Four Forms of Information[1]

Terrence C. Stewart (tcstewar@chat.carleton.ca)
Institute of Interdisciplinary Studies: Cognitive Science
Carleton University, Ottawa, ON K1S 5B6 Canada

Andrew Brook (abrook@ccs.carleton.ca)
Institute of Interdisciplinary Studies: Cognitive Science
Carleton University, Ottawa, ON K1S 5B6 Canada

Abstract

Notwithstanding the fact that cognitive scientists are supposed to be unified by their view that cognition is information processing, four different definitions of the word 'information' are described herein. These are found in the various fields which cognitive science research draws upon. The definitions are not minor modifications of each other; they talk about very different things in very different ways. The same set of events can be described as information processing for four very different reasons, which each identify different aspects of the situation. However, a model is presented which ties these four together and illuminates their important distinctions. The ultimate goal is better understanding and communication between cognitive scientists.

Introduction

One of the common claims about cognitive science is that it is the study of information processing (Dawson, 1998). From this point of view, cognition *is* information processing, and that fact is what allows various researchers from other disciplines (philosophy, psychology, linguistics, computer science, etc.) to say that they are all doing cognitive science. This 'information processing' metaphor gives these researchers a common ground for communication, allowing them to see that they are studying the same thing from various different perspectives.

Of course, in order to be in agreement as to what the 'information processing' metaphor means, these researchers should also be in agreement as to what 'information' means. Here, unfortunately, is where this seeming harmony shows some rather serious signs of division.

This paper describes different definitions of 'information', and thus implies that there are many different conceptions as to what 'information processing' is. Furthermore, these differences are subtle enough that two people can seem to agree that something is 'information', but actually be in complete disagreement as to what it is about the situation that makes it 'information'. It is therefore important to take a close look at these different usages of the word.

Type 1: Reduction in Local Uncertainty

The first use of the word 'information' to be examined is one strongly associated with Information Theory. The advantage to this usage is that it is very precisely specified, since it is based in mathematics. This definition is that the amount of information (measured in bits) gained by observing an event is the negative log (base two) of the probability of that event occurring. So, if you flip a coin and discover that it has landed on heads, then you have gained one bit of information. If you flip three coins, and discover that they land heads, tails, heads, in that order, then you have gained three bits of information. If you roll a six-sided die and discover that it landed on a value less than 10, then you gain zero bits of information (Cover and Thomas, 1999).

There are a few important limitations implied by this usage. The first and most obvious is that it requires the prior knowledge of the probabilities of the various events occurring. This confines its use to situations amenable to statistical analysis. The second is that it is talking specifically about the *amount* of information in a message. This is a quantitative measure, and is limited to looking at a signal that is being passed from one place to another. So, for example, it could be used by cognitive scientists who want to measure the amount of information entering the eye and compare that to the amount of information transmitted down the optic nerve. To do the first task, we would look at the possible states of the cones and rods in the eye, do some statistical analysis to determine the probabilities of these various states, and multiply by the total number of rods and cones. For the second task, we do a similar process on the axons in a cross-section of the optic nerve. By comparing these numbers, we can see how much information is lost (or filtered out) between these two points.

[1]Carleton University Cognitive Science Technical Report 2003-06
http://www.carleton.ca/ics/techReports.html

The reason for calling this usage a "reduction in local uncertainty" is to highlight the fact that we are talking about reducing uncertainty *in the signal itself*. We have something very specific and immediate (the signal), and we do not know what the signal is going to be until it arrives. The rods and cones in the eye do not know what types of photons are going to hit them, so they are in some sense uncertain about the future. Once the photon does arrive, the uncertainty is reduced, and this is captured numerically as the amount of 'information' that has been received. The cells at the terminus of the optic nerve do not know what the signal from the nerve is going to be, but when it is received, they gain that amount of information. Within this sense of information, there is no room for reference to the outside world. Here, we are not interested in the deductions we may or may not be able to make from the signal itself. A different sense of information (type 4) will deal with this *aboutness* quality.

We can use this term to talk of storing or sending information, and whenever there is a quantity involved. For example, a standard computer floppy disk can contain 1.4MB (11,661,312 bits) of information. However, if I take a disk and fill it with *either* all zeros or all ones, but I don't tell you which, and then I give the disk to you to look at, you are only gaining *one* bit of information from the disk, not 11,661,312. This highlights the fact that this sense of information is not an objectively quantifiable one. The amount of information gained is dependent on one's prior knowledge of the probabilities of various different signals arriving.

