To appear in In L. Floridi (Ed.) The Routledge Handbook of Philosophy of Information. London: Routledge.

Chapter 29: Information-Theoretic Philosophy of Mind

Jason Winning and William Bechtel

1. Introduction

The branch of philosophy known as "philosophy of mind," at its core, is concerned with two closely related questions: What sort of thing is a mind or mental state? And how are these related to the non-mental? We know some of the characteristic activities of minds: thinking, remembering, dreaming, imagining, etc. We consider certain types of things to be *mental*, i.e., to exist only in minds, as mental "states" or properties, or "mental contents": thoughts, memories, desires, emotions, and what philosophers refer to as "qualia" (that aspect of an experience that one refers to by the phrase "what it is like to undergo it"; see chapter 30). We also consider certain kinds of things to be non-mental, such as rocks, tables, and rain drops. These do not have minds, do not constitute minds, and cannot exist in minds as mental states or properties. Minds, and the states they can undergo, are *mental*; things that cannot be minds and cannot be undergone by minds are *non-mental*.

But is there a third category of thing, that can exist either inside or outside of minds? Consider the question of how *communication* is possible between minds. Warren Weaver, one of the pioneers of mathematical communication theory, defined communication in mental terms as "all of the procedures by which one mind may affect another" (1949: 2). At least in some instances, communication might be described as a process whereby the contents of mental states (such as thoughts or hopes) can be *sent* from one mind to another. In order for this to occur, there must be some *medium* by which the contents are conveyed—by which the mental contents are able to be sent between minds; naturalistic philosophers do not believe that a semantic content can merely float in space-time by itself, or that the contents of mental states can jump immediately from one mind to another without crossing some type of non-mental medium. These "contents," of course, are not strictly speaking *mental* contents when they are being conveyed outside of a mind; instead of residing in a mind, they reside within (or as it is usually put, are "carried by" or "tokened by") non-mental *vehicles*. 'Vehicle' is a term introduced by Dennett (1969) to refer to any given manifestation of a content; the vehicle makes it possible for the content to be tokened (instantiated) at a given place and time. Such vehicles might include black ink on a page, or patterns of electricity running along a wire. We also have a word to refer to such tokenized contents while they are being conveyed from place to place, or even when they are being stored or processed outside of minds: *information*. Information is here being understood broadly, as that which has a vehicle and a content.

A further prefatory terminological clarification is in order. The concepts of informational *content* and *vehicle* are also closely related to the concept of *representation*. Sometimes authors use the term 'representation' to refer to the relation that holds between *any* vehicle and its content, some restrict its use to the relation holding for a particular category of these, and some use it to refer to either the content or the vehicle standing in such a relationship. The following remarks by Dretske are pertinent here:

> There are representational vehicles—the objects, events, or conditions that represent—and representational contents—the conditions or situations the vehicle represents as being so. In speaking about representations, then, we must be clear whether we are talking about content or vehicle, about what is represented or the representation itself. It makes a big difference. In the case of mental representations, the vehicle (a belief or an experience) is in the head. Content—what is believed and experienced—is (typically) not. … When I speak of representations I shall always mean representational vehicle. (2003: 68)

There is a large debate about the nature of the *representation* relation and how it is (conceptually or otherwise) related to content. Here, we wish to largely sidestep this debate; we will simply follow Dretske in using 'representation' as synonymous with 'representational vehicle' and we will consider representations to be one species of vehicle, possibly among others.

We also wish to sidestep the internalism/externalism debate about mental content. The issue of whether some mental content is "outside the head" is independent of the vehicle/content distinction. Even if mental contents are themselves mind-external, or if they are fixed by mind-external facts, we can still make sense of those contents being communicated to another mind, or of the mental vehicles in another mind taking up that same content. Either way, such a process may involve this content being encoded and transmitted by means of mind-external informational vehicles.

2. Communication, Control, and Computation

Karl Pribram has claimed that information is a notion that is "neutral to the mind/brain dichotomy" (1986: 507). As such, it can be a powerful bridge principle in philosophy of mind's quest to answer the questions about the identity of the mental and the relation between the mental and the non-mental. But it is a bridge that has only become available in the last century or so as a result of the development of technologies for sending, generating, processing, storing, and consuming information. These advancements created the practical need, in the mid-20th century, for scientific theories of communication (e.g., Shannon 1948), control (e.g., Wiener 1948), and computation (e.g., Turing 1936).

