

Teleosemantics and Natural Information

John Symons

Forthcoming in Floridi, L. (ed.) *Routledge Companion to Philosophy of Information*

1. Introduction

We human beings tend, rightly or wrongly, to explain the intelligent behavior we observe among our fellow organisms by reference to their mental states; their thoughts, feelings, and the like. Our understanding of the bear's desire for honey, for example, allows us to loosely predict what he would do in its presence and to explain to one another why he investigates holes in trees.

Generally speaking, when we think about mental states like belief and desire we assume that they are *directed towards*, or that they are *about* things. Philosophers call this *aboutness* relation 'intentionality'. Alexius Meinong (1899) and others noted that those things towards which beliefs and desires are directed are not always physically present to the organism and in some cases might not even exist at all. My desire to own a regular, seven-sided polyhedron, for example, has no object because the kind of polyhedron that I covet is mathematically impossible. Here, mental states have a puzzling relationship to their objects. In the case of the absent honey, for example, many philosophers would argue that the bear's desire is at least partly composed of or directed towards something like a mental representation of the honey. To use philosophers' jargon the bear's representation of honey forms the content of this particular desire. The content of my desire for an impossible object is a little harder to pin down.

The enterprise of explaining intelligent behavior is generally thought to call on a range of philosophically problematic notions. Most philosophers will assume that the concepts of intentionality, representation, and content force themselves on us as soon as we reflect on the nature of intelligent behavior.

In addition to the intentionality of mental states, philosophers have also been concerned with the ways in which representations track aspects of the organism's environment. When we see intelligent behavior we often have good reasons to believe that the organism's mental representations are related to aspects of its environment in a regular manner. Bears want honey because it is tasty and its tastiness is the body's way of indicating that it is a good source of carbohydrate for hungry bears. Similarly, my belief that Motan is on the mat tracks the state of affairs in front of me at this instant and allows me to navigate my room in ways that are good both for me and for my cat. Reflecting on this tracking relationship, some philosophers (eg. Daniel Dennett 1987, 75) take it as obvious that natural selection would favor organisms with true representations of the world. Building on this apparently straightforward observation, other philosophers, (like Nicholas Shea

2007, 404) argue that any adequate theory of mind must account for the truth of representations underlying successful behavior.¹

For many philosophers and scientists the best way to make sense of the tracking relationship between mental representations of objects and the objects themselves (for example, the bear's representation of honey and honey itself) is in terms of information. The idea is that representations somehow carry information about the objects in the world in ways that make them useful for organisms. On this account, an organism that exhibits intelligent behavior does so via its sensitivity to the information about its environment that is contained in its representations. Sensitivity, along with the ability to manipulate useful information, allows an organism to adapt its behavior in intelligent ways and to thrive thereby.²

Fred Dretske's *Knowledge and the Flow of Information* (1981) provided the most influential early theory of how this might work. Dretske's informational semantics was inspired by the mathematical theory of information and builds on the idea that representations are best understood as causally related in some way to their objects.³ For other scientists and philosophers, non-informational accounts of signaling, communication, and perception can suffice to explain intelligent behavior. The debate over whether to understand representation and intelligent behavior in informational terms or in some other way is lively and ongoing.⁴ This chapter explores part of this debate, specifically, the relationship between teleosemantic theories of mental content and the concept of information.

Teleosemantics is a prominent and influential theory that was largely developed in the 1980s and 90s by Ruth Millikan and David Papineau.⁵ Very roughly, teleosemantics explains the meaningfulness of representation in terms of biological functions. The role of information in biologically oriented theories like teleosemantics is a topic of great controversy.⁶ Reading the literature on these questions is complicated by diversity of jargon and by differing uses of the same terms. However, a much deeper obstacle is the presence of competing conceptions of explanation in the philosophical literature. Philosophers differ sharply with respect to what counts as a good explanation of notions like normativity, meaning, and representation. The initial challenge for readers is to determine what the various interlocutors take to be the demands, standards, and commitments that we ought to adopt with respect to explanations in the philosophy and science of mind.

