

Lunch Uncertain

Review of: Floridi, Luciano (2011) *The Philosophy of Information* (Oxford)

Stevan Harnad

Chaire de recherche du Canada en sciences cognitives

Institut des sciences cognitives (ISC)

Université du Québec à Montréal

Montréal, Québec, Canada H3C 3P8

<http://www.crsc.uqam.ca>

&

School of Electronics and Computer Science

University of Southampton

Highfield, Southampton

SO17 1BJ UNITED KINGDOM

<http://users.ecs.soton.ac.uk/harnad/>

ABSTRACT: The usual way to try to ground knowing according to contemporary theory of knowledge is: We know something if (1) it's true, (2) we believe it, and (3) we believe it for the "right" reasons. Floridi proposes a better way. His grounding is based partly on probability theory, and partly on a question/answer network of verbal and behavioural interactions evolving in time. This is rather like modeling the data-exchange between a data-seeker who needs to know which button to press on a food-dispenser and a data-knower who already knows the correct number. The success criterion, hence the grounding, is whether the seeker's probability of lunch is indeed increasing (hence uncertainty is decreasing) as a result of the interaction. Floridi also suggests that his philosophy of information casts some light on the problem of consciousness. I'm not so sure.

The sceptical philosophers worried about uncertainty. They wanted to know what (if anything) we can know for sure: not only what is true, but what we can be *certain* is true. Their standard for certainty was logic and mathematics: things that can be proved to be true "on pain of contradiction" are necessarily true, so we can be certain they're true. But the only things that are necessarily true on pain of contradiction are formal truths, such as "If all humans are mortal and if Socrates is a human, then Socrates is mortal." You could replace that with "If all A are B and if x is A then x is B", and still be sure that's true for any A, B, and x whatever. But that doesn't tell you whether there are such things as humans, Socrates or mortality, let alone whether Socrates was mortal.

So what about "Socrates existed"? True, but not necessarily true; so we can't know for sure. Maybe the historical evidence is wrong. What if you saw Socrates with your own eyes? But how can you be sure of what you see? Maybe it was someone else, or maybe you were hallucinating.

There's the same uncertainty about scientific truths: Can we know for sure that apples fall, as the "law" of universal gravitation tells us? No. Maybe the law of universal gravitation is false: We can't prove it's true. Even if all apples have been falling unfailingly every time anyone ever looked, how can you know for sure they'll keep falling tomorrow? How can you even know for sure that there are apples, or anyone in the world, or the world, or a tomorrow?

That's scepticism. And if you insist that knowing means knowing for sure, then (with one prominent exception to which I will return at the end of this review) the only things we can know are those that are necessarily true, on pain of contradiction. Everything else that's true is only probably true, according to the available evidence so far; it's not certain. If something is indeed true, and you believe it's true, then you are right, especially if you believe it for the right reasons. But you don't really "know" it, because you can never be sure. Without a proof it's only probably true. —Not to worry, though, because probability on the basis of the available evidence to date, though short of certainty, looks as if it's good enough for government work (and science, and Darwinian survival).

What is evidence? It's data; and data are just physical events, with a certain "shape". Suppose I am on a long distance phone call. I say "Socrates is mortal." That acoustic event causes other events in the telephone receiver, which are then propagated along the cable, till they reach the receiver and ear of the person (or device) on the other end of the line. Claude Shannon's classical theory of information is about how reliably the data on one end of the line reach the other end. The theory says nothing about the meaning (if any) of the data, just about how well the data get through. The theory does give a definition of *information* that goes a little beyond data-transmission, however: information is whatever reduces uncertainty among alternative outcome possibilities. Certainty would be zero uncertainty, but only logical necessity can provide that. So everything else (with the one exception that I promised I would mention at the end: be patient) is only about reducing uncertainty, not eliminating it completely.

An example: suppose your daily lunch depends on being able to press the right button on a food-dispenser. There are six buttons, only one of them delivers food, you don't know which one, it changes randomly each day, and you get only one chance a day to press. Your uncertainty is 5/6, hence your chance of a daily meal is 1/6. If I tell you the food-button is an odd number today (and it's true), then I've reduced today's uncertainty to 1/2. If I say it's an odd number less than 5, I've reduced it to 1/3, and so on. If your probability of getting lunch has indeed been increased by getting those data, then you've been informed. What if I tell you the food button is number 3? Have I reduced your uncertainty to 0? Not quite, because there's still the sceptics' caution: I may be mistaken or lying, or there may be no food (or world) at all. But we've agreed to set aside that sort of fussiness: probability on the available evidence to date is close enough.