2.1. Communication

With the advent of the telegraph and telephone as technologies for transmitting messages, questions arose as to the requirements for reliable transmission. Pioneers in the field of mathematical communication theory such as Harry Nyquist, R. V. Hartley, and Claude Shannon realized the need for exact measurement of information transmission capacity, regardless of the specific technology involved. In doing so they abstracted away from the semantic interpretation of the messages themselves and introduced the crucial conceptual distinction between semantic meaning (sometimes referred to as "semantic information"; see chapter 6) and the statistically quantifiable reduction of uncertainty that is achieved in a given communication process. By isolating the latter, important mathematical results from thermodynamics and statistical mechanics could be applied directly to problems faced by communications engineers.[1]

Shannon's (1948) mathematical communication theory (which many authors simply refer to as *information theory*, see chapter 4) was very influential on scientists and philosophers, including those mentioned in the following sections. In particular, it helped convince them that informational processes are as much at home in a purely naturalistic theory as any physical phenomena. It is no surprise that the most influential philosophical treatment of information, Dretske's 1981, discussed in a later section, draws heavily from Shannon's work.

2.2. Control

An important use of information is to enable one system to control another system. A basic but highly potent form of control is negative feedback, whereby the activity of a system is corrected in light of information received back by the system about the results of the activity. For example, as you move towards a target you acquire information about the discrepancy between your current location and the target and use that to correct your motion. The potency of negative feedback was independently recognized in engineering and in biology. A particularly prominent mechanical engineering example was James Watt's invention in the 1780s of the centrifugal governor (Figure 29.1) to control the speed of a steam engine. When the flywheel turned too fast, centrifugal force would cause the angle arms attached to a spindle to extend up and out and, via a linkage mechanism, close the valve. When it moved too slowly, the arms would drop and cause the valve to open.<FIGURE 29.1 HERE>

In biology, Claude Bernard, in attempting to answer the challenge of vitalists who thought that only a non-physical force could explain the ability of organisms to resist destruction by physical processes, introduced the idea that organisms maintain the constancy of their internal environment by responding to conditions that disrupt them, a process Walter Cannon termed *homeostasis*. Drawing upon his own use of negative feedback in designing targeting systems for weapons, as well as Cannon's application of negative feedback in explaining how animals maintain homeostasis, Norbert Wiener (1948) elevated the principle of negative feedback to a general design principle for achieving control in both biology and engineering. He coined the term *cybernetics* to describe this framework and inspired an interdisciplinary movement involving mathematicians, engineers, neurophysiologists, and psychologists, amongst others, devoted to understanding the control of activities through negative feedback.

2.3. Computation

The term *computation* originally applied to the activity of humans in computing sums by applying rules to symbols written on a page. Given the tediousness of the activity, there has been a long history of humans attempting to automate the process. An important milestone was Charles Babbage's design of the difference engine to compute polynomial functions, and the analytical engine, which could be programmed by punch cards to execute any computation. Babbage did not succeed in building either of these devices. Alan Turing (1936) proposed an extremely simple machine that could be configured to compute any function. It consists of an indefinitely long tape on which a string of symbols (e.g., 0s and 1s) are written and a read/write device that can move along the tape while reading from it or writing to it. The read/write device can enter into a finite number of states and has rules specifying what to do when, in a given state, it reads a particular symbol from the tape. A *Turing machine* (as it came to be called; see chapters 3 and 10), is not practical for actual computing, but in the decade after Turing's theoretical work he and several others created the first electronic computers. John von Neumann developed the architecture that has become standard in contemporary computers and developed a means of encoding programs (representations) that specify operations to be performed on other representations.

3. Cognitive Science and the Cognitive Revolution in Psychology

In recognition of the generality of their application to domains other than the mind, we have introduced the concepts of communication, control, and computation without reference to how they were applied to the mind. But for many pioneers developing these concepts, understanding the mind was an important goal. It is important to emphasize that characterizing what goes on in the mind as comparable to what goes on in communication between people, employing processes like those that had been developed for the control of physical systems and computation in physical machines, was a major conceptual development. This was facilitated by the fact that information was not itself characterized in terms of the physical medium in which it was implemented. The recognition of the import of these ideas for philosophy of mind was mediated by developments in the empirical science of psychology and the developing interdisciplinary cluster that eventually became known as cognitive science.

