

Bateson did not define "information" as "a difference that makes a difference"

(And he would have been rather silly if he had.)

Aaron Sloman

Last updated: 24 Jan 2011

Installed: 22 Jan 2011

This file is

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/information-difference.html>

also available, possibly badly formatted, as a PDF file:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/information-difference.pdf>

What follows is based on section 2.3 of this book chapter:

<http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#905>

Aaron Sloman,

What's information, for an organism or intelligent machine?

How can a machine or organism mean?,

in **Information and Computation,**

Eds. Gordana Dodig-Crnkovic and Mark Burgin,

World Scientific Publishers, 2011, New Jersey, pp 393--438

Introduction: the Myth

It is widely believed that the polymath Gregory Bateson defined "information" as "a difference that makes a difference". I think this is a myth, and he did no such thing.

This alleged definition is often quoted with approval by thinkers of different backgrounds, as can be seen by searching for occurrences of the phrase "a difference that makes a difference" in conjunction with "information". Sometimes the definition is attributed to others, presumably because they have quoted or used it.

Obviously the phrase "a difference that makes a difference" resonates powerfully with many people. Perhaps this is because it is a pointer to a very common and important kind of complexity, in which systems are composed of linked, tightly coupled, sub-systems whose causal relationships have the property that any event in one sub-system (e.g. some property, value or relationship changing, or a part being added or remove) has effects in other subsystems, possibly ripples of effects spreading out through the whole system.

Examples include a pebble hitting the previously flat surface of a pond, a fly wriggling in a spider-web, the speed of rotation of a cog wheel in some machine changing because of a change in friction, pushing a button causing an electric circuit to be closed, and a large collection of interacting electronic devices activated, an army being galvanised into battle by a light signal flashed on a hill-top, a news item causing share-prices all round the world to begin to fall, a rumour spreading quickly through a community and causing a mob to attack a house.

In all these cases it is reasonable to say that some information flowed round the system, initiated by the initial change (a difference occurring) and that the intermediate stages of such propagation depend on new intermediate changes/differences producing new effects elsewhere (including positive and negative feedback loops in some systems).

It is possible for a change or difference that is not temporal but spatial to have effects. For example, a geologist surveying some terrain may notice a transition across a boundary, which suggests the possibility of some desirable material or substance being available somewhere underground on one side of the boundary; or a farmer noticing a boundary separating two kinds of soil may be caused to sow crops only on one side of the boundary. In these cases, the **static** change produces **temporal** changes as a result of the occurrence of detection or observation of the static change: i.e. the trigger is temporal, even though what is triggered depends on something non-temporal. So here the original difference need not actually make a difference to anything -- whether it does or not will depend on something else: a happening.

Bateson could deal with quibble by replacing "A difference that **makes a difference**" with "A difference that **can make a difference**".

Despite all these interesting facts, and this minor repair, there are several problems with the proposed definition of "information".

- First of all, insofar as the physical universe is a connected whole, it is true of every temporal change that it has effects (i.e. makes a difference to something), and true of every spatial change that it can have effects. So if every difference is information, why do we need the concept of "information" in addition to the concept of "difference"?

A partial answer might be that we need two distinct concepts in order to avoid the circularity that would be manifest in attempting to define "difference" as "a difference that makes a difference", or "change" as "a change that causes changes". By having two words, one being defined and one used in the definition we avoid circularity. But what has been achieved?

- The proposed definition might make sense if we were trying to discuss what can be propagated through a network of causally connected mechanisms. Obviously in some cases energy is propagated. In other cases, for instance in a chemical plant, or the circulatory system of an animal or plant, matter is propagated (including both nutrients and waste). But our concept of information is distinct from our notions of matter and of energy, although information, matter and energy can interact, for instance when information received causes some large machine or other physical system to start moving. Matter beginning to move requires energy to be dissipated. All that can be triggered by information, as in the battle example, above.
- Moreover a serious flaw in the attempt to define "information" in terms of "change" or "difference" is that when we talk about information normally we are thinking of information as being about, or referring to, something other than the information. A recipe in a cook book provides information about how to make a cake, but the making of the cake is not the same thing as the information about making the cake, although a suitably knowledgeable cook watching the process could gain information from it.

In the cases where changes are propagated through a connected system, something may use detected changes as bearers of information about something else, but that does not make the changes themselves the information.

In the vast majority of cases, information bearers and information contents are distinct. Exceptions might be the sight of obviously wet paint on a wall providing information that there is wet paint on the wall, or the visible motion of an object providing information that the object is moving. However, in general the content of information is something other than the bearer of the information, even though all bearers of information about something other than themselves also provide information about themselves. For instance the english word "battle" when written can be used to provide information about the location of a battle, but it necessarily also provides information about its own spelling, though not its own pronunciation.

We could consider revising what Bateson appeared to be attempting to define , from "information" to "information-bearer".

"An information-bearer is a difference that makes (or can make) a difference"

However, this still leaves unexplained what information is: what information is "borne", or "carried", or "expressed" by an information bearer. For that, we need an explanation of what it is that can refer or denote, successfully or unsuccessfully, what it is that can be true or false, or inconsistent, what it is that can answer a question, or be the content of a decision or an instruction or command about what to do.