Type 2: Exploitable Regularity

The next usage of 'information' deals with the idea that while most processes in nature are concerned mostly with the flow of energy (rain, gravity, friction, etc.), living creatures, and human beings in particular, are controlled by the flow of 'information'. This idea is grounded in the physics of entropy. The laws of thermodynamics imply that everything in the universe tends towards decay and breakdown, yet living creatures are able to maintain a stable internal structure. The reason they are able to do this is by exploiting regularities in the environment. So, in order to maintain your structure, you must obtain ordered materials from your environment (by eating and drinking). Importantly, it is not energy that you require from your environment; energy enters and leaves your body at roughly the same rate. Otherwise, the law of conservation of energy would imply that the amount of energy in your body would continually build up. However, the amount of order (a.k.a. 'information') in the materials you ingest is significantly greater than the amount of order in the materials that leave your body. Food is a very highly structured compound, with regularities that your body exploits to maintain itself.

It is this sense of information that is used by Chris Langton in his original Game of Life analysis. By using simple simulated universes with different "physical laws", he showed that complex "life-like" behaviour arose only in situations where the physical laws led to reasonably stable (regular) environments. However, if the laws made things too regular, then the complex structures would be unable to extract the order from their environments in order to maintain themselves. It is in this sense that the "flow of information" is important (Levy, 1992).

At this point, we can see that this sense of information is almost the exact opposite of the "reduction of local uncertainty" definition. Here, creatures make use of predictable patterns in the environment. Thus, a regularly patterned environment is more useful to an organism than a random one, since there is more order in it. However, by the first definition of information, the random input would contain more information. If I fill a disk with the results of my flipping a coin many times, there is much more information present (in the first sense) than if I fill it with a repeated pattern (such as all ones). However, there is much more order (the second sense of information) in the disk with a repeated pattern.

This sort of information is also discussed in chaos theory. In this case, we can find statements such as "When they spoke of systems generating information, they thought about the spontaneous generation of patterns in the world" (Gleick, 1998). Here, again, patterns are equated with information, although in this case the pattern is a highly complex one, which would be resistant to statistical analysis. For example, the digits in the decimal expansion of pi are a very particular pattern, even though they are seemingly random. In any case, for both chaos theorists and dynamic systems researchers, 'information' is equated with order, or anti-entropy.

Type 3: Data to be Manipulated and Transformed

The next sense of 'information' is associated mostly with cognitive psychology and its "information-processing model" of cognition (Solso, 1998). Here, we see the belief that cognition can be broken down into a series of interacting components, each of which take some input and produce some output. This is called information processing, and the inputs and outputs are referred to as information.

There is a clear distinction between this sense of information and the first sense. In the "reduction of local uncertainty" definition, information was something contained in the signal. Here, we are using the term 'information' to refer to the signal itself.

Another term for this would perhaps be 'data', as that term stresses the fact that we are interested in the objective truth about the state of the signal. What is being transmitted from one component in the brain to another is an objective physical truth (as much as any other truth about the real world).

An example of this sense of information would be our vision system. We take in input from our senses, and use a variety of different neural structures to produce some output. In this case, the output is a spatial model of our environment. In the intermediate steps, the data is massaged in various ways by various components to form colour data, stereoscopic disparity data, high-frequency detail, and so on. This whole process is referred to as information processing, not just the final result. The final result is 'information' in a different sense (type 4).

It might be said that this definition is most closely related to the sense referred to in the phrase "information technology". With this term, we are interested in the process of manipulation of data by each component of a computer system.

Type 4: Aboutness

We now finally turn to the most problematic aspect of information: *aboutness*. This is the idea that the signal we are receiving gives us information *about* something in the external world. For example, the signal entering a person's eyes gives that person information about the things around them: their colour, their shape, their motion, and so on.