Empirical inquiries into both brain and mind developed in the late 19th and early 20th centuries independently of the information-theoretical perspective. Pioneering research in the brain sciences involved the discovery that neurons are distinct cells that conduct electrical signals and deployed techniques such as lesion and microstimulation to link brain regions with behavioral activities. Psychology developed as a distinct discipline employing a variety of different approaches to study of the mind—act psychology (Brentano), functional psychology (James), structuralist psychology (Wundt, Titchner), cultural-historical psychology (Vygotsky, Luria), Gestalt psychology (Köhler), and genetic epistemology (Piaget). In part due to the lack of a clear framework for objectively characterizing activity in the mind, in North America these various fledgling approaches to psychology were overtaken by John Watson and other behaviorists who sought to make the new discipline appropriately empirical by bypassing the mind altogether and focusing on behavior. It might be said that, in rejecting the project of going inside the mind, behaviorism impeded the development of an information-theoretic understanding of mind. But William Aspray suggests that, in another respect, behaviorism actually fostered the development:

> Behaviorist psychology, by concentrating on behavior and not consciousness, helped to break down the distinction between the mental behavior of humans and the information processing of lower animals and machines. This step assisted the acceptance of a unified theory of information processors, whether in humans or machines. (1985: 128)

Physiological psychologists were among the first to draw upon information-theoretical ideas of computation and control. Neurophysiologist Warren McCulloch began a collaboration with Walter Pitts, a self-trained logician, that began by showing how idealized neurons could, when appropriately networked, implement all the connectives of sentential logic, and by extension could be viewed as carrying out any operation of which a Turing machine was capable

(McCulloch and Pitts 1943). In subsequent work they moved beyond viewing brain activity as implementing sentential logic, and focused on statistical properties of networks that would allow them to recognize related stimuli as the same. They also began to relate their networks to the architecture of specific brain regions, a project that was soon taken up by other researchers such as David Marr.

Psychophysics, which addressed questions such as the relation between the magnitude of a stimulus and the psychological experience of the stimulus, was another field of psychology that was not dominated by behaviorism in North America and it was in Stevens' psychophysics laboratory at Harvard that George Miller invoked information theory in his dissertation research on radio jamming. To present his classified research to wider audience, he shifted the focus to how noise affects the intelligibility of speech and asking why some messages are more robust to noise (a phenomenon Shannon [1948] himself went on to explicate in terms of the redundancy between parts of a message). Initially Miller tried to situate his work within the context of behaviorism using the label 'statistical behavioristics', but he quickly moved beyond the bounds of behaviorism and began to speak of cognition. His work remained influenced by information theory and one of his most cited papers (Miller 1956) was an investigation of the capacity limits affecting cognitive performance (e.g., limits to the information that can be stored in short-term memory). In a landmark book he, together with Karl Pribram and Eugene Galanter (1960), addressed purposive action and developed a computational conception of how agents could execute plans through a process of comparing a representation of the current state with that of a goal and implementing operations until no differences remained between the representations.

Miller was not alone in applying information-theoretic perspectives to understanding the mind. Donald Broadbent, for example, investigated how people are able to focus their attention on a particular conversation, but quickly shift to tracking a different conversation if they hear their own name. Broadbent invoked notions of information channels and filtering to explain such phenomena. By the late 1960s the range of research approaches applying information-theoretic ideas to understanding the mind reached a sufficient threshold that a new name was coined for this form of psychology—*cognitive psychology*. Ulric Neisser, in the book that gave this name to the new field, summarizes the importance of information to it:

> There were cognitive theorists long before the advent of the computer. Bartlett, whose influence on my own thinking will become obvious in later chapters, is a case in point. But, in the eyes of many psychologists, a theory which dealt with cognitive transformations, memory schemata, and the like was not about anything. One could understand theories that dealt with overt movements, or with physiology; one could even understand (and deplore) theories which dealt with the content of consciousness; but what kind of a thing is a schema? If memory consists of transformations, what is transformed? So long as cognitive psychology literally did not know what it was talking about, there was always a danger that it was talking about nothing at all. This is no longer a serious risk. Information is what is transformed, and the structured pattern of its transformations is what we want to understand. (1967: 8)

Appropriately, an alternative name for the field is *information-processing psychology*.

Psychology and neuroscience were not the only disciplines that began to employ information-theoretic ideas to explain mental phenomena in the 1950s. Linguistics, like psychology, had been dominated by behaviorism, but linguists such as Zellig Harris soon took up the challenge of characterizing the vast range of syntactic forms available in language that they confronted. To bring coherence to the range of forms, a framework was introduced in which the different syntactic forms were arrived at by applying formal operations to symbol strings. Harris, for example, hypothesized transformations of kernel sentences to account for passive and active versions of a sentence. This approach was further developed by Harris' student Noam Chomsky, who not only developed a number of transformational grammars in the attempt to account for the forms possible in actual languages, but also demonstrated that such grammars required computational power equivalent to a Turing machine.