And lurking in the background to all these questions is the problem that "a difference" suggests something discrete: a step-change, as does Bateson's use (echoing Shannon) of the phrases "bit of information" and "elementary unit of information", suggesting that information is built out of indivisible chunks that are combined to create larger items. This does not square well with the common sense idea that information can be about things that vary continuously, such as pressure, or distance, or speed, or direction, or closeness to danger, and suggesting that there is no smallest information difference between two states, such as two possible velocities or locations for the same object.

Perhaps his answer would be that even if the **content** of some item of information can vary continuously, there isn't continuous variation between not having and having the information. It may be that if the information is complex it could be acquired in steps, but there would be some minimal first step, and each time the content of the information is increased (as opposed to merely being changed) there is a minimal possible increase, which is indivisible. If so, that minimal increment could not be acquired in parts or stages.

All that might be a way of defending Bateson's talk of bits or elementary units but I don't know whether he ever wrote something with precisely that interpretation.
(Comments from Bateson scholars welcome: I am not one.)

What Bateson Actually Wrote

While working on the "What's information" paper referenced at the top of this file I was mystified as to how someone as intelligent as I knew Bateson to be could have written something so obviously problematic and unhelpful. Since I had a copy of a collection of his papers I began to search for the reported definition. But as far as I can find, the report is actually a **mis**-report, for the quoted definition is not what he wrote. What I found he actually had written was something much more sensible.

In the 1972, Chandler Paperback edition of **Steps to an Ecology of Mind: Collected Essays in Anthropology**, Bateson described "a bit of information" and later "the elementary unit of information" as "a difference that makes a difference".

He does this in at least two of the essays, namely in "The Cybernetics of 'Self': A Theory of Alcoholism" and in "Form Substance and Difference". Notice that there is a difference between attempting to define (or say something definitive about) the word "information" and attempting to do it for more complex phrases like "a bit of information" and "the elementary unit of information", which he seems to take as different labels for the same thing, which he describes as "a difference that makes a difference". Similar or equivalent wording, with "information" always qualified as illustrated here occurs in several places in the book.

In all the contexts that I found, he is talking NOT about information in general but about an ITEM or UNIT or PIECE of information as a difference that makes a difference.

So it looks as if he accepted the assumption that information increments (or decrements) must be discontinuous, and that there is a minimal discontinuity -- one of the interpretations suggested above.

It seems that this remark is what is widely misquoted, and misrepresented as being about "information" rather than merely being about "a bit/unit of information".

In saying this sort of thing, Bateson seems to be thinking of any item of information as essentially a collection of "differences" that are propagated along channels. This is far too simplistic -- and perhaps too influenced by low level descriptions of computers and brains, though as indicated above, it may be a useful first approximation to a characterisation of the sort of thing that is capable of being used as a bearer of information, where the information itself could be expressed or carried by alternative structures: the information is not necessarily linked to a unique mode of expression, since different bearers for the same information content might be preferable in different contexts.

This still leaves the problem of saying what information is, which I shall not do here, because the answer is long and complex, as explained in the "What's information" paper cited above, which attempts to show that it is at best possible only to define "information" **implicitly**, by presenting a complete theory about information and its role in many systems.

This kind of implicit definition of deep and complex concepts is the only possibility for many scientific concepts, including "matter" and "energy" -- which is why "symbol-grounding" theory (another name for "concept empiricism"), is false, as explained in [this presentation](#).

The "[What's information](#)" paper attempts to present substantial portions of such a theory, though the task is not completed. In particular section 3.2 explains how theories can implicitly define the concepts they use and relates this to defining "information".

More specifically, what it means for **B** to express **I** for **U** in context **C** cannot be given any simple definition, in part because it is a generic polymorphic concept, which can be instantiated in different ways in different contexts.

Some people try to specify the meaning by saying **U** uses **B** to "stand for" or "stand in for" **I**. For instance, in an interesting contribution Barbara Webb writes "The term 'representation' is used in many senses, but is generally understood as a process in which something is used to stand in for something else, as in the use of the symbol '**I**' to stand for the author of this article"

That sort of definition of "representation" is either circular, if standing in for is the same thing as referring to, or else false, if "standing in for" means "being used in place of". There are all sorts of things you can do with information that you would never do with what it refers to and vice versa. You can eat food, but not information about food. Even if you choose to eat a piece of paper on which "food" is written that is usually irrelevant to your use of the word to refer to food.

Information about X is normally used for quite different purposes from the purposes for which X is used. For example, the information can be used for drawing inferences, specifying something to be prevented, or constructed, and many more. Information about a possible disaster can be very useful and therefore desirable, unlike the disaster itself.

So the notion of standing for, or standing in for is the wrong notion to use to explain information content. It is a very bad metaphor (based on some person or object taking the place of another in some process or situation), even though its use is very common.

We can make more progress by considering ways in which information can be used. If I give you the information that wet weather is approaching, you cannot use the information to wet anything. But you can use it to decide to take an umbrella when you go out, or, if you are a farmer you may use it as a reason for accelerating harvesting. The falling rain cannot so be used: by the time the rain is available it is too late to save the crops.

The same information can be used in different ways in different contexts or at different times. The relationship between information content and information use is not a simple one.

Comments, criticisms and suggestions welcome.

Maintained by [Aaron Sloman](#)
[School of Computer Science](#)
[The University of Birmingham](#)