compete with each other on the basis of two quality requirements: *fit for purpose* and *correct* (right) *first time*;
e) the agent's interactions with the system provide the competitive context within which incorrect models are revised or abandoned (Bayesian learning);
f) ultimately, semantic information is the outcome of the agent's active and constructive interpretation of the system that is the referent/source of the relevant data, not of its passive representation.

Let me now show how the elements may be put together with a praxical approach (the pragmatic "bridge" just mentioned) to give rise to a full picture.

10. The Constructionist Interpretation of Perception and Testimony

Imagine a very early stage where there is no difference between agent and system, or sender and receiver, informer and informee. We may assume the presence of only an environment, in which physical structures occur more or less dynamically, that is, there are patterns of physical differences understood as asymmetries or lack of uniformities. There is no specific name for such "data in the wild". One may refer to them as *dedomena*, that is, "data" in Greek (note that our word "data" comes from the Latin translation of a work by Euclid entitled *Dedomena*). Dedomena are not to be confused with *environmental information*. They are pure data, that is, data before they are interpreted or subject to cognitive processing. They are not experienced directly, but their presence is empirically inferred from, and required by, experience, since they are what has to be there in the world for our information about the world to be possible at all. So dedomena are whatever lack of uniformity in the world is the source of (what looks to an informational agent like Alice) data. Try to imagine the photons that will generate the perception of a yellow light flashing before they are perceived as a yellow light flashing. Such data might be flowing around, but they are not signals yet, as there are no senders or receivers.

Such an initial stage is where there are environmental data and patterns that might be exploitable as information by the right sort of agents for their purposes, before there is any kind of communication. Therefore, it is also the stage (see Figure 1) where Shannon's classic model of communication may easily be misleading, if applied too early. In the relata (sender, receiver) vs. channel and communication process (message), it is the message that comes logically first, in the form of physical data as potentially exploitable constraining affordances.

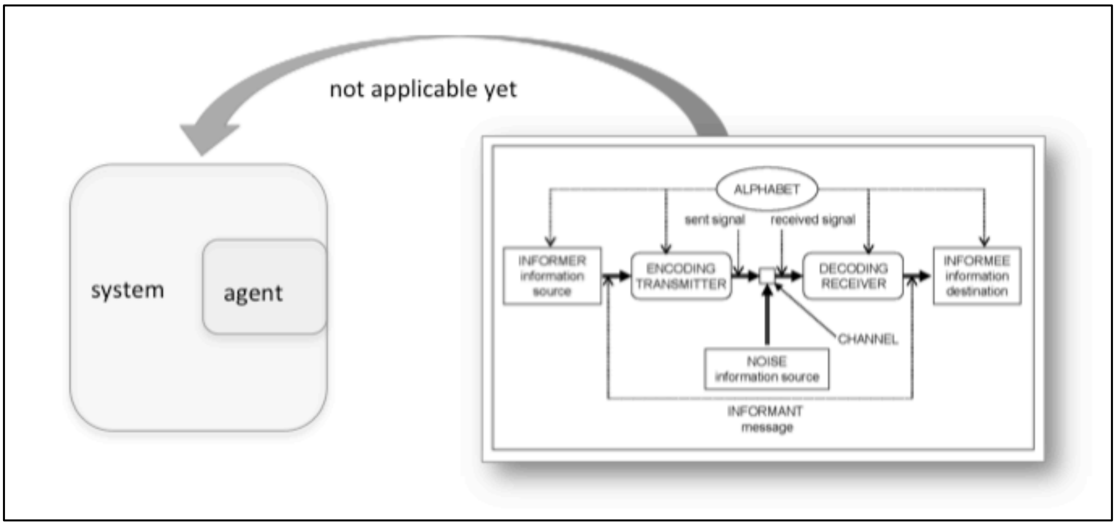

Figure 1 Shannon's Model of Communication not Applicable

Once some structures in the environment become encapsulated through a *corporeal membrane*, such encapsulation of part of the environment allows the separation of the interior of an agent from the external world. The ontological function of the membrane is to work as a hardwired divide between the inside, the individual structure or agent, and the outside, now the environment *for* the agent. Its negentropic function is to enable the agent to interact with the environment to its own advantage and withstand, evolutionarily, the second law of thermodynamics, for as long and as well as possible. The epistemological function of the membrane is that of being selectively-permeable, thus enabling the agent to have some minimal variety of degrees of inputs and outputs with respect to the environment. At this stage, data are *transduceable* physical patterns, that is, physical signals now seen as *broadcast* by other structures or agents in the environment, which are captured by the permeable membrane of the agent. The body is a barrier that protects the stability of the agent (physical homeostasis). A good example is a sunflower.

We move from pre-cognitive to post-cognitive agents once data become *encodable* resources, exploitable by agents through some language broadly conceived (sounds, visual patterns, gestures, smells, behaviours, etc.). Patterns and flows of data/differences, which were before quantities without direction (scalars), broadcast by (what it is still improper to interpret as) sources not targeting any particular receiver (e.g. the sun generating heat and light, or the earth generating a magnetic field), acquire a direction, from sender to receiver (become vectors), and an interpretation (e.g., noises become sounds interpreted as alarms), thus being exploitable as signals. From now on, latent Shannon information becomes manifest and Shannon's classic communication model applies. This shift requires a *cognitive membrane*, or *bodily interface*, which allows the encapsulation of data (some form of memory) for processing and communication. The body as an interface or cognitive membrane is a semi-hardwired (because configurable through learning) divide between the cognitive agent and its environment, that is, a barrier that further detaches the agent from its surroundings, and allows it to exploit data processing and communication in its struggle against entropy.

