TERRENCE W. DEACON
University of California, Berkeley

8

What is missing from theories of information?

Theories of information that attempt to sort out problems concerning the status and efficacy of its content – as it is understood in thoughts, meanings, signs, intended actions, and so forth – have so far failed to resolve a crucial dilemma: how what is represented could possibly have physical consequences. The legacy of this has been played out in various skeptical paradigms that either conclude that content is fundamentally relativistic, holistic, and ungrounded or else is merely epiphenomenal and ineffectual except for its arbitrary correlation with the physical properties of the signs that convey it. In this chapter I argue that the apparent conundrums that make this notion controversial arise because we begin our deliberations with the fallacious assumption that in order for the content of information to have any genuine real world consequences it must have substantial properties, and so must correspond to something present in some form or other. By contrast, I will show that this assumption is invalid and is the ultimate origin of these absurd skeptical consequences.

The crucial property of content that must be taken into account is exactly the opposite: its absence. But how is it possible for a specific absence to have definite causal consequences? A crucial clue is provided by Claude Shannon's analysis of information in terms of constraint on the entropy (possible variety) of signs/signals (Shannon, 1948; Shannon and Weaver, 1949). In other words, the capacity to convey information is dependent on a relationship to something that is specifically not produced. But such a change in the Shannon entropy of a physical medium is also necessarily a physical change, and this must be a product of extrinsic work. In addition, contra Shannon, even when there is no change in Shannon entropy where a change is possible, this can be informative because it indicates the absence of some specific form of extrinsic influence. Both conditions are determined with respect to a potential influence on the form of the semiotic medium that is extrinsic to it. These explicit extrinsic sources of potential signal constraint constitute the ground of the referential capacity by which information is defined. Moreover, I will argue that a process capable of interpreting the reference implicit in such signal constraints depends on coupling the context-dependent signal generation process to a specific extrinsic source of that constraint that is relevant to the existence and maintenance of this interpretive process. Such a linking relationship is provided by a self-sustaining non-equilibrium thermodynamic process, such as that characterizing living organisms. Such a process is capable of interpreting something as information about something else because such systems are necessarily open to and dependent on a precise correlation between intrinsic dynamics and extrinsic conditions. Thus constraints exhibited in certain substrates, which are in some way correlated both with maintenance of this unstable condition and with non-intrinsic or absent conditions that are relevant to this maintenance, can "matter" to its persistence. Failure to explain this relationship to absence is why past efforts to reduce information to a measurable physical substrate (Shannon information) or to treat it as a primitive nonphysical phenomenon (versions of phenomenology) both led to absurd consequences.

8.1 INTRODUCTION

It is often said that we are living in the "information age," but although we use the concept of information almost daily without confusion, and we build machinery (computers) and network systems to move, analyze, and store it, I believe that we still do not really know what *it* is. The ubiquity of the concept of information in our present age is largely a consequence of the invention, perfection, and widespread use of computers and related devices. In our everyday lives information is a necessity and a commodity. We routinely measure the information capacity of silicon, magnetic, or laser data-storage devices, and find ourselves guarding, sharing, or selling information. The importance of understanding and managing information has penetrated to the most technical and most mundane realms of daily life. Molecular biologists have recently mapped the molecular information "contained" in the human genome, and the melding of computer technology with molecular biology has spawned an entirely new field, dubbed bioinformatics, that promises radical new medical technologies and unprecedented threats to privacy. Even household users of internet communication are sensitive to the information bandwidth of the cable and wireless networks that they depend on for connection to the outside world.

It is my contention, however, that we are currently working with a set of assumptions about information that are just barely sufficient to handle the tracking of its most minimal physical and logical attributes, but which are insufficient to understand either its defining representational character or its pragmatic consequences. For more than half a century we have known how to measure the information-conveying capacity of any given communication medium, and yet we cannot give an account of how this relates to the content that this signal may or may not represent. These are serious shortcomings that impede progress in a broad range of endeavors, from the study of basic biological processes to the analysis of global economics.

It is a disturbing fact that, despite the centrality of the concept in our daily lives, we are entirely lacking a clear physical account that explains how information about some abstract concept can have massive and sometimes devastating physical consequences. Consider the concept of "patriotism." Despite the fact that there is no specific physical object or process that constitutes the content of this word, and nothing intrinsic to the sound of the word or its production by a brain that involves more than a tiny amount of energy, its use can contribute to the release of vast amounts of energy unleashed to destroy life and demolish buildings (as in warfare). This is evidence that we are both woefully ignorant of a fundamental causal principle in the universe and in desperate need of such a theory.