¹ Alvin Plantinga has taken the opposite position with respect to the relationship between the truth of beliefs and their evolutionary history arguing that if evolutionary explanations are correct then the function of evolved capacities (like the capacity to believe) is survival and not the production of true beliefs. (1993, 218)

² See Adams 2003 for a discussion of what he calls the informational turn in philosophy of mind. He helpfully describes the difference between non-naturalistic theories of information in Bar-Hillel and Carnap (1952) from naturalistic accounts of information as applied to philosophy of mind in the 1980s.

³ Causal theories of representation in their modern form were first defended in Stampe. (1977) The appeal of causal theories is straightforward. Since at least some mental representations should track events and objects outside the mind then presumably there should exist some causal connection representations and non-mental states of affairs. For an introductory discussion of causal theories of representation see Adams and Aizawa. (2010)

⁴ See Stegmann (2013, 1-35) for a comprehensive introduction to the role of information in animal communication and signaling.

⁵ While David Papineau is also responsible for a prominent version of teleosemantics (See for example his 1984; 1987; 1997), in this chapter I concentrate on Millikan's version of the view for the purpose of examining the debate about the role of information in any theory of content. Criticisms of Millikan's view and her responses to those criticisms provide a better venue for thinking about information in this context. Papineau's views differ from Millikan's in ways that are described by Karen Neander in her excellent Stanford Encyclopedia of Philosophy article *Teleological Theories of Mental Content* (2012).

⁶ See Neander (2012) for a complete overview of the controversies surrounding teleological theories.

In addition to determining the kinds of explanations that philosophers are after, the second challenge involves comparing and evaluating these differing views of what is required for an explanation to count as a good explanation. Part of the purpose of this chapter is to help readers determine for themselves which style of explanation most closely matches their own.

The role of information in these debates is generally determined by the explanatory demands that the philosopher assumes are in play. If, for example, like Fodor and Pylyshyn, we regard properties like systematicity as crucial for any acceptable theory of mind we will have a different view of the role information plays in explanation than if, for example, we see mind in behaviorist, or dynamicist terms.⁷ Similarly, if, for example, like Shea or Dennett we see true representations as playing a central role in intelligent behavior, we will have a very different view of the role of information than might be the case for philosophers like Stephen Stich (1983) or Christopher Stephens (2001). Stich and Stephens have defended ways of thinking about the mind that do not assume a direct relationship between intelligent behavior and true belief.

While this chapter focuses on teleosemantics and its critics, this debate is an instance of a broader set of questions concerning competing conceptions of explanations. The chapter describes the landscape of problems that teleosemantics was intended to solve before going on to present Ruth Millikan's central arguments. From there, criticisms of her position are described and some responses are sketched. Nicholas Shea's criticisms of Millikan takes center stage. He argues for the importance of information (an input condition) in any theory of representation and contends that Millikan's emphasis on the role of selection in the appearance of the agent's capacities (the output conditions) leaves her theory vulnerable to objections. I will examine the success of Millikan's response to critics below and will describe her view of information as it appears in her more recent work. I argue that Millikan's emphasis on the properties of the consumer of representations can be reconciled in a relatively straightforward manner with the importance of information so as to preserve the theoretical virtues of her approach while acknowledging the intuitive appeal of information-based theories.

2. The purpose of teleosemantics

The tendency to explain behavior in terms of conscious mental states is difficult for most of us to resist. Even in cases where we are aware that no mental state is in play, or where we understand that mental states have no relevant causal relationships to the behaviors or phenomena of interest, we find ourselves ascribing thought and feeling to the organisms and other intelligent systems we encounter. It is our deeply engrained habit to engage not only with the natural world but with artifacts and with the worlds of fiction as though they are filled with meaning, purpose, and thought.

⁷ See Symons (2001) for my account of differing views of explanation in the study of representation. See also Symons and Calvo (2014) for our discussion of the role of systematicity arguments in these debates.