At this stage (see Figure 2), sensorimotor interactions through bodily interfaces are best understood as interactions at a given set of levels of abstractions or LoAs (gradient of abstraction, (Floridi 2008)), where LoAs are hardwired as sensory receptors. Note that, according to this reconstruction, there are no signals (let alone information) in the wild: data as signals are elicited by the nomic interactions between types of systems and types of agents. This is not relativism but *relationism* (if the difference is unclear, consider the concept of food: not everything is food, but food is understandable only relationally, by understanding the nature of both the consumed substance and the consuming agent). Agents are further de-coupled from their environments, with different embodiments determining different types of epistemic agents, which are able to interact informationally with their environments and other agents through their bodily interfaces. Thus, each type of agent is a type of LoAs implementation. Same type same LoAs. Wittgenstein's lion and Nagel's bat are incommensurable LoAs. The stability (cognitive homeostasis) now concerns the internal data within the agent and their codification: memory and language.

The emergence of natural signals as meaningful for an agent is the stage where the praxical solution to the SGP is applicable. To oversimplify, the semantic value (meaning) of the signals is the state in which they put the receiving agent (cf. adverbial theory of perception and Grice's comments below). A good example is a bird on the sunflower.

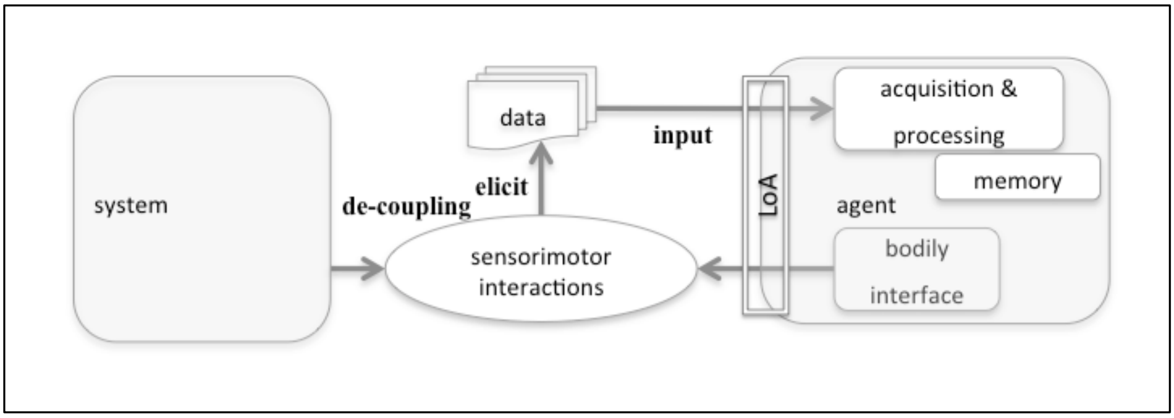

Figure 2 The Interactions Agent-System elicit Data as Signals at a Bodily LoAs

The elicited data, in Figure 2, understood now as signals, may have both a source *s* (a sender) and a referent *r* that is in some state φ (this is what the data are taken to be "about"), see Figure 3. Clearly, if there is no referent *r*, then we have a *virtual system*: the model (the interpreted data) generates its referent, as in a computer game. If there is no source *s*, then the agent is in a state of complete ignorance, where "complete" means that the state of ignorance itself is unknown. If both *r* and *s* are present, then, in most cases of communication, including perception and testimony, $s \neq r$. This simply means that the data come *from* a source interacting with the agent (the photons coming from the yellow light), but they are not *about* the source (the photons are not "about" the yellow light), which is not their referent, not least because they are the outcome of the cognitive interactions and negotiations between agent and the data source, although we shall see that there is a plausible sense in which ordinary perception works correctly when it interprets $s = r$. If $s = r$, testimony becomes an unusual case of self-confession. Francis Bacon was perhaps the first to rely on this feature in order to speak metaphorically of the scientific inquiry as a questioning procedure addressed to Nature, an informational interpretation of the logic of discovery that we consider rather common nowadays (Sintonen 1990). However, even conceding that all this, including the praxical solution of the SGP, is correct, it still falls short of providing a full account of Alice's perception of the yellow light flashing as indicating that the dishwasher is running out of salt. For the latter is a conventional meaning, and Grice was right in distinguishing it from natural meaning (Grice 1957), as I shall explain below. At this stage, the best one can do, without begging the question, is to show how Alice may be "put in a yellow state", as it were, by a yellow light flashing. According to the praxical solution of the SGP, there is a plausible sense in which Alice may be said to see "yellowly", but she cannot be said to see "salt-in-the-dishwasher-runnig-outly", not without presupposing what needs to be accounted for.

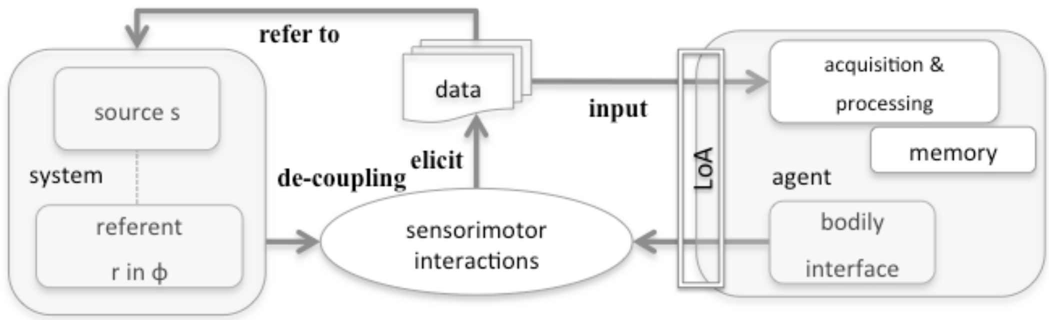

Figure 3 Data Source and Referent

Pace Skyrms, the naturalist tradition—which seeks to account for non-natural meanings by reducing them entirely to natural ones through signalling or information theory (Skyrms 2010)—provides the right beginning but seems to be unable to deliver the whole story. Using a different example, it struggles to explain why the same sound is perceived as (mind, not merely believed to be) a song by some and as a national anthem by others. Non-natural (conventional, artificial, synthetic) semantics seems to require more than natural semantics to emerge. If this were not the case, we would have already made at least some successful steps in the realization of classic AI, where the frame problem is just a specific instance of the SGP (Harnad 1993). We have not (Floridi, Taddeo, and Turilli 2009). Indeed, the whole project of information or signal processing as sufficient for the development of a full-blown semantics runs into the semantics-from-syntax fallacy, seen above. The usual reply to such an objection consists in asking for more, indefinitely: more time, more complexity, more processing, more "add your resource". In AI, this has often and conveniently translated into more funding. Yet, in the same way as we are reminded in AI that climbing to the top of a tree is not the first step towards the moon, but the end of the journey, no matter how many more resources may become available, so, in naturalistic theories of meaning and of meaningful perception, accounting for the communication procedure among birds, bees, monkeys, or indeed robots, is not the first but the last chapter in the book of natural semantics. It is where things start becoming interesting if difficult, not where one may accept a "…and so on" clause or some hand-waving.