In many ways, we are in a position analogous to the early-nineteenth-century physicists in the heyday of the industrial age (with its explosive development of self-powered machines for transportation, industry, timekeeping, etc.), whose conception of energy was still framed in terms of ethereal substances, such as "caloric," "phlogiston," and the "élan vital" that were presumably transferred from place to place to animate machines and organisms. The colloquial notion of information is likewise conceived of in substance-like terms, as for example when we describe movement, storage, or sales of information. The development of the general concept of energy took many decades to clarify, even though the exploitation of energy was a defining feature of that era. The concept was ultimately demystified by recognizing that energy was not a substance, but rather a constant dynamical parameter that was transformed and yet conserved in processes of induced change. The conceptions of energy as ineffable ether or as special substance were abandoned for a dynamical relational account. With this reframing, many once-mysterious phenomena became subject to exact analysis and the basis for myriad new technologies.

Similarly, I argue that in order to develop a full scientific understanding of information we will be required to give up thinking about it, even metaphorically, as some artifact or commodity. To make sense of the implicit representational function that distinguishes information from other merely physical relationships, we will need to find a precise way to characterize its defining non-intrinsic feature – its referential content – and show how it can be causally efficacious despite its physical absence. The enigmatic status of this relationship was eloquently, if enigmatically, framed by Brentano's use of the term "inexistence" when describing mental phenomena.

> Every mental phenomenon is characterized by what the Scholastics of the Middle Ages called the intentional (or mental) inexistence of an object, and what we might call, though not wholly unambiguously, reference to a content, direction toward an object (which is not to be understood here as meaning a thing), or immanent objectivity.
>
> This intentional inexistence is characteristic exclusively of mental phenomena. No physical phenomenon exhibits anything like it. We can, therefore, define mental phenomena by saying that they are those phenomena which contain an object intentionally within themselves (Brentano, 1874).

As I will argue below, both the engineer's identification of information with the reduction of signal uncertainty and the intuitively attractive phenomenological conception of information as an irreducible "aboutness" relationship that is "always already there," simply take this enigmatic relationship to something not-quite-existent for granted. The first takes it for granted but then proceeds to bracket it from consideration to deal with physically measurable features of the informational medium. The second treats it as an unanalyzed primitive, and brackets its necessary physicality and efficacy from consideration in order to focus on intrinsic attributes. Neither characterization provides criteria that explicitly distinguish merely physical or logical relationships from those that convey information.

The concept of information is a central unifying concept in the sciences. It plays crucial roles in physics, computation and control theory, biology, cognitive neuroscience, and of course the social sciences. It is, however, employed somewhat differently in each field, to the extent that the aspects of the concept that are most relevant to each may be almost entirely non-overlapping. More seriously, the most precise and technical definition used in communication engineering, computational theory, and quantum physics completely ignores those features that distinguish information from any other causal relationship. This promiscuity threatens to make the concept of information either so amorphous that it provides no insight into the physical distinctiveness of living and mental relationships, or else licenses a retreat into a sort of methodological dualism.

Ultimately, the concept of information has been a victim of a philosophical impasse that has a long and contentious history: the problem of specifying the ontological status of the representations or contents of our thoughts. The problem that lingers behind definitions of information boils down to a simple question: How can the content (aka meaning, reference, significant aboutness) of a sign or thought have any causal efficacy in the world if it is by definition not intrinsic to whatever physical object or process represents it? In other words, there is a paradox implicit in representational relationships. The content of a sign or signal is not an intrinsic property of whatever physically constitutes it. Rather, exactly the opposite is the case. The property of something that warrants calling something information, in the usual sense, is that it is something that the sign or signal conveying it is not. I will refer to this as "the absent content problem." Classic conundrums about the nature of thought and meaning all trace their origin to this simple and obvious fact.

This relationship has often been framed as a mapping or correspondence between a sign or idea in the mind and this something else, which is not present. As countless critics have pointed out, however, this superficially reasonable account fails to identify any features of this relationship that distinguish it from other merely physical relationships. Consider the classic example of a wax impression left by a signet ring in wax. Except for the mind that interprets it, the wax impression is just wax, the ring is just a metallic form, and their conjunction at a time when the wax was still warm and malleable was just a physical event in which one object alters another when they are brought into contact. In these facts there is nothing to distinguish it from any other physical interaction. Something more makes the wax impression a sign that conveys information. It must be interpreted by someone. Unfortunately, this obvious answer is ultimately circular. What we invoke with an interpreting mind is just what we hope to explain. The process we call interpretation is the generation of mental signs interpreting extrinsic signs. So we are left with the same problem inside as outside the mental world. The problem of specifying how a specific content is both not physically present and yet inheres in some way in the sign and interpretive process is no better grounded in neurological processes than it is outside of brains.