Daniel Dennett offered an explanation of how we got into this situation. Specifically, Dennett explains how this interpretive strategy -- what he calls the *intentional stance* -- emerged in our species over the course of natural history.⁸ Dennett's approach is intended to show how we became natural mindreaders. However, even given a naturalistic explanation of the *ability* to read the world as filled with minds, questions remain. We can still ask, for example, how the success of the intentional stance can be reconciled with the widely held view of the physical world as devoid of purposes, meanings, and the like? Why should adopting the intentional stance provide an advantage to a species in a world without meaningful mental states? More deeply, we can also ask what it means for our thoughts to represent reality correctly or incorrectly. What does it mean to apply the normative judgment of rightness or wrongness to the representational aspects of mental life? When we say that an animal or a person made a mistake, that they have a false belief, or are in error in some other way, how do we reconcile our normative claim with the non-normative facts of science?

Millikan's theory specifically addresses the problem of normativity. The view was articulated in detail in her book *Language, Thought, and Other Biological Categories* (1984) and elaborated upon in her (1993) and (2004). While teleosemantics appears in the context of traditional philosophical debates concerning mind and language, and must be understood in those terms, it approaches the question of mental representation from a distinctively biological perspective.

The biological orientation of Millikan's methodology owes much to Dennett's philosophy of mind. However, there are significant differences. As with Dennett, the action of natural selection on agents and systems plays a central explanatory role in her view. Unlike Dennett, she rejects the view that we use something like the intentional stance to *predict* the behaviors of others arguing instead that standard folk psychology provides explanations of intelligent behavior in retrospect. (2004 21-24) Another difference between Dennett and Millikan is her focus on explaining the normative aspect of mental representation.⁹ As discussed below, the possibility that mental representations can misrepresent the world is essential, she argues, to the nature of representation. Thus, at its heart, the goal of teleosemantics is to provide an explanation of how a system's use of representations can be correct or incorrect.

From a certain perspective, it appears fundamentally misguided to say that a simple organism or even a non-cognitive system can misuse a representation, or that it can objectively fail to achieve a purpose. The critic's concern here is simple - we might project rightness or wrongness onto the use of representations in non-human contexts, but to claim that there is some non-human dependent sense, in

⁸ See Dennett (1989). See also Symons *On Dennett* (2002) for an explanation of the relationship of the intentional stance to traditional questions in philosophy of mind.

⁹ Dennett largely follows Millikan in her account of the normative aspects of representation. In his criticisms of the notion of intrinsic intentionality he takes Millikan to have correctly addressed the issue via an appeal to natural selection. (Dennett 1988, 200)

which an animal can be wrong about its representations is confusedly anthropomorphic.

Millikan sees the normative aspect of representation as natural and objective and she aims to explain it in terms of the biological notion of function. Her theory assumes that the emergence of mental representation is a process that takes place over the course of natural history and that it can be explained in scientifically respectable terms. According to its proponents, teleosemantics explains the normative aspect of mental states biologically, in the same way an evolutionary biologist would explain the function of the heart or the shape of a finch's beak.

On her view, representational systems (like most other biological systems) should be explained as products of natural selection. The function that these systems were selected for; their so-called *proper function*, is to allow the animal to relate representationally to particular aspects of the environment. An animal's representation of berries, for example, is part of (or the product of) a system that allows the animal to detect food in its environment. This system exists insofar as it provided an evolutionary advantage to that animal's ancestors. She argues that this advantage serves to explain why the representational system does what it does. Essentially, the core of teleosemantics is the simple contention that an animal's representation of, say, berries, is meaningfully related to berries because of the evolutionary history that gave rise to the process or mechanism in the animal's brain that produces that representation.

In very broad terms, Millikan's view is that a representation of a berry can serve a representational role in virtue of being produced by a mechanism that has the right kind of proper function. In general, the function of the mechanism that produces the representation is the detection of specific aspect of the environment for the agent. A great deal rests on the idea of proper functions in her theory and I will discuss this idea in detail below. *Going wrong*, in this context, means *failing to perform the function*. To use her example, a representation can fail to perform its function while still being a representation in the same way that a can opener can fail to perform its function without ceasing to be a can opener. Knowing that a can opener is failing requires knowledge of the purpose for which it was made. Similarly, determining rightness and wrongness in the context of representations requires knowledge of their purposes.