This is hardly news. More than half a century ago, Grice had already identified and exposed the shortcomings of such a naturalism:

> I want first to consider briefly, and reject, what I might term a causal type of answer to the question, "What is non-natural meaning [Grice actually uses the abbreviation meaning$_{NN}$] ?" We might try to say, for instance, more or less with C. L. Stevenson,[5] that for *x* to non-naturally mean [mean$_{NN}$] something, *x* must have (roughly) a tendency to produce in an audience some attitude (cognitive or otherwise) [in the praxical solution this is expressed in terms of putting an agent in the specific, correlated, internal state] and a tendency, in the case of a speaker, to be produced by that attitude, these tendencies being dependent on 'an elaborate process of conditioning attending the use of the sign in communication.'[6] This clearly will not do. (Grice 1957), p. 379.

Grice goes on to explain several reasons why such a naturalisation will "clearly not do". They were as clear and hardly refutable then as they are now, but if the reader remains unconvinced, let me add a further consideration. The irreducibility of non-natural meanings to natural ones is not just a matter of scientific results and philosophical arguments. The view that just more data or signals processing, without appeal to any further variable, may somehow lead to the development of higher-level, non-natural semantics—what Grice describes as, at best, a circular reasoning ("We might just as well say, 'X has non-natural meaning [meaning$_{NN}$], if it is used in communication,' which, though true, is not helpful"), see the fallacy above of obtaining semantics from semantics—also runs against a specific result in information theory, one that indicates that natural data input and processing is necessary but

[5] Grice adds here a footnote to (Stevenson 1944), ch. iii.

[6] Grice adds here a footnote to (Stevenson 1944), p. 57.

insufficient to generate the meaning for perceptions and concepts such as "the dish washer is running out of salt". This is the data processing theorem (DPT).

The DPT concerns the quantity of mutual information between signals, messages or data. *Mutual information*[7] is, together with the *inverse relation principle*,[8] the *covariance model*,[9] and Shannon's *communication model*, one of the pillars of any information-based project for the full naturalisation of semantics. The DPT states that data processing tends to *decrease* information. Here is an informal summary. Suppose three systems S_1, S_2, and, S_3 are such that X is the output of S_1 and the input of S_2, Y is the output of S_2 and the input of S_3, and Z is the output of S_3, as illustrated in Figure 4, then:

DPT) if the random variables depend on each other, that is, $[X \Rightarrow Y \Rightarrow Z]$; and

if $[X \Rightarrow Y \Rightarrow Z]$ is a Markov chain;

then the mutual information I satisfies the following condition: $I(X; Y) \geq I(X; Z)$.

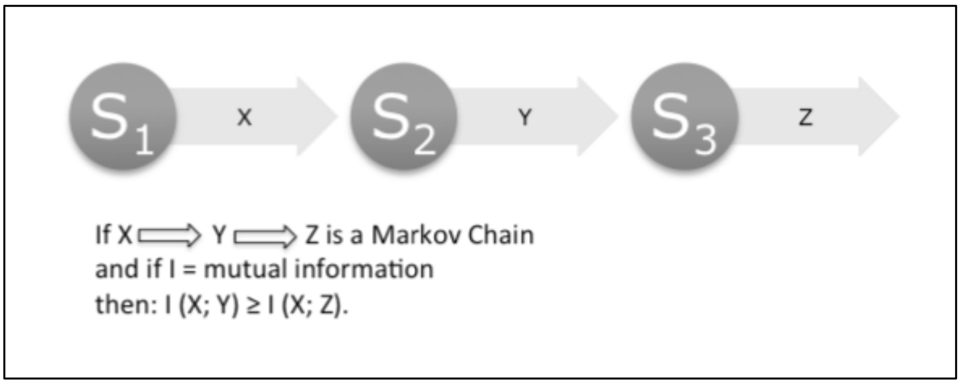

Figure 4 The Data Processing Theorem

[7] Mutual information, indicated as $I(X; Y)$, is a measure of how dependent two random variables X and Y are, for example, the dependency between the information X = the dish washer is running out of salt (that is, the average reduction in uncertainty or the expected reduction in yes/no questions needed to guess X) and the information Y = the low salt yellow light indicator is flashing. The higher the dependence is the higher the degree of mutual information is. Mutual information satisfies the properties $I(X; Y) = I(Y; X)$; $I(X; Y) \geq 0$; if X and Y are independent, then $I(X; Y) = 0$; highest I when X = Y (ideal, noiseless channel).

[8] The principle states that there is an inverse relation between the probability of *p*—where *p* my be a proposition, a sentence of a given language, a situation, or a possible world—and the amount of semantic information carried by *p*. Thus, a biased coin provides increasingly less information the more likely one of its outcomes is. The principle, though very plausible, runs into two problems, the "scandal of deduction" (D'Agostino and Floridi 2009; Hintikka 1973) and the "Bar-Hillel-Carnap Paradox"(Floridi 2005).

[9] The model states that if two systems *a* and *b* are coupled in such a way that *a*'s being (of type, or in state) *F* is correlated to *b* being (of type, or in state) *G*, then such correlation carries for the observer of *a* the information that *b* is *G*. For example, the dishwasher's yellow light (*a*) flashing (*F*) is triggered by, and hence is informative about, the dishwasher (*b*) running out of salt (*G*) for an observer *O*, like Alice, informed about the correlation. See (Barwise and Seligman 1997), (Dretske 1999), (Floridi 2010).