8.2 MEANINGS OF INFORMATION

There is, additionally, a troublesome ambiguity in the term "information" that tends to confuse the analysis. This term is used to talk about a number of different kinds of relationships, and often interchangeably without discerning between them. It can refer to the sign or signal features themselves, irrespective of any meaning or reference, as in the information content in bits (binary digits) of the computer file encoding this chapter. This is sometimes called syntactic information. It can refer to what these bits refer to, as in the ideas I am hoping to communicate. This is sometimes called semantic information. And it can refer to that aspect of these ideas that is news to you the reader, and thus not merely redundant as it might be to experts in the field. This is sometimes called pragmatic information. Currently, the first of these meanings has grown in prominence, mostly as a result of our contemporary familiarity with, and dependence on, computing.

This document was created on multiple computers, and in the process I have shared draft versions back and forth with colleagues by sending the information over the Internet. But what exactly was sent? The answer is: a series of high and low voltage values, crudely analogous to Morse code, organized into sets of eight 1s and 0s that together code for an alpha-numeric character or some operation with respect to these. The computers at each end of the process are set up to encode and decode this sequence of shifts in voltage in the same way. For the user, this is accomplished invisibly. All we see are letters arranged on the screen. Clearly using the terms above we can say that syntactic information is being exchanged back and forth, in the form of numbers of distinguishable signals, and hopefully these signals also convey semantic and pragmatic information as well.

But would we still call it information in any of these senses if there were no people involved? And to make it a bit more like science fiction, would we still call it information if by wild coincidence a large collection of molecules just spontaneously came together to make two computers organized just this way, sending signals back and forth identical to those I have recently sent?[1] This would certainly not qualify as semantic or pragmatic information. What if the signal was composed of randomly generated gibberish rather than English? Would this exclude it from even being described as syntactic information? Presumably, for any finite length of randomly generated signals a suitable code could be defined that would change it into English text. Does this possibility change anything? Ultimately, there is nothing intrinsic to such a string of signals to distinguish it from encrypted English. Is the potential to be translated into a form that could be assigned a meaning sufficient to make it information? If so, then any signal, from any source, however produced and sent, would qualify as syntactic information.

In current computer technology, the rapidly flipping voltages that constitute the operations of a computer can be translated to and from the pattern of microscopic pits burned into a plastic disk or to and from the magnetically modified pattern of iron particles embedded in the surface of a spinning hard disk. But what if we happened upon a naturally occurring pattern of burn pits or magnetized iron atoms in a rock corresponding exactly to these patterns. Would these constitute the same information? Although they might be described as identical syntactic information they would not be likely to convey identical semantic information. Can it be information at all if it derives from a random process? Actually, yes. The chemical reactions caused by unknown molecules in a water sample being tested for contamination, or the radio waves reaching Earth from an exploding star, are typical of the sorts of signals that scientists are able to use as information. Both the patterns that we deliberately create in order to convey an idea and those we discover in nature, can convey information. Ultimately, this demonstrates that almost anything can qualify as information in the syntactic sense, because this is only an assessment of the potential to inform.

This most basic notion of information corresponds to the contemporary theory of information, originally dubbed the "mathematical theory of communication" by its discoverer Claude Shannon (1948). As we will see, Shannon's definition is more precise than this, but in essence it shows us why it is possible to treat any physical distinction as *potential* information, whether made by humans or the product of some mindless natural process. It specifies what features of

[1] The absurdity of this happening should also tell us something about the complexity hiding behind the notion of information being used here.

a physical material or process are necessary for it to serve as a medium for transmitting or storing information. So, in this sense, when we use the term "information" to refer to signals transmitted over a wire, or ink marks on paper, or the physical arrangement of objects at a crime scene, we are using a sort of shorthand. Without these physical features there is no information, but we are actually referring to something more than the physical pattern; something not present to which these present features correspond.

Identifying the features of physical processes that are the necessary conditions for something to be able to provide information helps to make sense of the enterprise of the natural sciences. Scientific observation and experiment are directed to the task of probing the patterns of things to discover information about how they came to have the properties they have. And the same physical objects or events can yield new information with each change of interpretive apparatus. This open-endedness is the result of there usually being vastly more information potential in natural phenomena than can ever be interpreted. To a brilliant sleuth equipped with the latest tools for materials analysis and DNA testing, almost everything at a crime scene can become a source of information. But being able to specify what physical properties can potentially serve to provide information does not help us to discern how it is that they can be about something else.