As we shall see, Millikan's theory of representation is non-informational in the sense of not being what had counted as a standard informational theory of mind prior to her work. Specifically, her theory differs in important ways from Fred Dretske's informational theory of mind as articulated in his (1981) and (1988). In particular, when readers call her view non-informational, what they mean is that the theory is built around the selective pressures governing the users or consumers of representations rather than on the transmission and reception of information from the environment to the organism or system. An explanation that starts from the characteristics of users, rather than from the flow of information has some strategic and theoretical advantages, as we shall see, but it also makes her version of teleosemantics vulnerable to non-trivial criticisms.

3. Meaning and representation

Even a very minimal level of meaningfulness; the semantic property of thoughts, for example, poses a conceptual challenge. When I think of San Francisco, my thought is meaningful insofar as it is about San Francisco. But what is it for my thought of San Francisco to be *about* a city? How does a physical structure or process in a brain relate to a city in a way that somehow carries meaning?

Conceptual problems concerning meaning and reference occupied analytic philosophers of language for much of the Twentieth Century and this tradition has had considerable influence on the neighboring sub-disciplines of philosophy of mind and philosophy of psychology. Philosophers of language created theories of reference in order to provide a set of necessary and sufficient conditions for characterizing the referential relationship between terms in a language and objects in the world. Attempts to provide a naturalistic account of mental representations were shaped by the powerful influence of philosophy of language.

Continuing the tradition of philosophy of mind and language, Millikan keeps the focus on the strangeness of the representational relation. However, unlike traditional philosophers she crafted biological explanations that target the most conceptually problematic aspect of this strangeness, namely the normativity of representation. This section describes (in simplified terms) the traditional answers to the problem of representation so as to understand the purpose and character of Millikan's view.

3.1 Conditions for folk psychological explanations

Let's begin with our commonsense view of psychological explanation before considering the normative aspects of representation. Any of us, reading a novel or watching a cartoon will slip into interpreting the characters as intentional agents with inner lives that explain their actions: Roadrunner saw that Wiley Coyote was approaching, believed that he was dangerous, and wanted to escape. Referring to Roadrunner's beliefs, his desires, what he sees, or what he fears in order to explain behavior exemplifies what Dennett called "folk psychology". Folk psychological explanation is generally reliable (by some standards), but it is also thoroughly anthropomorphic (by any standard).

Folk psychological explanatory strategies depend on our confidence that the patterns we mention in our explanations are real enough to help us predict the behavior of reasonably intelligent organisms. In addition to assuming that complex organisms have the kind of inner system of beliefs, desires, and actions that we find in ourselves, we also assume that the organism must, in some sense, get the world right in order to thrive. In this respect, folk psychological explanations incorporate normative judgments essentially.

As we saw above, when we consider an animal getting the world right and thriving as a result, it is natural to think in informational terms. Intuitively, we

assume that organisms are sensitive to information in their environments and that they manipulate and respond to environmental information for their purposes. This assumption that intelligent systems can be right or wrong about aspects of their environment is at the heart of folk psychological practice.

These two facets of our folk psychological theorizing, namely the belief-desire-action model of the organism's inner life on the one hand and the idea that it is sensitive to environmental information on the other, are not as tightly coupled as one might assume. One very important feature of Millikan's view, for example is that the consumers of representation need not be cognitive systems in the traditional sense. So how does one distinguish between the representational or informational features of an organism and its cognitive capacities?