This means that the average amount of information that Z conveys about X is usually less than, and at most equal to, the average amount of information that Y conveys about X. Of course, larger degrees of mutual information correspond to greater degrees of statistical dependence between a system's input and output, in our example between X and Z. Indeed, we shall see below that such mutual information can reach total equivalence between variables. But the introduction of further n stages of data processing can never increase the statistical dependence between the variables, and is likely to decrease it. In short, if one does not have such and such information at the input, data processing, as formulated above, is not going to generate it. If one obtains it at the end of it, either it was already there since the beginning (see above Grice's comment about the true but unhelpful solution, or the fallacy of semantics-from-semantics), or it has been surreptitiously introduced by something other than the data processing itself. For example, the Markov chain has been broken[10] (for a classic example of such a "break" in human communication consider the "crying wolf" scenario).[11]

Intuitively, the reader may recall the game of Chinese Whispers, in which a first player in a line whispers a message to the player next in line, who whispers it as accurately as possible to the next player, and so on, until the last player communicates the message to the whole group. The players are the equivalent of S_1, S_2, S_3, and so forth, and their whispered messages are X, Y, Z, and so forth. The longer the chain of speakers is, the less likely it becomes that the final message will resemble the initial one. With a slogan more memorable but less accurate: data processing tends to destroy information, it certainly cannot increase it. Asking for more signal processing is not going to solve the problem of escalating natural semantics to non-natural semantics. At best, it can only deliver more natural semantics. Some other factor must be at play.

Let us take stock. We saw that, if perception is a data providing process, then we need to be able to explain how such data become meaningful and hence suitable for generating information (well-formed, meaningful and truthful data) and then knowledge (accounted

[10] This has been suggested as a solution to the problem of enriching the semantic value of computer visualizations, in (Chen and Floridi forthcoming). For a simple and balanced introduction to the limits of Markov Chains in animal communication see (Bregman and Gentner 2010), pp. 370-371.

[11] This is Aesop's famous tale: a young shepherd repeatedly raises a false alarm (by crying wolf) in order to trick nearby villagers into thinking that a wolf is attacking his flock. When a wolf actually does attack it, his signal goes unnoticed because the villagers, who *remember* all the times this was a false alarm, no longer believe him, and the flock is destroyed.

information). There are ways, such as the praxical solution to the SGP, to show how the meaning of some perceptual data may be naturalised. However, many, if not most, of our perceptions, deal with non-natural meanings. There was a time when Alice saw a piece of cloth with some coloured patterns on it, but it is now impossible for her not to see the Union Jack, or Old Glory, or the Tricolore, etc. In this case, signalling, the agent's data processing and internal states, and the memory of such states, are necessary but insufficient conditions to account for the emergence of non-natural semantics. We also saw that the development of cognitive agents should be interpreted in terms of an increasing distance from their environment. Despite this—and despite the fact that data as signals are elicited by the interactions between the agent and the system, and hence should be taken as negotiations with, rather than representations of, the system—we usually correctly assume that, in the best (but also very common) circumstances, even when non-natural meanings are in question, the end-product of the agent's perception of the system is a faithful grasp of the state of the system by that agent. When Alice sees the yellow light flashing and perceives that the dishwasher is running out of salt, that is normally the case: she is correct and the dishwasher is indeed running out of salt. The same holds true for her perception of the national anthem, the flag of her country, the red traffic light at the crossroads, and so forth. Mistakes are the exception in Alice's cognitive life. She would not be here as a species, if they were the rule. Or, to put it in more Davidsonian or Dennettian terms, the more mistakes she made the less she would be interpretable as a cognitive agent.[12] In order to solve this strange predicament, according to which perception as data provider both decouples the agent from her environment through more and more non-natural meanings and couples her to it successfully and indeed in a way that is cognitively superior to any other species, we need to adopt a different perspective and move from a naturalistic to a constructionist view (Floridi 2011). Here is how we may do it.

It seems incontrovertible that human agents do not merely use natural meanings but constantly *repurpose* them for other epistemic, communicative and semanticising goals. Alice not only sees the yellow light flashing, she also repurposes it to mean, in the kitchen appliance context, that the dishwasher is running out of salt. Conventional or non-natural meanings are the outcome of such repurposing. The cognitive strategy of using, converting or modifying

[12] I owe this point to one of the anonymous referees, who actually remarked: "Of course the more Dennettian or Davidsonian view is that 'she wouldn't be [interpretable as] cognitive at all, if mistakes were the rule'. I agree, but I think it would be a matter of degrees rather than threshold.

data/signals for a purpose or function other than their original, natural one, to fit a new use— think of Peirce's distinction between icon, index and symbol—is very cost-effective and can be reiterated endlessly: a cloth becomes a flag, which becomes a country, which becomes a foe to burn, or something to be proud of and wear as qualifying one's identity, and so forth. By repurposing perceptual data, human agents like Alice actually use them as resources to interact with the world, with themselves (see narrative theories of the self, (Floridi 2011)) and among themselves more richly, innovatively, inventively, indeed more intelligently, than any other kind of agents we know, which are unable to go beyond natural semantics. And since "repurpose" may be just another word for "hacking", a simple and more colourful way of putting the previous point is by saying that *humans are natural-born data hackers* (see Figure 5).