During this period as well some pioneers in the new field of computer science began to explore whether an appropriately configured computer could be intelligent and established the subfield of artificial intelligence (see chapters 10 and 11). Especially noteworthy was Newell and Simon's (1956) program Logic Theorist, which constructed proofs, some of them novel, of theorems from *Principia Mathematica*. In subsequent decades these researchers focused on developing a general approach to problem solving and created a programming architecture known as a *production system* that was particularly useful for implementing programs designed to mimic human intelligence (Newell and Simon 1972). The representations used in production systems are modeled on propositions in language and the approach is generally referred to as *symbolic* AI. Other theorists found propositional representations to be too inflexible to account for cognitive abilities and developed larger-scale knowledge structures known as *schemas* or *scripts*, or broke altogether from symbolic approaches and pursued a more neutrally inspired architecture in which the fundamental activity is not application of rules but recognition of patterns through parallel constraint satisfaction.

Although the name *cognitive science* would not be introduced until the 1970s, these efforts in linguistics, psychology, artificial intelligence, and neuroscience provided the foundation for what became a highly productive interdisciplinary endeavor.[2] As the applications of information-theoretic approaches were being developed in these various disciplines, philosophers, especially Hillary Putnam and Jerry Fodor, drew upon them in articulating a new philosophical stance on the nature of mind (for an overview, see Bechtel 1988). In the early 20th century dualism and identity theory seemed to be the only live options for characterizing mind: mental states involved either a non-material substance or were identical to material states such as brain processes. To capture the ways the new cognitive disciplines invoked informational states and operations on them, Putnam and Fodor introduced what came to be known as *functionalism*, according to which mental states are states in an information-processing system that might be realized in different physical processes (and thus, are not identical with any of them). The theorists to whom we turn in the next section all worked within this functionalist perspective.

4. Information Enters Philosophy of Mind

As it was brought into philosophy of mind, the notion of information became intertwined with other concepts that already had a long history of use by philosophers. First, information is closely related to another concept that has enjoyed prominence in philosophical discussions since the Ancient Greeks: *knowledge*. Second, the notions of mental *contents*, mental *representations*, and *ideas* had been central to philosophy of mind since the 17th century. Third, the notion of information has become highly relevant to the problem of *intentionality*, which was given its modern formulation by 19th century philosopher Franz Brentano. Fourth, at the beginning of the 20th century Frege articulated his highly influential distinction between *sense and reference* that provided an account of meaning in non-mental terms. The task of integrating the notion of information arising from scientific fields with these preexisting philosophical notions has proved to be far from trivial.

We turn first to the problem of intentionality. Brentano had claimed that intentionality was the mark of the mental—mental states are *about* other states, including ones that do not exist. Because it might lack a *relatum*, Brentano characterized intentionality not as a true relation but as *relation-like.* Some philosophers turned to information theory for suggestions: if the mind processes information, then perhaps that could explain the content of mental states. However, as we noted above, Shannon's formal treatment of communication did not address *how* communicated messages carried information; he simply treated it as a basic fact that "frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities" (1948: 379). Shannon himself reserved the term *information* for the reduction of uncertainty that was crucial for his analysis of communication.

The two senses of *information*, however, were often conflated. The use of information-bearing states or representations in psychological and AI theories suggests that these accounts can explain the semantic sense of information. In a particularly influential paper, John Searle (1980) argued that such attempts confound the crucial distinction between intrinsic and derived intentionality: any meaning or aboutness that a message, state in a computer, signal running across a wire, etc. can have is dependent on the minded beings that give a message meaning or respond to it. Searle emphasized that the same applies to symbol strings in a computer—they are dependent for their meaning on the programmer. Trying to ground the aboutness of mental states in semantic information, goes Searle's argument, is putting the cart before the horse: the aboutness of information is itself grounded in those mental states. Not all philosophers have accepted Searle's contention that intentionality must be dependent on minds. The theories we discuss next can be looked at as attempts to respond to this line of argument by advancing accounts of the intrinsic aboutness of information in non-mental terms.[3]

4.1. Dretske

Perhaps because Shannon's specialized usage of the term 'information'—i.e., to refer to the amount of reduction of uncertainty about the occurrence of potential events that may be afforded by a signal—had become widespread by the 1950s, many philosophers who have been concerned with accounting for intentionality have tended to use the terms 'content' and 'representation' with much greater frequency. Dretske (1981), however, advanced an account, built on the foundation of Shannon's sense of reducing uncertainty, of semantic content. Important for understanding the relation between Dretske's theory and Shannon's is that instead of using 'information' to mean something distinct from content, as Shannon did, Dretske used ‘information’ to mean something distinct from representation. This was a distinction between kinds of vehicles: information carriers (he also used the term ‘indicators’) versus representations.