One could, for example, imagine an organism whose behavior is tightly coupled to environmental information in a reliable and adaptive manner without that organism having a complex cognitive mechanism involving the interplay of beliefs and desires. A simple system might make use of the position of a light source or the distribution of some chemical in the environment in a way that is adaptive without having beliefs about the light or desires to approach the light.¹⁰

On the other hand, one can also imagine agents with complex cognitive inner lives being systematically mistaken with respect to environmental information. It is likely, for example, that much of our own interplay of adult human beliefs, desires, and actions is epiphenomenal with respect to the actual mechanisms guiding our behavior. We think our actions are motivated by beliefs and desires of a particular kind when in fact they are sometimes caused by mechanisms to which we do not have conscious access. If this is the case, even occasionally, then complex cognition and environmental information could also be uncoupled along these lines.

The possibility that these two aspects of folk psychological explanation can be uncoupled will be important when we assess criticisms of teleosemantics below.¹¹

3.2 Explaining representation

A natural starting point for any effort to explain representation in naturalistic terms is to assume that there is an isomorphism between the representation and what is being represented. A representation of the coyote is about the coyote, one might think, in virtue of sharing the relevant kind of structure or form with it. On reflection, we soon recognize that the fact of simply sharing some features in common with some object is not sufficient for something to be a representation of that object. After all, we could discover countless isomorphisms between some brain process and objects in the world. The representation of the coyote may be isomorphic with the coyote, but it is also, from another perspective, isomorphic with a dog, a horse, a cat, etc. Isomorphism alone is not enough for a satisfying theory of representation.

¹⁰ Think, for example of Braitenberg's vehicles.

¹¹ In fact, there are a variety of different ways that this decoupling can be articulated that have bearing on the relationship between content and information. Luciano Floridi has recently argued for a constructionist account of how meaningful representation results from repurposing what he calls "natural data" generated by the senses. (forthcoming)

Another way we might link representation and the thing represented is via a causal connection between the two. In that way, the intentional content of a representation would be explicable in virtue of a causal chain connecting the thing represented and the representation; the representation of the coyote is about the coyote because it was caused by the coyote. It is natural, from this perspective, to think that the causal connection between the representation and the represented object can somehow provide a respectable explanation of the intentional content of the representation. Unfortunately, causal relations alone are not sufficient for characterizing the representational relationship. The reason is as follows: many things contribute to causing a representation that have nothing to do with what it is that the representation is about. Roadrunner's being delayed by a train causes him to see Coyote's trap. But the content of the visual representation is the trap, not the myriad causal factors, including the train, his parent's romance, the presence of oxygen in his immediate environment, etc., that contributed to Roadrunner having the representation he does.

So, the causal relationship has to be elaborated somewhat in order for us to have an explanation of why it is that Roadrunner's representation of the trap is a representation of the trap. The usual way of elaborating this relationship is in terms of the notion of information. Specifically, the idea is that the representation carries information about its cause. The most influential early proponent of this way of thinking about representations was Fred Dretske. He argued that representation is a causal-informational relation.

Information, on Dretske's view, is "an objective commodity" (1981 vii) prior to, and in some sense serving as the condition for intelligent behavior. Information is not (as is often argued) the result of an interpretive process of some sort. Treating information as an objective feature of the natural world was an improvement over thinking of representation in terms of isomorphism or causal relations alone. However, other philosophers soon pointed out that Dretske's account is still insufficient. The reason is simple: many things that are not representations are casually related to other things in ways that can be understood to carry information. For an experienced hunter, battered vegetation, or a trough in the soil on the forest floor, carry information about his prey. But a causal-informational connection alone does not make marks on the forest floor representations of the prey. So, what, in addition to causal informational connection makes a representation a representation and not merely an indicator of some other connected feature of the natural world.

At this point, in the effort to understand the difference between indicators and representations, we see the appearance of teleological theories of content. As mentioned above, Millikan and others recognized that the challenging feature of this problem involves normativity. Of course, in one sense this recognition was not novel. Understanding how representations can fail to represent (or can represent non-existent objects) has long posed a difficulty. However, teleosemantics focuses sharply on cases of error, arguing that philosophical theories should have the resources to explain the occasions where Roadrunner is mistaken about his representations. Roadrunner swerves to avoid the black circle on the road ahead thinking that it is a hole when it is in fact a piece of paper. This is the kind of

misrepresentation that teleosemantic theorists take to be the most puzzling and philosophically distinctive feature of representation. They hope to account for the normative features of intentional content while simultaneously providing a naturalistic explanation of that content.