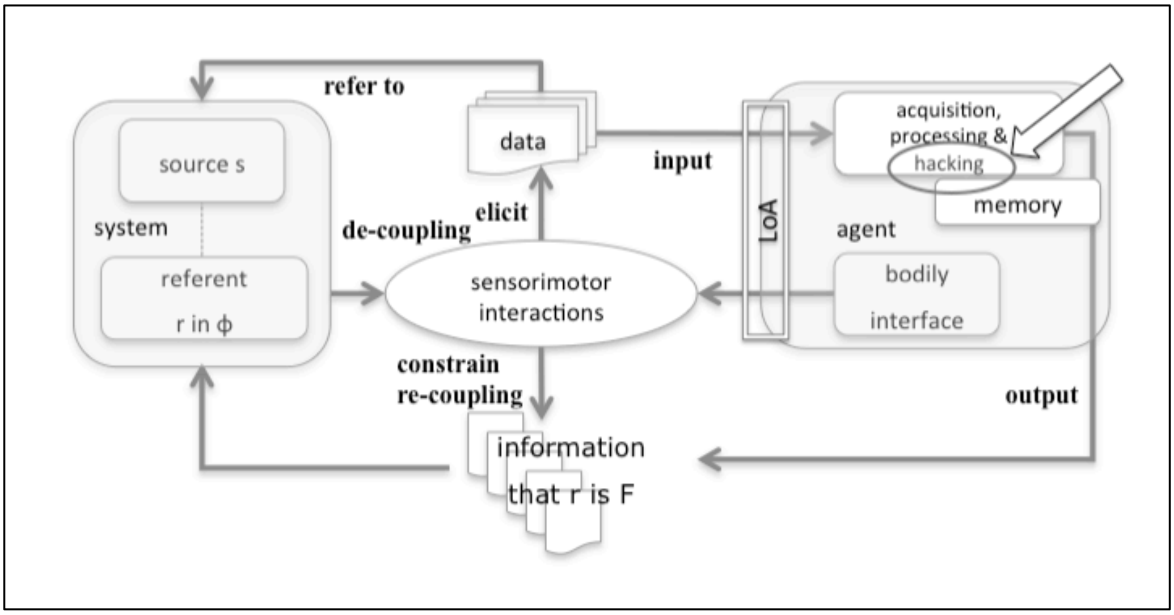

Figure 5 Non-natural Meanings as Data Hacking

A conception of Alice as a cognitive mirror or a representationalist mechanism is simply wrong. It is certainly inconsistent with our best neuroscience:

> Increasingly, the brain reveals itself proactive in its interface with external reality. In the past, our conception of the brain changed from that of a mirror to that of an interpreter. Several current lines of research—in fields such as memory, motivation and attentional orienting—now begin to cast the brain as a predictor. The results of experience are integrated over various timescales in order to anticipate events relevant

> to the current task goals and motivational state of the individual and to tune the relevant perceptual and motor machinery accordingly. In particular, research on attentional orienting has shown how signals coding predictions about the location, identity or simple features of relevant events can influence several stages of neural processing. Recent evidence shows that these predictions are not restricted simply to the contents of events but also extend to their anticipated timing. (Nobre, Correa, and Coull 2007), p. 465.

The philosophical hypothesis about repurposing—the view that humans are natural-born data hackers—may be convincing, but there is a final problem. For unless Alice's repurposing of natural data and signals is somehow successfully constrained, its outcome could be indistinguishable from the delusions or hallucinations of a mentally ill person. The similarity between semantic repurposing and mental disorder is an important point, to which I shall return in the conclusion. At the moment, it is clear that, while the hypothesis of a data hacking process may solve the problem of understanding how the data acquired through perception may move from natural to non-natural meanings, it does not, in itself, say anything about the fact that through perception human agents interact very successfully with each other and their environments, and are re-coupled to the world in the most realistic, indubitable, "what you see is what you get" kind of way. What does re-couple Alice to the world, once she is decoupled from it by her data hacking? Recall that, normally, outside Hume's studio, Alice has no doubts whatsoever about the fact that the yellow light flashing is exactly what the world is like, and that the world is such that it now contains her dishwasher running out of salt. She is right, and her successful interactions with the world show her to be right, but this is a problem because, at the moment, the creativity offered by data hacking fails to explain the cognitive and pragmatic success of her naïve and commonsensical realism. Indeed it works counter to it. What seems to be required is a re-interpretation of the representational correspondence between (i) Alice's perception of the world and (ii) the way the world is, as (iii) *retro-fitness*, in the following sense. In normal and common circumstances, say in Alice's kitchen, data/signals (the yellow light flashing) sent by, and referring to, system r (the dishwasher) in state φ (running out of salt) become the information (model) that r is F (the dishwasher is running out of salt) by being processed in such a way that the mutual information $I\ (\varphi(r);\ F(r)) = \text{MAX}$. This ensures complete correspondence between the perception and the perceived, which Alice enjoys cognitively and on which her actions are successfully based practically. However, such a complete correspondence seems possible only if

a) either: *s* (the source) = *r* (referent), *F*(*r*) is a faithful representation of φ(*r*), and there is a noiseless ideal channel between the system-referent and the agent;
b) or: *F*(*r*) constitutes φ(*r*).

None of the three conditions in (a) seems really satisfiable, so (b) remains the only option, but this means that perceptual information is the *output* of the perceptual data processing/hacking not the input (see Figure 6).