Instead of treating a cognitive agent as the sender of a message, he treated objects or events in the world as the sender and the cognitive agent as the receiver (Adams 2003: 472). Hence, he treated communication as a causal process but instead of thinking of the effect as a message that carries the same information as the cause, (drawing from Grice’s earlier notion of *natural meaning*) Dretske viewed the effect as carrying information *about* the cause. Being an epistemologist, the account was also well-suited to his goal of accounting for knowledge: on his view, knowledge was the result of the flow of information from the world into the knowing subject (Dretske 1981). This set the basic framework of Dretske’s account, but he recognized several ways in which causation by itself was insufficient to account for cognitive states with specific contents. First, causes carry information about more things than what is usually taken to be the semantic content of the effect—for example, a cause carries information about all the intermediate steps in a causal chain. Dretske proposes to explicate this in terms of the direction of causal dependencies: “S gives primary representation to property B (relative to property g) = S’s representation of something’s being g depends on the informational relationship between B and g, but not vice versa” (1981: 160).[4]

Moreover, messages should be able to be mistakenly tokened without affecting their semantic content. Dretske (1981) appeals to the process of acquisition to explain this—as the causal linkage between sensory stimulus and the message internal to the cognitive agent develops, the internal message can be generated without the sensory stimulus. He later abandoned this view, however, in favor of a teleological view of representations as those content vehicles whose *function* it is to carry information about what is represented (Dretske 1988; a version of this view was advocated earlier by Millikan 1984; see also chapter 22). Dretske does not believe that his appeal to function prevents his theory from achieving his goal of grounding aboutness “in an objective, mind- (and language-) independent notion of information” (2009: 381):

> What is important for the purposes of information-theoretic semantics is that there be a set of circumstances, or perhaps a kind of history, that, independent of human interests, grounds descriptions of animals and their parts as ill, sick, broken, damaged, injured, diseased, defective, flawed, infected, contaminated, or malfunctioning. If the truth of these descriptions is independent of our interests and purposes, then there is a way natural systems are supposed to be, or supposed to behave, that is independent of how we conceive them. (Dretske 2009: 387)

The most common criticism of Dretske’s theory poses what is known as the “disjunction problem.” In a case of misrepresentation, a content vehicle A, whose function it is to indicate B, instead indicates C (by being caused by state of affairs C instead of state of affairs B). If A-tokens occasionally carry information C, then in virtue of what is it true that the function of A is to indicate condition B, rather than the disjunctive condition B-or-C? Dretske’s theory must be able to give an answer to this that does not itself rely on human purposes and intentions in order to account for misrepresentation. A large literature has been devoted to this problem; see a recent discussion of it in Dretske (2009: 387–9).

4.2. MacKay

Other information-based approaches to semantic content began not with the mathematical theory of communication but with the cyberneticists' conception of negative feedback as enabling physical systems to be goal-directed by taking in information about the effects of their actions in determining new actions. Donald MacKay characterizes information as "that which alters" an agent's "total state of readiness for adaptive or goal-directed activity" (1969: 60).[5] This approach, as further developed in the work of Dennett, treats information not quite as mind-independent but as in fact dependent on the (implicit or explicit) goals of the animal:

> … since a stimulus, as a physical event, can have no intrinsic significance but only what accrues to it in virtue of the brain's discrimination, the problem-ridden picture of a stimulus being *recognized by* an animal, meaning something *to* the animal, prior to the animal's determining what to do about the stimulus, is a conceptual mistake. (Dennett 1969: 75–6)

MacKay argues that a message is only *significant* or *meaningful* insofar as receipt of the message would alter the *conditional readiness* of action. Behavior plays a key role in his account of intentionality, but he is clear about how his account is to be distinguished from behaviorism: "What has been affected by your understanding of [a] message is not necessarily what you do—as some behaviourists have suggested—but rather what you would be ready to do if given (relevant) circumstances arose" (1969: 22). The idea here is that mindedness, information, and goal-directedness are interdependent naturalistic phenomena. It is a physical fact about the animal that it has certain needs and certain capacities for acting to meet these needs, and is therefore motivated to act in certain ways under certain conditions. In virtue of such facts, stimuli become relevant to determining how to behave in virtue of information they carry about the effects of such behavior, and thereby acquire significance. For MacKay, this foundational significance ultimately underwrites the aboutness of information, and the aboutness of mental states, as a whole.