8.3 LOCATING THE INFORMATION IN INFORMATION PROCESSING

Shannon's analysis of the logical properties that determine the information capacity of a physical medium helps make sense of the concept of information in computer theory. In the most general sense, the possibility of computing depends on being able to assign referential value to some feature of a physical process and to map a specific logical operation to some physical manipulation of that feature with respect to others that have also been assigned reference. In this sense, one can even talk about arbitrary mechanical operations (or other physical processes) as potential computers. Just as one could come up with a coding scheme that could interpret an arbitrary sequence of signals as an English sentence, so it is possible to find a mapping between an arbitrary physical process and some symbol manipulation process. Of course some physical processes and mechanical (or electronic, or quantum) devices are better than others for this, especially when we desire flexibility in possible mappings. This mapping relationship – assigning reference – is crucial for distinguishing between computing and other merely physical processes. All physical processes are potential computations, but no physical process is intrinsically a computation.

It is in this most general sense that we can also describe mental processes as computing. And yet missing from this analogy is precisely the mapping relationship that distinguishes thought from computing. There is no separate homunculus external to the computation that assigns reference to the relevant physical differences of neural dynamics. Computation is generally described in terms of a syntactic conception of information, and yet it implicitly presupposes a semantic conception, although it can give no account of it. It is sometimes assumed that this referential mapping can be provided by the equivalent of robotic embodiment, so that the input and output of the computing is grounded in physical world constraints. But this can also be seen as an extension of the physical mechanism underlying the computation to include causal events "outside" some arbitrarily defined boundary separating the "computing device" and the physical environment.

Describing both physical and mental relationships in computational terms is only problematic if this presupposition of mapping is ignored or assumed to be intrinsic. The result is either eliminative reduction or cryptic panpsychism, respectively. In either case, if any physical event is considered to be a computation and the mind is merely a special purpose computer, then the mind–body problem dissolves. But there is a troubling implication to this collapse of the concept of information to its syntactic meaning only. In such a uniformly informational universe there is no meaning, purpose, value, or agency. In this informational cosmology, networks of informational causality are still just as blindly mechanical as in any Laplacian universe.

To escape this deflationary view of an information universe blindly mechanistically computing us, many have turned to quantum physics to loosen the bonds of mechanistic determinism, both in discussions of consciousness and in terms of information processes. Quantum events appear postponable until they are observed, and quantum objects can be both independent and correlated (entangled) at the same time. Thus notions of causality and of information about that causality appear to be inextricably linked at this level of scale.

For example, in the dominant (although not the only) interpretation of quantum mechanics, events in the world at the quantum level become real (in the sense of being susceptible to classical analysis) only when they are measured. Before such an intervention, no explicit single state exists, only a field of potentiality. This is exemplified by the famous Schrödinger's cat paradox, in which the prior death of a cat in a closed box is dependent on an as-yet-unmeasured quantum state. In this interpretation it is presumed that neither macroscopic state exists until the quantum event is later measured (that is, by an observer). Measurement information about this quantum state is thus treated as a fundamental causal determinant of the transition from quantum indeterminacy to classical determinism. Similarly, another strange quantum phenomenon – quantum entanglement – shares other features with the correspondence relationship associated with information. In simplified terms, this involves an apparent instantaneous correspondence of measurement consequences between particles that are separated and non-interacting. So one might also argue that it exemplifies a sort of intrinsic representation relationship.

These are counterintuitive phenomena, which challenge our normal conceptions of causality, but do they explain the higher-order senses of information? Unfortunately, they do not actually resolve the paradox of the absent content. These features in the quantum realm (for example, superposition, entanglement) resemble correspondence-mapping relationships. Thus we might be tempted to see this as a referential relationship that is intrinsic to the quantum physical relationship. But physical correlation is not aboutness. Whereas measurements of particles that affect measurements of other particles exist in something like a correspondence relationship, this alone does not make one *about* the other, except to an external observer interpreting it. The aboutness does not exist in the interstices between indeterminate quantum events any more than between the gears of a clock, because it is not an intrinsic feature. So in both classical and quantum computation only the syntactic concept of information is invoked. There is nothing intrinsic to computation processes that distinguish them from other physical processes, and nothing intrinsic to the quantum or classical physical features that are manipulated in computations that make them about other features of the world.

8.4 IS INFORMATION PHYSICAL?

There is something correct about the link between information and the fabric of causal processes of the world – whether determinate or intrinsically statistical – but there is something missing too. There is something more that we assume when we describe something as information, and indeed something absent from the physical processes and patterns that we recognize as conveying (though not fully constituting) information.