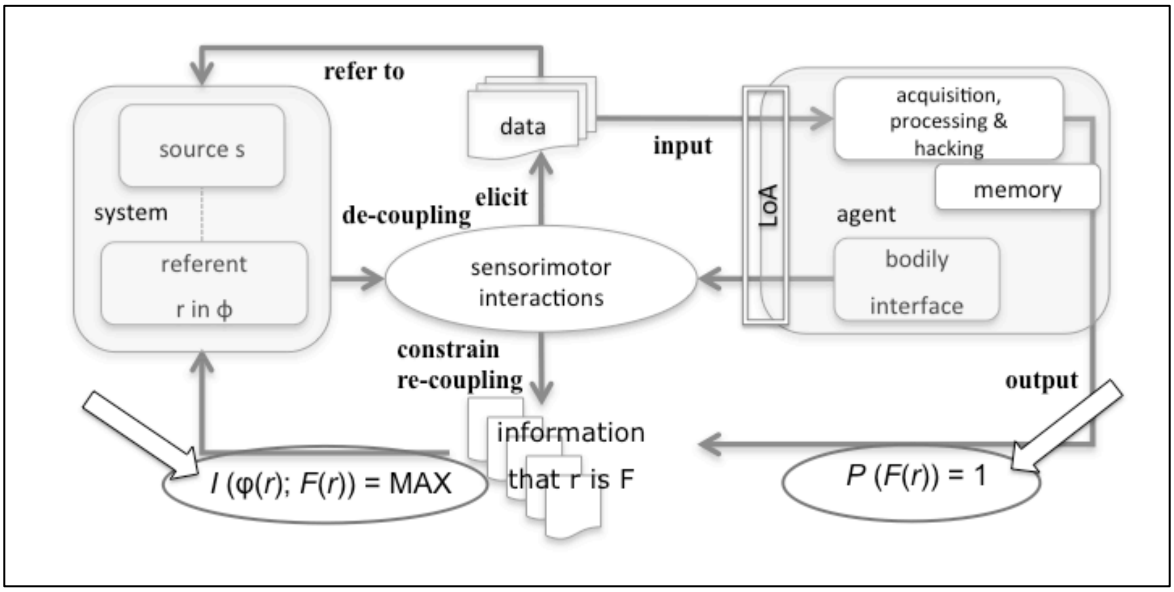

Figure 6 The Construction of Perceptual Information

Such output—the models of the world that the data (constraining affordances) hacking generates—competes for success on the basis of *fit for purpose* and *correct* (right) *first time*. Given her constraining affordances (data), Alice generates an endless amount of interpretations of the world: some of them are correct because they respect and make the most of her data and are evolutionarily selected to ground and improve her interactions with her environment and other agents ("the dishwasher needs more salt"), some are simply mistaken ("the dishwasher is not getting enough water"), many are simply innocuous and unrestrained by the available data ("the yellow light flashing means I will not get the job"). Following the inverse relation principle, Alice will be in no state of surprisal (Shannon's term) with regard to her perceptions: she cannot be informed *about* them (mind, not *by* them) because to her the probability *P* of her perception (*F*(*r*)) is 1. New information, of the kind exemplified by the yellow light flashing, is the exception, certainly not the norm. To put it in Kantian terms, perceptual information about the world is the world, and the world-information by default has probability 1 for those who perceive it. The system is the source/referent of the data, but the interpreted data, properly understood as semanticised constraining affordances, do not represent the system, no more than radio signals represent the radio sending it. To a realist this may sound a bit suspicious, so let me haste to add that we are standing on the shoulder of a gigantic champion of realism, Helmholtz. Here is how he put it:

> Our sensations are effects brought forth in our organs by means of exterior causes, and how such an effect manifests itself depends of course quite essentially on the nature of the apparatus on which the cause operates. Insofar as the quality of our sensations gives us *information* about the peculiarities of the exterior process that excites it, it can count as a *sign* of that process, *but not as a picture*. For one expects of a picture some sort of similarity with the pictured object […]. But a sign need have no similarity of any sort whatever with that of which it is the sign. The relation between them is only that the same object, working its effects in the same way, produces the same sign, and that unequal signs always correspond to unequal causes. To the popular view, which naively and complacently assumes the full truth of the pictures that our senses give us of things, this remainder of similarity that we recognise may seem rather paltry. In truth it is not; with its aid something of the greatest significance can be achieved: the representation of the regularities in the processes of the real world […]. So even if our sense impressions in their qualities are only signs, whose special nature depends wholly on our internal organisation, they are nonetheless not to be dismissed as empty appearance, but are in fact a sign of something, whether this is something existing or something occurring; and what is most important, they can picture the law of this occurring. (Helmholtz 1995), pp. 347-8, cited by (Carus 2007), pp. 117-8)

Helmholtz published this in 1878. Had he written it after Shannon, he would have spoken of data providers.

In all this, testimony as information transmission, not yet generation, and as a by-product of perception, which allows further semantic hacking, plays a final and crucial role.

So far, the analysis has been developed by considering only a single agent. Of course, this is an untenable simplification. Alice is part of a community of speakers and epistemic agents. Most of the semantics she enjoys and controls is inherited. She read on the manual that the yellow light flashing nonnaturally means (or, to put it *à la* Grice means$_{NN}$) that the dishwasher is running out of salt. It is here that testimony—understood as the main mechanism through which agents learn and share a language as well as information and hence can constitute a multiagent system (a community of speakers)—plays a fundamental role (see Figure 7). For testimony is what enables the development of language as the main cognitive tool to hack natural meanings, thus allowing the Lamarckian evolution of hacked data through generations (cumulative learning).