4.3. Sayre

In contrast to Dretske and MacKay, Kenneth Sayre proposes a process in which Shannon's quantitative sense of information is turned into semantic information. He uses the term 'info(t)' for the "technical" sense of information that Shannon's theory focused on, i.e., that of reduction of uncertainty. 'Info(s)', on the other hand, denotes a signal that bears not merely a statistical significance but a genuine semantic content. His project is to show how info(t) can be transformed into info(s) and Sayre focuses on the visual system to show how this is done. An organism may use info(t) in order to get around and engage in adaptive behavior on a limited basis, but much more adaptive behavior is possible with info(s), that is, when a state of the system has as its *content* the perceptual object.

Central to Sayre's approach is his characterization of organisms in which info(s) states arise. He characterizes organisms as adapting to changes in their environment, either through evolution or learning, but more fundamentally, needing to do this because ultimately they are *organized systems* that must resist dissipation into their environment. In thermodynamic terms introduced by Schrödinger, living organisms are low entropy systems relative to their environment, and to maintain themselves as such, they must extract energy from their environment. Sayre quotes Schrödinger in describing "a living organism as a device 'for sucking orderliness from its environment'" which Schrödinger characterized as negative entropy (1967: 79, quoted in Sayre 1986: 128). In order to perform this crucial activity, organisms need not only a way of incorporating energy from outside themselves but also identifying sources of energy. They do this by receiving info(t) from distal objects at their senses.

The sensory processing system then generates info(s) as it selectively reduces info(t), extracting and retaining only info(t) about a distal object O that is relevant to the actions the organism takes with respect to O. It is by "locking onto" the dynamically changing info(t) that is relevant to responding behaviorally to object O, and so *tracking* O, that the organism generates info(s) that is about O. Looked at in this way, perception is not just a passive process of aligning internal states of the organism with external stimuli, but a process that organisms use to *focus on* objects in their environment with which they need to interact appropriately if they are to continue to live. At some points Sayre speaks of a "perceptual-behavioral control loop," where the behavioral component initiates actions with respect to the object in the environment about which the perceptual system carried info(t). It is through such loops that the info(t) states become info(s) states—they become states the organism uses to engage with the external object. What underlies this transformation is that organisms are active systems that can alter processes such as perception so as to provide the appropriate info(t) for engaging in successful actions in the world. As we will see below, this approach foreshadowed to some extent the later approach of Neo-Gibsonians such as Chemero, who have embraced non-representational information-processing accounts of perception.

5. The Many Facets of Information in Philosophy of Mind

Informational theories of intentionality were vigorously debated through the 1990s but interest in the problem has waned since then. Although not couched in information-theoretic terms, recent philosophical discussions of intentionality and vehicles have significant implications for information-theoretic accounts. Accordingly, we discuss some more recent developments in philosophy of mind and their relevance to the topic of information.

The topic of the nature of intentionality/aboutness has itself grown to be very large and contentious, and has recently been the focus of in-depth investigation (e.g., Crane 2013; Yablo 2014). One important upshot is the recognition that there are a number of kinds of aboutness; this implies that for purposes of philosophy of mind, there is a corresponding sense of 'information content' for each of these kinds. For example, "A carries information about B" might mean that A refers to B (i.e., by indicating that B is present), or that B enters into the meaning of A (for example, A might be a story about B, a fictional or even a logically contradictory object). The former case invokes an extensional type of aboutness, the latter an intensional type. Information content can therefore be broken down into intensional and extensional content.

As an illustration of further ways to divide aboutness, consider Martin Davies' distinction between "subdoxastic aboutness" and other kinds of aboutness:

> Subdoxastic aboutness is distinct from attitude aboutness since, like experiential aboutness, it is a kind of non-conceptualised content. (Indeed, it is between those two kinds of non-conceptualised content that Evans (1982: 158) distinguishes in terms of serving 'as the input to a thinking, concept-applying, and reasoning system'.) Subdoxastic aboutness is distinct from linguistic aboutness since it is not derived. Subdoxastic aboutness is also distinct from indicator aboutness, since it allows for the possibility of misrepresentation. We can say, for example, that a state of the auditory processing system represents the presence of a sound coming from the left even though there is not in fact any sound coming from the left. And finally, subdoxastic aboutness is quite unlike experiential aboutness—the fourth variety of aboutness—since it is not tied to consciousness. (1995: 16)

One lesson to draw from this is that it may be an oversimplification to talk as if there is a single notion of "semantic information" (or info(s) in Sayre's terminology).

A second point that is emerging from recent work in philosophy of mind and cognitive science involves the notion of a content *vehicle*. Philosophers have pointed out that vehicles also come in more than one kind. Dennett (1983) for example proposed that information may be manifested non-representationally in simple storage formats—"brute storage" in his words—which do not allow any kind of information processing during storage (only retrieval). Note that the question of what types of vehicles may exist is orthogonal to the question of what types of contents might exist.