This 'aboutness' seems to be the most contentious aspect of defining information. In many works (notably in philosophy, and especially in informational semantics), this property of 'aboutness' is the key issue at hand. Here we talk about 'natural meaning', where seeing tracks in the snow gives you the 'information' that an animal has passed this way. This is 'natural' since the relationship between the signal and the conclusion is governed by the laws of the natural environment. We also have 'non-natural meaning', where seeing the words "Fido is a dog" gives you the information that the creature picked out by the name 'Fido' is a member of the canine species. This is non-natural due to its dependence on human conventions for language (Dretske, 1999).

It may be instructive to further examine this last example. All four definitions of information that we have seen thus far would say that seeing the sentence "Fido is a dog" involves gaining information. However, they all mean very different things by this. For the first sense, we can look at the probability of seeing the sentence "Fido is a dog", as compared to all other sentences. In this case, we could conclude that seeing the sentence "Fido is a quadruped" gives you more information than seeing the sentence "Fido is a dog", since "Fido is a dog" is a more common sentence and is shorter (example from Cohen and Stewart, 1994). For the second sense of 'information', we could say that information is being used here because the very act of reading and interpreting the sentence is making use of regularities in the environment (such as the rules of grammar which make up the sentence, and the rules for what various letters look like). In the third sense, we are dealing with information since we are taking data in through our eyes, manipulating it, sending it to various parts of the brain, and having it enter some sort of long-term storage (memory), from whence it can modify future actions. But none of these first three cases have anything to do with the creature in the real world which we are gaining information *about*. After all, it is clear that "Fido is a dog" tells us more than "Fido is a quadruped", since it gives us more predictive power.

This is not to say that this fourth definition has nothing to do with Information Theory itself. Indeed, the framing of Information Theory makes it clear that there is an original source for the signal. However, it seems a rather unwieldy and unlikely task to apply probability theory to the complex causal relationship between the real-world creature 'Fido' and the visual or auditory signal "Fido is a dog" that reaches a person's senses.

This definition also has a runique property: it allows you to say that some information as *false*. The other senses cannot do this, since there is nothing for them to be false about. There can be a loss of information in these other senses, but only with this fourth sense is it meaningful to say that some bit of information is wrong. With this definition, information is representing something else, and we can evaluate the truth or falsehood of the information by seeing how good a job it does at this task of representation.

Bringing It All Together

Throughout this paper, we have seen four different definitions of the word 'information' commonly found in cognitive science literature. Of course, these definitions are not completely unrelated to one another; there is some sort of common thread that links them together, so that it does make a certain amount of sense to call them all 'information'. It would be beneficial to be able to see how these usages of the term relate to one another. The following diagram is an attempt to depict the relationship between these various different ideas.

In this diagram, the solid arrows stand for actual transmission of signals, which can be easily analysed in terms of probability and information theory (type 1). This includes all of the solid arrows in the diagram. The clearest of these is between the environment and the agent itself. These are the signals being picked up by the cognizing agent through its senses.