The search for a link between information and physical causality in general requires that we identify a precise physically measurable correlate of information. This is necessary in order to solve engineering problems that involve information systems and to address scientific issues concerning the assessment of information processes in natural systems. A first solution to these practical challenges was formally provided by a Bell Labs researcher in the 1940s: Shannon and Weaver (1949). His "mathematical theory of communication" demonstrated that the capacity of a communication medium to convey or store information could be precisely measured, and that even informational error-correction can be accomplished without any reference to informational content. This laid the foundation for all of modern information technology, but it also left us with a deflationary theory of information, from which content, reference, and significance are excluded and irrelevant.

Claude Shannon's introduction of a statistical approach to the analysis of signals and their capacity to carry information has stood the test of time with respect to any practical accounting of how much information a given medium can be expected to store or convey to an interpreter. Unfortunately, because this analysis excluded any reference to problems of defining content or significance, it has led to the rather different uses of the term that we have been struggling to define and which are often a source of theoretical sleight of hand. By bracketing issues of reference and significance Shannon was able to provide an unambiguous, interpretation-free measure of what might be called the information-bearing capacity (as distinguished from information itself). Not only does this work for human-made communication processes, it also usefully conveys the potentiality of any physical distinction to provide information, such as might be discovered by scientific experiment or detective work. But for this purpose it had to stop short of conveying any sense of how information could come to be about something. And there is good reason to have avoided this. For different interpreters or for different scientific instruments the same physical distinction can provide information about different things, or can be irrelevant and uninformative. What something is about and how this relationship is mediated are explicitly a function of external relations and thus are not able to be mapped to any intrinsic properties.

Information is by definition something in relation to something else, but in colloquial use the term can either refer to what is conveyed or what provides the conveyance. If, as in Shannon's sense, it refers only to the latter, then its aboutness and its significance are assumed potentialities but are temporarily ignored. The danger of being inexplicit about this bracketing of interpretive context is that one can treat the sign as though it is intrinsically significant, irrespective of anything else, and thus end up reducing intentionality to mere physics, or else imagine that physical distinctions are intrinsically informational rather than informational only *post hoc*, that is, when interpreted.

However, although Shannon's conception of information totally ignores the issue of what information is *about*, or even that it is about anything, his analysis nevertheless provides an important clue for dealing with the absent-content problem, specifically by showing that absence could have a function at all. This clue is provided by Shannon's negative characterization of information. Shannon's measure of the potential information conveyed by a given message received via a given communication channel is inseparable from the range of signals that could have been received but were not. More precisely, Shannon information is defined as the amount of uncertainty that is removed with the receipt of a given signal. So to measure information requires comparing the potential variety of signals that could have been transmitted to what was transmitted. Perhaps the most important contribution of this analysis was his recognition that measuring the potential signal variety was mathematically analogous to measuring the entropy of a physical system, such as an ideal gas. Following the advice of the mathematician John von Neumann, he decided to call this variety of possible states the "entropy" of the signal medium (or "channel," as he described it, using the model of a communication channel between a sender and recipient of a message). This decision, like the decision to define information capacity with respect to signal constraint, has led to many confused debates about the physical correlates of information. But these analogies are also important hints for expanding the concept of information so that it can again embrace the very features that had to be excluded from this engineering analysis.

8.5 TWO ENTROPIES

By defining information in negative terms with respect to the potential variety of what could have occurred, Shannon has inadvertently pointed us toward the relevant physical property of the signal or sign medium that provides access to what it is about, which is also a negative attribute. The key to seeing this link is simply recognizing that the representing medium, whatever form it takes, is necessarily a physical medium. It is something present that is taken to be about something not immediately present. The reduction in signal entropy has the potential to carry information because it reflects the consequences of physical work and thus the openness of this physical signal medium to extrinsic influence. In thermodynamic terms (Boltzmann, 1866), a change in the state of a physical system that would not otherwise occur is inevitably characterized by a local reduction in its physical entropy (I will describe this as "Boltzmann entropy" to distinguish it from Shannon entropy) resulting from work done on that system from outside. According to Shannon, the information-bearing capacity of a signal is proportional to the improbability of its current physical state. But an information medium is a physical medium, and a physical system in an improbable state reflects the effects of prior physical work, which perturbed it from some more probable state or states. In this way, the Shannon information embodied in signal constraints implicitly represents this work. It is in this sense that Shannon entropy is intimately related to Boltzmann entropy. It is explicitly the change in the Boltzmann entropy of the medium that is the basis of signal reference, because this is necessarily a reflection of some extrinsic influence (Deacon, 2007, 2008).