This issue has come to the fore in recent debates about perception. Earlier proponents of ecological approaches such as James Gibson had argued against information-processing accounts of processes such as visual perception by maintaining that information needed for action (what he termed *affordances*) already existed in the light. Organisms merely need to pick up this information; they do not need to process it. What marks the approach of Neo-Gibsonians such as Chemero (2009) is that it embraces Gibson's non-representational direct realism, while acknowledging the importance of information processing. For Chemero there is no contradiction here because information processing in direct perception does not require representations in the mind, a point he also expresses by saying it does not require "computation" or "mental gymnastics":

> Action changes the information available to an animal's perceptual systems, and sometimes the action actually generates information. Thus there is a sense in which perception-action as studied by radical embodied cognitive scientists involves information processing, but it is a variety of information processing that does not involve mental gymnastics. (2009: 127)

Discussions about information often assume that all informational content is declarative. Another contribution of recent theorizing questions this. Not all messages communicated between people are assertions: they can also be imperative or interrogative, for example. Belnap (1990) has argued that the semantics of imperative and interrogative sentences cannot be analyzed in terms of assertoric content, and that truth or falsity would not be applicable to such sentences. Further, Belnap along with a number of other philosophers has proposed many-valued logics. These developments, as well as parallel advances in analog and non-binary digital computation, point to the need for richer accounts of semantic information.

Finally, for a piece of information to exist, a content token must become, as it were, *attached* to a vehicle token; they become informationally bound to one another. What are the necessary and sufficient conditions for this attachment to occur? Multiple answers are possible since there are many *schemes* under which such attachment might occur. A might mean B merely by arbitrary stipulation (as in games or stories), it might mean B by means of being linked causally to B in the right way (as in Dretske's indication relation), or it might mean B by being historically linked in the right way (as in Millikan's biosemantics). There is the possibility of natively endowed content, or (as a Kantian might be inclined to say) content imposed by a constitutive faculty of the mind. Each of these might be referred to as a distinct *content tokening scheme*: an ordering principle that determines which contents get tokened or "attached" to which vehicles and under what conditions.[6] The issue of tokening schemes is importantly different from understanding the nature of content itself, a point that has also been made by Rick Grush: "a theory of content, by contrast, need not concern itself with how or why contents are carried by this or that vehicle—rather, it is concerned with what contents are" (1998). Providing a reductionistic account of what it is for A to be *about* B in one of Davies' senses, for example, would not be the same thing as providing a reductionistic account of the potential schemes for tokening such contents.

Final Reflections

This brief review reveals that philosophers of mind have adopted multiple perspectives on information. To make progress, future philosophical discussions of information need to acknowledge the multi-faceted nature of information and differentiate at least five dimensions (Table 29.1). If an account treats a physically real thing as a piece of information, it must address at least three questions: What is the intensional content? What is the vehicle? And what is the content tokening scheme? Depending on the type of information, it may need to address two further questions: What is its truth value? And what does it refer to in the world (i.e., what is its extensional content)? As our brief review of philosophical approaches makes clear, not only are there multiple answers to each question, there are multiple *categories* of answers for each question. Continued cross-pollination between philosophy of mind and information-related fields will require both to acknowledge the interrelatedness and multi-dimensionality of information-related notions.<TABLE 29.1 HERE>

References

Adams, F. (2003) "The Informational Turn in Philosophy," *Mind and Machines* 13: 471–501.

Aspray, W. (1985) "The Scientific Conceptualization of Information: A Survey," *Annals of the History of Computing* 7, no. 2: 117–40.

Bechtel, W. (1988) *Philosophy of Mind: An Overview for Cognitive Science*, Hillsdale, NJ: Erlbaum.

Bechtel, W., Graham, G., and Abrahamsen, A. (1998) "The Life of Cognitive Science," in W. Bechtel and G. Graham (eds), *A Companion to Cognitive Science*, Oxford: Basil Blackwell, pp. 2–102.

Belnap, N. (1990) "Declaratives Are Not Enough," *Philosophical Studies* 59, no. 1: 1–30.