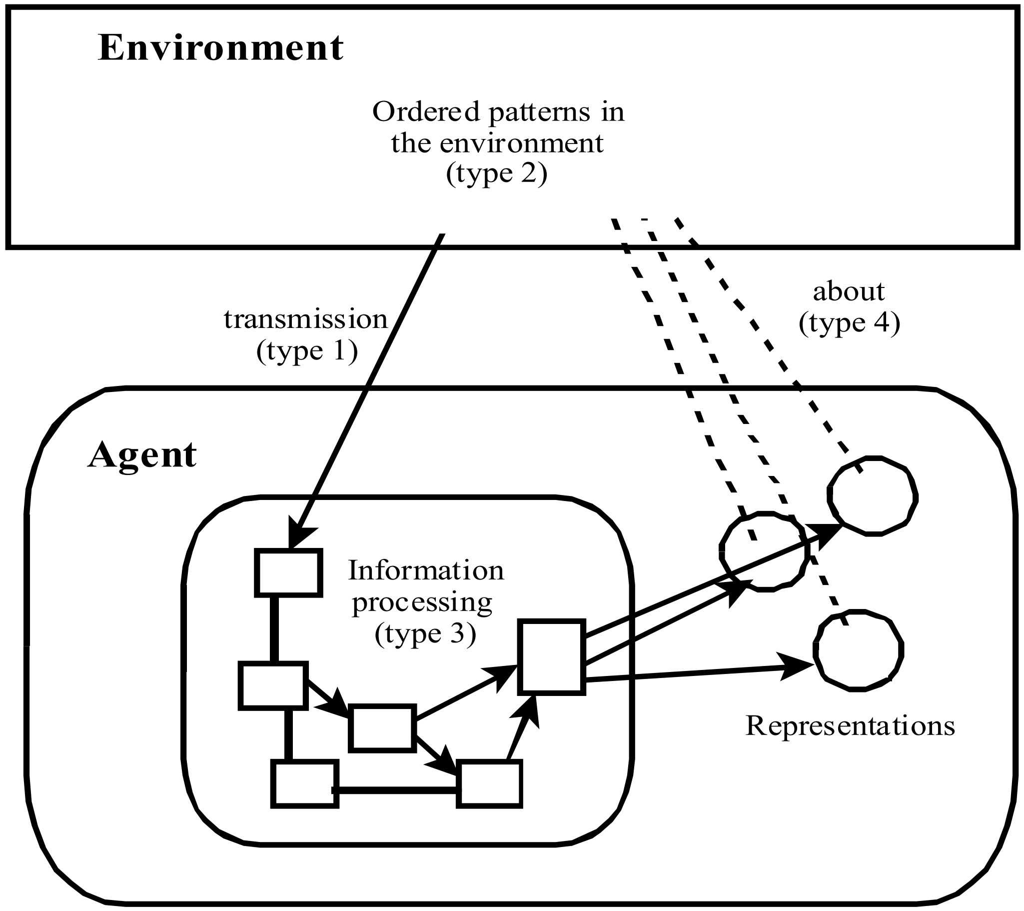

Figure 1: A framework for the four forms of information.

The reason that these signals are meaningful or useful to the agent is that these signals are generated by a pattern-filled, rule-governed external world. Furthermore, the agent itself must be a highly structured artifact in order to make use of these signals in the first place. It is here that we capture the type-2 sense of information. The transmitted signals come from an environment controlled by the regularities of physical laws.

The internal structure of the agent that interprets these signals can be described as an information-processing system (type 3). However, in this case, the information processing means *data*-processing, or *signal*-processing, since it is direct manipulations on the sensory data itself. This, in turn, produces data that can be said to be representations of the external world, since their purpose is to say something *about* the world. This is indicated in the diagram by dotted lines, to show that there is no direct causal relationship between the representations and the features of outside world that are being represented.

This lack of direct relationship is a rather important feature of type-4 information. First of all, it allows for the information to represent the world incorrectly. After all, if our internal notions about the world were always directly causally linked to the things they are attempting to represent, then we could never be mistaken about something. This is clearly not acceptable for a theory of human cognition.

A second advantage to the lack of direct relationship is that it allows representations of things that do not, in fact, exist. Reading a fictional novel gives us information (type 4) about events and persons which were never actually in the real world. Furthermore, the real world may not, in fact, separate itself into the concepts and categories that we use to represent it.

A Final Word

Information (type 2) exists in this paper, as it is a reasonably well-formed English-language document. By allowing your eyes to focus on this paper, some information (type 1) has been transmitted to you. Reading the document has clearly involved some information processing (type 3) occurring in your brain. It remains to be seen if the information (type 4) that may have thus been produced is actually useful. The hope is that this depiction of the different sense of information will improve the reader's awareness of these issues, and can be used to allow some researchers to make better sense of the usage of the word 'information' by fellow cognitive scientists.

References

Cohen, J. & Stewart, I. (1995). *The Collapse of Chaos*. London: Penguin Books.

Cover, T.M., & J.A. Thomas. (1999). Information Theory. In R. Wilson & F. Keil (Eds.), *MIT Encyclopedia of the Cognitive Sciences*. Cambridge, Mass.: MIT Press.

Dawson, M. (1998). *Understanding Cognitive Science*. Massachusetts: Blackwell.

Dretske, F. (1999). Informational Semantics. In R. Wilson & F. Keil (Eds.), *MIT Encyclopedia of the Cognitive Sciences*. Cambridge, Mass.: MIT Press.

Gleick, J. (1998). *Chaos: Making a New Science*. London: Penguin.

Levy, S. (1992). *Artificial Life: The Quest for a New Creation*. New York: Pantheon Books.

Solso, R. (1998). *Cognitive Psychology*, 5th edition. Boston: Allyn & Bacon.