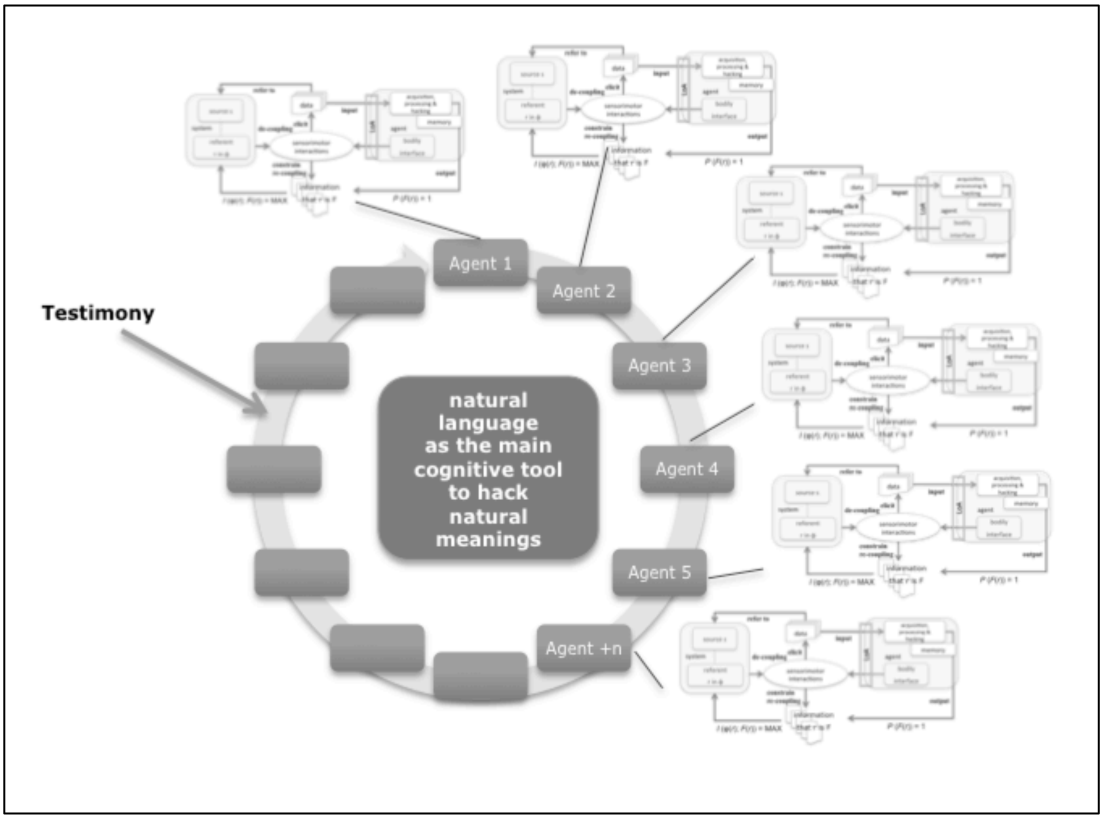

Figure 7 Perception, Testimony, and Language

11. Conclusion: the Beautiful Mistake

Much more could and should be said about the previous two answers to the questions outlined in section one. The data hacking hypothesis is only the beginning. In order to facilitate further steps ahead, in this conclusion I shall only summarise a few salient points.

We saw that perception is a complex process through which constraining affordances (data) are negotiated, acquired, elaborated, and repurposed by epistemic and semantic agents like Alice and Bob in order to make sense of, and interact successfully with, their environments both naturally and nonnaturally. The interpretation of perception as a decoupling and then re-coupling process of the cognitive agent is coherent with the development of language, through testimony, as mainly a cognitive tool (Deacon 1997; Schilhab, Stjernfelt, and Deacon 2012) rather than just a communication medium, and with the emergence of *consciousness* (Floridi 2005) and the construction of a sense of *personal identity* (Floridi 2011) as part of a progressive detachment of the agent from the world. Mentally healthy humans are different from animals because they are nonnaturally de-coupled from the here-and-now by their data hacking, and they are different from the mentally ill because the same data hacking re-couples them to the here-and-now inventively and successfully. Contrary to animals we construct semantic artefacts. Contrary to the mentally ill, our semantic artefacts work *correctly*.[13] From an evolutionary perspective, we are uniquely different from, and more successful than, other species not because of a plus but because of a minus; namely the (both perceptual and linguistic) semantic incapacity of being absolutely and inseparably present, cognitively, where we are located, bodily. We cannot help experiencing the world *as*. This gap, this detachment or decoupling, this initial incapacity of being thoroughly absorbed by the world—which our intelligence and mental life then has to bridge through the development of language and our knowledge of the world—is what makes us special. Looking for our semantic, linguistic, mental, conscious *quid* is looking for an absence, for a gap, for a fissure. Indeed, it has been controversially argued in psychiatry[14] that the same

---

[13] For a theory of truth as correctness see (Floridi 2010).

[14] The theory that schizophrenia might be a consequence of the human evolution of language is scientifically associated to the research of Tim Crow, a professor of psychiatry at Oxford University. A close view, according to which schizophrenia contributed to the evolution of *homo sapiens*, was popularised rather controversially, by David Horrobin (Horrobin 2001). More recently, the publication of (Faulks 2005), a novel in which the theory is presented in a fictional scenario, "sparked an academic feud" (Thorpe 2005).

evolutionary causes lie behind our capacities to develop both language and mental illness. It seems that the price to be paid to be *Home sapiens sapiens* is that of being potentially the schizophrenic species as well. The fissure is double-edged, as it were.

Some 50,000 years ago, the *Homo* species finally snapped and began regularly, widespreadly, and consistently to distance itself from its environment through the development of a culture of tool- and weapon-making, art (sculpture, cave painting, body ornaments), travelling, long-distance trade, and burial rituals (Diamond 2006). We are not evolution's finest moment, the peak of the process, some kind of *Über*-animal, but nature's beautiful mistake. We are the odd ball in the natural set. It is because we are a bit less that we are so much more. And in the same way as a broken mechanism that manages to survive and evolve by repurposing itself is perfectly natural, and yet unlikely to be reproducible, it remains an open question whether we might ever be able to "break" our syntactic machines in such a way as to make them intelligent like us. Real, old-fashioned AI as we know it may remain unachievable not because it is physically impossible—Nature managed the trick with us—but because some mistakes may be perfectly natural and yet be forever unique. The beautiful mistake may well be Nature's *hapax legomenon.*

Philosophically, the accidental and yet resilient balance between informational decoupling and re-coupling, detachment and engagement, is better understood from a constructionist and non-representationalist perspective, than from a naturalistic and representationalist one. This means understanding our informational, linguistic, and cognitive activities as part of our creative manipulation of the data/signals (the cognitive constraining affordances) negotiated with the world. Epistemologically, this leads to the acknowledgment that knowledge by acquaintance is more fundamental, in terms of input, than know-how, and this, in turn, is more fundamental than knowledge-that, to rely on a useful but slightly inadequate Russellian terminology. The opposite is true when it comes to creativity and repurposing. It also means that we should be able to find a middle ground between naïve realism (which is not a philosophical position but really the final outcome of a complex process of successful construction) and relativism, reinforcing the relational analysis of many of our key informational concepts. Only by strengthening our understanding of such a third way, Plato's *metaxy* (see for example *Symposium* 203b-c), shall we be able to escape the usual dichotomies that haunt our philosophy (think of the naïve question whether colours are in the perceived or in the perceiver).

The world as we consciously experience it is the totality of its models as we unconsciously create them. It is the outcome of a constant construction and amendment of a stable interpretation. We are cognitive amphibians: as embodied and embedded physical agents, we live in, and interact with the world in itself, Kant's *noumena*. We eat and drink, handle and build, avoid and crash into *noumena*. But as informational organisms, we experience and inhabit the world as a semantic reality both afforded and constrained by the world in itself. Our ontology is entirely semantic, so we know the world when we are informed about it and are able to account for such information. For a knower is "the man who knows how to ask and answer questions" (Plato, *Cratylus*, 390c), giving an account, that is, about the information that she holds.