Boden, M. (2008) “Information, Computation, and Cognitive Science,” in P. Adriaans and J. van Benthem (eds), *Handbook of the Philosophy of Science. Volume 8: Philosophy of Information*, Amsterdam: North Holland, pp. 749–69.
Chemero, A. (2009) *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press.
Crane, T. (2013) *The Objects of Thought*, Oxford: Oxford University Press.
Davies, M. (1995) “Consciousness and the Varieties of Aboutness,” in C. Macdonald and G. Macdonald (eds), *Philosophy of Psychology: Debates on Psychological Explanation*, Oxford: Blackwell, pp. 356–92.
Dennett, D. (1969) *Content and Consciousness*, London: Routledge.
Dennett, D. (1983) “Styles of Mental Representation,” *Proceedings of the Aristotelian Society* 83: 213–26.
Dretske, F. (1981) *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.
Dretske, F. (1988) *Explaining Behavior: Reasons in a World of Causes*, Cambridge, MA: MIT Press.
Dretske, F. (2003) “Experience as Representation,” *Philosophical Issues* 13, no. 1: 67–82.
Dretske, F. (2009) “Information-Theoretic Semantics,” in B. P. McLaughlin, A. Beckermann, and S. Walter (eds), *Oxford Handbook of Philosophy of Mind*, Oxford: Oxford University Press, pp. 381–92.
Floridi, L. (2011). *The Philosophy of Information*, Oxford: Oxford University Press.
Grush, R. (1998) “Skill and Spatial Content,” *Electronic Journal of Analytic Philosophy* 6, no. 6.
Grush, R. (2002) “The Semantic Challenge to Computational Neuroscience,” in P. K. Machamer, R. Grush, and P. McLaughlin (eds), *Theory and Method in the Neurosciences*, Pittsburgh: University of Pittsburgh Press, pp. 155–72.
MacKay, D. M. (1969) *Information, Mechanism, and Meaning*, Cambridge, MA: MIT Press.
McCulloch, W. S., and Pitts, W. (1943) “A Logical Calculus of the Ideas Immanent in Nervous Activity,” *Bulletin of Mathematical Biophysics* 5: 115–33.
Miller, G. A. (1956) “The Magical Number Seven, Plus or Minus Two: Some Limits On Our Capacity for Processing Information,” *Psychological Review* 63: 81–97.
Miller, G. A., Galanter, E., and Pribram, K. (1960) *Plans and the Structure of Behavior*, New York: Holt.
Millikan, R. (1984) *Language, Thought, and Other Biological Categories*, Cambridge, MA: MIT Press.
Neisser, U. (1967) *Cognitive Psychology*, Englewood Cliffs, NJ: Prentice-Hall.
Newell, A., and Simon, H. A. (1956) “The Logic Theory Machine: A Complex Information Processing System,” *IRE Transactions on Information Theory* 2, no. 3: 61–79.
Newell, A., and Simon, H. A. (1972) *Human Problem Solving*, Englewood Cliffs, NJ: Prentice Hall.
Pribram, K. (1986) “The Cognitive Revolution and Mind/Brain Issues,” *American Psychologist* 41, no. 5: 507–20.
Sayre, K. (1986) “Intentionality and Information Processing: An Alternative Model for Cognitive Science,” *Behavioral and Brain Sciences* 9: 121–60.
Schrödinger, E. (1967) *What is Life?* Cambridge: Cambridge University Press.
Searle, J. R. (1980) “Minds, brains, and programs,” *Behavioral and Brain Sciences* 3: 417–57.
Shannon, C. E. (1948) “A Mathematical Theory of Communication,” *The Bell System Technical Journal* 27: 379–423, 623–56.

Turing, A. (1936) "On Computable Numbers, With an Application to the *Entscheidungsproblem*," *Proceedings of the London Mathematical Society* Series 2, 42: 230–65.

Weaver, W. (1949) "Recent Contributions to the Mathematical Theory of Communication," In C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*, Urbana: University of Illinois Press, pp. 1–28.

Wiener, N. (1948) *Cybernetics: Or Control and Communication in the Animal and the Machine*, Cambridge, MA: MIT Press.

Yablo, S. (2014) *Aboutness*, Princeton, NJ: Princeton University Press.

---

[1] See Aspray (1985) for more details on these conceptual developments.

[2] For a historical account of the development of cognitive science, see Bechtel, Abrahamsen, and Graham (1998). Boden (2008) examines the various roles the notion of information has played in cognitive science.

[3] The "symbol grounding problem," discussed in chapter 11, is a species of the more general problem of explaining the intentionality of informational states in non-mental terms. The "symbol grounding problem," in particular, is focused on *symbolic* states such as those found in a computer, and those invoked in, e.g., Fodor's "computational theory of mind."

[4] Adams (2003: 491) notes the similarity of Dretske's explanation to Fodor's account of asymmetric dependence between a cognitive structure and what it means, but in Fodor's account, the key connection is secured by the existence of a law connecting the referent with the message.

[5] Here MacKay's view bears a resemblance to Floridi's (2011, p. 164) "action-based semantics".

[6] Grush (2002) uses the term 'content assignation scheme'.