Luciano Floridi's book proposes a philosophy of information. Shannon's information theory is only about data-transmission and uncertainty reduction: but what about the *meaning* of the data? How do we know what data mean? How does that meaning get into

our brains? Here Floridi invokes the “symbol grounding problem,” according to which symbols (such as the ones in mathematics, computer programs, or the English language) are just meaningless data-strings, squiggles and squoggles. In data-processing they are rulefully manipulated on the basis of their shapes (or “syntax”) rather than their meanings (“semantics”) – except if the symbols are causally connected, somehow, to the things they are about. Think of a robot with the string of symbols “‘humans’ ‘are’ ‘mortal’” occurring inside its head. Those symbols are grounded if the robot, using its sensors, can recognize, respectively, humans and things that are mortal, and can do with them (and say of them) all the kinds of things that we humans can do and say. (Such a robot would pass the “Turing Test” if it could do all that -- not only with the symbols “humans” and “mortals,” but with everything a normal person can name, say and do.) Unless they are thus grounded, the robot’s symbols are as meaningless as the words on a page: they mean nothing to either the page or the robot, only to the reader (and it’s the symbols in the reader’s head that need to be grounded). But is grounding enough for meaning?

Floridi does not think that robotic grounding is enough if it is accomplished by neural networks that learn what’s what through trial and error, because such networks depend on external feedback for error-correction in much the same way that books depend on the external reader for meaning (though it’s not clear why the networks depend on external feedback any more or less than we do, whenever we learn things). The usual way to try to ground knowing according to contemporary theory of knowledge is: We know something if (1) it’s true, (2) we believe it, and (3) we believe it for the “right” reasons. Floridi proposes a better way. His grounding is based partly on probability theory, and partly on a question/answer network of verbal and behavioural interactions evolving in time. This is rather like modeling the data-exchange between the data-seeker who needs to know which button to press on the food-dispenser and the data-knower who already knows the correct number. (The data-knower, in turn, would have come to know the correct number through a previous question/answer network of data-exchanges along these same lines. And even before that, the question-answerer could have been Nature, correcting our errors with their negative consequences.) The success criterion, hence the grounding, is whether the seeker’s probability of lunch is indeed increasing (hence uncertainty is decreasing) as a result of the interaction. This is Floridi’s update on the standard view that to know something is to believe something true to be true for the “right” reasons.

Floridi also suggests that his philosophy of information casts some light on the problem of consciousness. I’m not so sure. Yet consciousness may indeed have something very fundamental to do with meaning and knowing, thanks to that one further certainty I promised: The other kind of truth we can know for sure, apart from truths that are formally necessary, is that it *feels like something* to know (or mean, or believe, or perceive, or do, or choose) something. Without feeling, we would just be grounded Turing robots, merely acting as if we believed, meant, knew, perceived, did or chose. Even for robots acting and interacting as a result of the right, Floridian, uncertainty-reducing reasons and causes, if they did not feel, all that would really be happening would be meaningless data-processing, modulating equally meaningless outcome probabilities – but with more lunches eventually being eaten. This would be grounded syntax; but still not semantics. If, however, the Floridian robots did indeed feel, then their

symbols could have not only grounding but meaning. That certainly does not mean that apart from the sure fact that they felt, everything else the feeling robots felt they knew for sure, they really knew for sure. Most of the truths that the feeling, meaning robots believed would (like our own) remain matters of probability, rather than certainty, hence not known for sure -- merely believed for the right reasons on basis of the data available to date, namely, whatever reduced or minimized uncertainty among the alternative outcome possibilities on which the robots' survival and success depended.

Portions of this lengthy book will be daunting for nontechnical readers, not versed in formal logic and probability theory. (The author has already written a shorter version especially for them: *Information - A Very Short Introduction*, OUP 2010.) But the nontechnical portions are understandable to everyone and provide plenty of food for thought, especially inspiring the reader to try to meet the challenge of solving the "Gettier problem": Finding reasons strong enough to upgrade one's believing that something is true to knowing for sure that it's true, by immunizing it from the possibility that one is only right by accident. The author thinks this problem is provably insoluble, hence systematic uncertainty reduction is the best we can do. The reader is free to keep trying (and to hope it's not like trying to square the circle).