However, the relationship is more subtle than just the consequence of physical work to change a signal medium. Although this relation to work is fundamental, referential information can be conveyed both by the effect of work and the evidence that no work has been done (Deacon, 2007, 2008). Thus, no news can be news that something anticipated has not yet occurred. This demonstrates that Shannon information and referential information are not equivalent. This is again because the signal constraint is not something located *in* the signal medium: it is rather a relationship between what is and what could have been its state at any given moment. A reduction in variability is a constraint, and a constraint is in this case not an intrinsic property but a relational property. It is defined with respect to what is not present. So implicitly, a physical system that exhibits constraint is in that configuration because of extrinsic influences – but likewise if the sign medium exhibits no constraint or change from some stable state, it can be inferred that there was no extrinsic influence doing work on it. The relationship of present to absent forms of a sign medium embodies the openness of that medium to extrinsic intervention, whether or not any interaction has occurred. Importantly, this also means that the possibility of change due to work, not its actual effect, is the signal feature on which reference depends. This is what allows absence itself, absence of change, or being in a highly probable state to be informative.

Consider, for example, a typo in a manuscript. It can be considered a reduction of referential information because it reflects a lapse in the constraint imposed by the language that is necessary to convey the intended message, and yet it is also information about the proficiency of the typist, information that might be useful to a prospective employer. Or consider a technician diagnosing the nature of a video hardware problem by observing the way the image has become distorted. What is signal and what is noise is not intrinsic to the sign medium, because this is a determination with respect to reference. But in either case the deviation from a predicted or expected state is taken to refer to an otherwise unobserved cause. Similarly, a sign that does not exhibit the effects of extrinsic influence – for example, setting a burglar alarm to detect motion – can equally well provide information that a possible event (a break-in) did *not* occur. Or consider the thank you note not sent, or the tax return not submitted on time. Here, even the absence of a communication is a communication that can carry significance and have dire consequences.

In all cases, however, the referential capacity of the informational vehicle is dependent on physical work that has, or could have, altered the state of some medium open to extrinsic modification. This tells us that the link between Shannon entropy and Boltzmann entropy is not mere analogy or formal parallelism. It is the ground of reference.

8.6 DARWINISM AND INTERPRETATION

Up to this point of the analysis it has been assumed that the relationships being described have involved signs and signals, and not merely physical events chosen at random. But in fact, *none* of the criteria specified thus far actually distinguishes events and objects that convey information from those that do not. They are *requirements* for something to be information about something else, but they do not in themselves constitute it. Shannon described the necessary conditions for something to have the potential to convey information: providing a syntactical conception. Even the linkage to physical work, while a necessary requirement for referential capacity, is like Shannon's criterion, only a necessary but not sufficient feature of reference: that is, a semantic conception of information. But of course, not just any alteration of entropy constitutes a reference relationship. Although any physical difference *can* be interpreted as information about something else – whether it is the state of the mud on someone's shoes or the presence and evenness of the microwave background radiation of the universe – this is not an intrinsic feature, but something entirely relative to how it is interpreted. This post-hoc dependency does not diminish the necessity of these attributes. It merely demonstrates that they are not sufficient. And yet, as we have seen, reference depends on the responsiveness of the information medium to physical change.

A physical difference becomes informational when it plays a modulatory role in a dynamic process. The potential to inform is dependent on the Shannon–Boltzmann criteria just discussed, but this potential is only actualized as it influences a specifically structured dynamical process. Although we commonly talk about this as an interpretive process, invoking folk psychological assumptions, this is still for the most part a mere promissory note for a missing theory about what dynamical organization is sufficient to constitute such a process. And this heuristic immediately becomes problematic when we attempt to expand the usage to domains such as molecular biology, in which the presence of a homuncular interpreter cannot be invoked.

A key insight into the conditions for interpretation was provided by Gregory Bateson in an oft-cited aphorism proposed as a characterization of information: Information is "a difference that makes a difference" (Bateson, 1972). It is no coincidence that this phrase also would be an apt description of mechanical or thermodynamic work. Implicit in this phrase is the notion that information can be used to change things. And in this sense it has the potential to control work. So putting this idea about the physical basis of interpretation together with the Boltzmannian criterion for referentiality we have: "a medium that is susceptible to being modified by physical work, which is used to modify the state of some other dynamical system because of that system's sensitivity to changes of this medium, and which is differentially capable of performing work with respect to such a change."

This is a complicated definition, but even so it lacks detail concerning the nature of such a dynamical architecture, and so it is still incomplete in a couple of important respects. These have to do with informational relevance and the notion of function from which the normative aspect of information arises: that is, the pragmatic conception of information. Before trying to address these limitations, however, we need to flesh out the requirements for a system with the capability to differentially perform work with respect to a given signal state. This is because these requirements will ultimately provide grounding for these additional features, and the basis upon which the possibility for specific reference can arise.