Acknowledgements

I discussed previous versions of this article as invited speaker at the "Séminaire, 'Philosophie de l'informatique, de la logique et de leurs interfaces'" (École Normale Supérieure, Paris, 3-4 April, 2012), the "Colloquium, Epistemology of Information and Communication (University of Lyon, Lyon, 8 April 2011), the "Third International Workshop on the Philosophy of Information" (Royal Flemish Academy of Belgium for Science and Arts, Paleis der Academiën, Brussels, 18-19 November 2010), and the "Research Seminar, ICTs in the Contemporary World" (London School of Economics, London, 2 November, 2010). I am indebted to Jean-Baptiste Joinet and the École Normale Supérieure; Fidelia Ibekwe-SanJuan and the University of Lyon; Patrick Allo, Giuseppe Primiero and the Royal Flemish Academy of Belgium for Science and Arts; Jannis Kallinikos and the London School of Economics; and to the participants in such meetings for the fruitful opportunity to discuss my ideas and to receive so much helpful feedback. Penny Driscoll kindly copyedited the final version. The two anonymous reviewers of the journal provided many insightful comments and suggestions for improvements, I owe them badly, I wish I could buy them a drink.

References

Barwise, Jon, and Jerry Seligman. 1997. *Information flow: the logic of distributed systems*. Cambridge: Cambridge University Press.

Bregman, M. R., and T. Q. Gentner. 2010. "Syntactically Complex Vocal Systems." In *Encyclopedia of Animal Behavior*, eds. M. D. Breed and J. Moore. Oxford: Elsevier. vol. 3, pp. 368-374.

Carus, André W. 2007. *Carnap and twentieth-century thought: explication as enlightenment*. Cambridge: Cambridge University Press.

Chen, Min, and Luciano Floridi. forthcoming. An Information Map for Visualisation. *Synthese*.

D'Agostino, Marcello , and Luciano Floridi. 2009. The Enduring Scandal of Deduction. Is Propositional Logic really Uninformative? *Synthese* 167 (2): 271-315.

Deacon, Terrence William. 1997. *The symbolic species: the co-evolution of language and the brain*. New York: W.W. Norton.

Diamond, Jared M. 2006. *The third chimpanzee: the evolution and future of the human animal*. 3rd (original 1991) ed. New York: Harper Perennial.

Dretske, Fred I. 1999. *Knowledge and the flow of information*. Stanford, Ca.: CSLI Publications; originally Oxford: Blackwell, 1981.

Faulks, Sebastian. 2005. *Human traces*. London: Hutchinson.

Floridi, Luciano. 1996. *Scepticism and the foundation of epistemology: a study in the metalogical fallacies*. Leiden; New York: Brill.

———. 2004. On the Logical Unsolvability of the Gettier Problem. *Synthese* 142 (1): 61-79.

———. 2005. Consciousness, Agents and the Knowledge Game. *Minds and Machines* 15 (3-4): 415-444.

———. 2005. Is Semantic Information Meaningful Data? *Philosophy and Phenomenological Research* 70 (2): 351-370.

———. 2008. The method of levels of abstraction. *Minds and Machines* 18 (3): 303-329.

———. 2010. *Information - A Very Short Introduction*. Oxford: Oxford University Press.

———. 2010. Semantic Information and The Correctness Theory of Truth. *Erkenntnis* 74 (2): 147-175.

———. 2011. A Defence of Constructionism: Philosophy as Conceptual Engineering. *Metaphilosophy* 42 (3): 282-304.

———. 2011. The Informational Nature of Personal Identity. *Minds and Machines* 21 (4): 549-566.

———. 2011. *The Philosophy of Information*. Oxford: Oxford University Press.

———. 2012. Semantic Information and The Network Theory of Account. *Synthese* 184 (3): 431-454.

Floridi, Luciano, Mariarosaria Taddeo, and Matteo Turilli. 2009. Turing's Imitation Game: Still a Challenge for Any Machine and Some Judges. *Minds and Machines* 19 (1): 145-150.

Grice, H. P. 1957. Meaning. *The Philosophical Review* 66 (3): 377-388.

Harnad, Stevan. 1990. The Symbol Grounding Problem *Physica*: 335-246.

———. 1993. Problems, Problems: The Frame Problem as a Symptom of the Symbol Grounding Problem. *Psycoloquy* 4 (34): on-line/unpaginated.

Helmholtz, Hermann von. 1995. *Science and culture: popular and philosophical essays*. Chicago: University of Chicago Press.

Hintikka, Jaakko. 1973. *Logic, language-games and information: Kantian themes in the philosophy of logic*. Oxford: Clarendon Press.

Horrobin, David F. 2001. *The madness of Adam and Eve: how schizophrenia shaped humanity*. London: Bantam.

Nobre, Anna Christina, Angel Correa, and Jennifer T. Coull. 2007. The Hazards of Time. *Current Opinion in Neurobiology* 17: 465-470.

Piazza, Tommaso. 2010. Perceptual Evidence and Information. *Knowledge, Technology & Policy* 23 (1): 75-95.

Schilhab, Theresa, Frederik Stjernfelt, and Terrence William Deacon. 2012. *The symbolic species evolved*. New York: Springer.

Sintonen, Matti. 1990. How to Put Questions to Nature. *Royal Institute of Philosophy Supplements* 27: 267-284.

Skyrms, Brian. 2010. *Signals evolution, learning, and information*. Oxford: Oxford University Press.

Stevenson, Charles L. 1944. *Ethics and language*. New Haven: Yale University Press.

Taddeo, Mariarosaria, and Luciano Floridi. 2005. Solving the symbol grounding problem: a critical review of fifteen years of research. *Journal of Experimental & Theoretical Artificial Intelligence* 17 (4): 419-445.

———. 2007. A Praxical Solution of the Symbol Grounding Problem. *Minds and Machines* 7 (4): 369-389.

Thorpe, Vanessa. 2005. How a novelist's twist sparked academic feud. *The Observer* (Sunday 21 August).