It is a simple rule of thermodynamics that to be capable of performing work a system must be in a non-equilibrium state. So, any full explanation of what constitutes an interpretive process must include a central role for non-equilibrium dynamics. But non-equilibrium conditions are inherently transient and self-undermining. For a non-equilibrium process to persist, it must rely on supportive environmental conditions (for example, a source of free energy and raw materials) to make up for this spontaneous degradation. In this sense, like the signal medium, it must be open to influences extrinsic to itself: for example, a larger thermodynamic context to which it is well fitted. Thus the presence of a system maintaining a persistent non-equilibrium state capable of performing work with respect to a source of information entails the presence of environmental conditions that promote it.

This is important for the constitution of an interpretive process for two additional reasons. First, the openness to context of an interpreting system cannot merely be sensitivity, as is the case for an information medium. Persistence of a non-equilibrium thermodynamic system requires a quite specific matching of dynamical organization with extrinsic supportive conditions. In other words, there must be a necessary correspondence in form between system and context. Second, for a dynamical non-equilibrium system to be persistent there must be some maintenance of self-similarity and thus boundedness to it. It must have a unit identity in at least a loose sense.

Understanding that a process capable of generating information necessarily involves non-equilibrium dynamics also provides a way to address the normativity issue that is implicit in a pragmatic conception of information. Normativity in its various aspects has been a non-trivial problem for correspondence and mapping theories of reference. Following Bickhard (1998, 2000, 2003), I would argue that the normativity that defines representational error is an emergent property of the relationship of the Shannon–Boltzmann referential relationship with respect to the organization of the non-equilibrium processes that interpret it. This follows because of the intrinsic dependence on specific environmental conditions that are required for such a dynamical system to persist. To the extent that a particular signal interpretation effectively contributes to this end, and thus aids the successful maintenance of this supportive correlation, then that particular interpretive response to that particular state of the signal medium will also persist. Of course, the opposite is also possible; hence the possibility of misinterpretation.

But even in the simplest case this presupposes a non-equilibrium process that is precisely organized with respect to both supportive environmental conditions and to some feature of that environment that tends to correlate with those conditions. The establishment of such a reliable relationship is then the transformation of an incidental physical relationship into an information relationship. It is probably the case that this matching of specific referential relationship with a specific dynamical modification of the capacity to perform work can be achieved spontaneously only by an evolutionary process. Not surprisingly, then, this analysis suggests that the generation of information in the full sense is an emergent property of life. Of course, this does not exclude the infinitely many ways that information can be generated and manipulated, indirectly, with respect to living processes. Yet these, too, must at least indirectly embody these same basic criteria. It is in this sense that both the syntactic (Shannon) and semantic (Shannon–Boltzmann) conceptions of information are ultimately dependent on a pragmatic (Shannon–Boltzmann–Darwin) conception (Deacon 2007, 2008). In this way, the process of evolution, in its most general form, can be understood as the necessary originative source for information. Where there is no evolutionary dynamic there is no information in the full sense of the concept.

8.7 INFORMATION EVOLVES

The claim that evolution itself constitutes an information creation process needs to be unpacked in order to complete this analysis. There is yet an additional curious – but in hindsight not unexpected – parallel between the Shannonian determination of information capacity and the evolutionary determination of fittedness. Both involve the relationship between a potential and a realized variety. Natural selection depends on the generation of functionally uncorrelated (aka "random") variety of forms (genotypes and phenotypes) followed by the reduction of this variety due to death or the failure to reproduce. In our analysis of the Shannon–Boltzmann relationship, the referential potential of a signal medium was shown to be a consequence of the way extrinsic factors reduce its potential variety (entropy). In the process of natural selection an analogous role is played by conditions in the environment that favor the reproduction and persistence of some variants and not others. It is in this sense that we feel justified in claiming that the traits that are present in any given generation of organisms are adaptations to (favorably correlated with) certain of those conditions. Metaphorically speaking, they could be said to be "about" those conditions.

There are deep disanalogies, however, that are crucial to explaining why this process generates new information. First, the initial variety is not signal variety, not merely variations of some passive substrate. The variety that is subject to the constraining influence of natural selection involves variation of processes and structures associated with the interpretive system itself. Second, the non-equilibrium dynamics of organisms competing with one another to extract or sequester resources is the source of the work that is the basis for this reduction in population "entropy." And third, the variety that is subject to selection is being generated anew in each generation by virtue of what in Shannonian terms would be considered noise (that is, mutations and recombinations) introduced into the signal (that is, genetic inheritance). Thus what amounts to uncorrelated corruptions of the genetic signal and incidental physical attributes can become information to the extent that they result in variations of the interpretive–adaptive process that happen to embody correlated predictive correspondences between the dynamics of interpretation and the supportive conditions enabling this interpretation. The capability of the Darwinian process to generate new information about organism–environment (and by extension the interpreter–context) interdependency is the ultimate demonstration of the post-hoc nature of information. This evolutionary transformation of noise into information is the ultimate demonstration that what makes something information is not intrinsic to any features of the information-conveying medium itself. It is irreducibly relational and systemic, and at every level of analysis dependent on a relationship to something not present.

8.8 CONCLUSIONS

The "intentional inexistence" of the content of a thought, the imagined significance of a coincidental event, the meaning of a reading from a scientific instrument, the portent of the pattern of tea leaves, and so on, really is something that is not there. In this sense the Cartesian-derived notion that the content of mind is without extension, whereas the brain processes that realize this content do have extension, is at least partly correct. But to say that this absent content is extensionless is not quite right. The non-produced signal (that is, reduced entropy) that is the basis for Shannonian informative capacity, the non-present work that was or was not the basis for the reference of this signal, and the interpretive options (organism trait variations) selected in an evolutionary process, all have a definite negative extension in the sense that something specific and explicit is missing. In other words, like the space within a container, these are absences that are useful because of the way what is present can exemplify them.

The nearly universal tendency to attribute intentional phenomena to a disembodied realm is a reflection of this negative defining feature, but the apparent paradoxes this creates with respect to the physical efficacy of informational content is the result of misinterpreting this negative feature as though it is in some way substantial in a separate disembodied realm. The modern shift to abandoning all consideration of intentionality in definitions of information, as the concept has come to be used in the sciences, in order to focus entirely on the material–logical attributes of signal differences has correspondingly stripped the concept of its distinctive value and has led to a reduction of information relationships to relationships of physical difference. As a result this most common and undeniable feature of our existence is often treated as though it is epiphenomenal. Even the recent efforts to reframe intentionality with respect to its embodiment, effectively recapitulates a cryptic form of dualism in terms of a variant of dual aspect theory. But avoiding addressing the "inexistence" problem in these ways guarantees that the real-world efficacy of information remains inexplicable.

Like so many other "hard problems" in philosophy, I believe that this one, too, appears to have been a function of asking the wrong sort of questions. Talking about cognition in terms of the mind–brain – implying a metaphysically primitive identity – or talking about mind as the software of the brain – implying that mental content can be reduced to syntactic relationships embodied in and mapped to neural mechanics – both miss the point. The content that constitutes mind is not *in* the brain, nor is it *embodied* in neuronal processes in bodies interacting with the outside world. It is, in a precisely definable sense, that which determines which variations of neural signaling processes are *not* occurring, and that which will in a round-about and indirect way help reinforce and perpetuate the patterns of neural activity that are occurring. Informational content distinguishes semiosis from mere physical difference. And it has its influence on worldly events by virtue of the quite precise way that it is *not* present. Attempts to attribute a quasi-substantial quality to information or to reduce it to some specific physical property are not only doomed to incompleteness, they ultimately ignore its most fundamental distinctive characteristic.

REFERENCES

Bateson, G. (1972). *Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology*. Chicago, IL: University of Chicago Press.

Bickhard, M. H. (1998). Levels of representationality. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(2): 179–215.

Bickhard, M. H. (2000). Autonomy, function, and representation. *Communication and Cognition – Artificial Intelligence*, 17(3–4): 111–131.

Bickhard, M. H. (2003). The biological emergence of representation. In *Emergence and Reduction: Proceedings of the 29th Annual Symposium of the Jean Piaget Society*, eds T. Brown and L. Smith. Hillsdale, NJ: Erlbaum, 105–131.

Boltzmann, L. (1866). The Second Law of Thermodynamics. Reprinted in *Ludwig Boltzmann: Theoretical Physics and Philosophical Problems, Selected Writings*, ed. B. McGuinness, trans. P. Foulkes (1974). Dordrecht: Reidel Publishing Co., 13–32.

Brentano, F. (1874). *Psychology From an Empirical Standpoint*. London: Routledge & Kegan Paul, 88–89.

Deacon, T. (2007). Shannon-Boltzmann-Darwin: Redefining information. Part 1. *Cognitive Semiotics*, 1: 123–148.

Deacon, T. (2008). Shannon-Boltzmann-Darwin: Redefining information. Part 2. *Cognitive Semiotics*, 2: 167–194.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27: 279–423, 623–656.

Shannon, C., and Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.