4. Proper Functions

How is it that we can categorize a physical process in an animal's nervous system using normative notions like truth or falsity? As we saw above, teleosemantic explanations make reference to the evolutionary history of organisms that use mental representations. They explain content in terms of biological purposes. For Millikan, the relevant purposes are the proper functions of the organism's representational systems. The proper function of a system is simply what it ought to do by virtue of what it is – its *raison d'être*. Any one system can serve a variety of functions, not all of these functions will count as the system's proper function. So, for example, one of the functions of the heart is to serve as a symbol of love in some cultures. This would *not* count as a proper function of the heart, neither would the sound it makes, its weight, what it tastes like in a stew, etc. Instead, we would ordinarily say that the proper function of the heart is to help circulate blood in the organism's body. Other functions are accidental by-products of the function for which it was selected.

Millikan's strategy then is straightforward: first, she recasts intentionality in terms of purposes. Then, she provides a biological account of purposes. Not all purposes will be relevant. On Millikan's view, representational relations are established via the functional properties of some mechanism in the animal. Those functions are selected for via learning or via the evolutionary history of the species. What specifically gets selected are the functions of the representation producing mechanisms in the animal. The selected function of a representation of a berry is to represent berries for the user of the representation. In cases where the user mistakes something else for a berry, the representation is not functioning correctly. Her argument for the idea of false representation as failed representation works by analogy. There are other perfectly natural ways in which we talk about failure of function that have a similar conceptual flavor, most notably in cases involving artifacts. The following is an example of how her argument by analogy works:

False representations are representations, yet they fail to represent. How can that be? It can be in the same way that something can be a can opener but be too dull and hence fail to open cans [...] They are "representations" in the sense that the biological functions of the cognitive systems that made them was to make them represent things. Falsehood is thus explained by the fact that purposes often go unfulfilled. (2004 64-5)

The analogy with failure in artifacts is unconvincing without further elaboration.¹² The first challenge is to make sense of the idea of purposes playing any role in an ostensibly naturalistic theory? In her book, *Varieties of Meaning*, Millikan acknowledges that talk of biological or sub-personal purposes is likely to be interpreted as a kind of metaphorical extension of ordinary person-level purposes and meanings. The success of her project requires that she respond to this criticism.

Her view is that sub personal or biological purposes are not fundamentally different in kind from the intentions or purposes that we find in the personal level. Instead, she argues that the purposes of the whole adult human person emerge from sub personal level processes and mechanisms. Thus, she contends, when we say that the eye is “*meant* to close automatically” in the presence of some foreign object that comes too close, this sense of *meant to* is not fundamentally different from what we mean when we say “I didn’t *mean to* blink”. (2004, 3)

She can argue for the continuity between these two examples by first arguing that the lack of awareness of a purpose does not entail that some action is not purposive. This is a plausible contention. We can, after all, imagine cultivating a pattern of behavior for some purpose such that it becomes habitual, or perhaps we come to forget the original purpose. Alternatively, we could imagine tricking someone into performing some action for some purpose through behaviorist-style reinforcement along the lines of Millikan’s examples in her (2004, 10). The fact that some purpose is hidden from introspective view does not make it any less of a purpose. She presents reasons that are intended to undermine the assumption that there is a principled distinction between so-called *real* purposes (associated with conscious agency) and the kinds of purposes that we might ascribe to parts of an agent or to some other non-cognitive system.

Blurring the distinction between purposes at the level of the whole person and purposes of parts of the person is only one component of her argument. For teleosemantics to serve as a naturalistic theory of the normative aspects of representation, it must provide a naturalistic theory of purposes. For philosophers like Millikan and Dennett such an account is already available and widely accepted, namely the theory of natural selection. The challenges of circularity, regress, and anthropomorphism are solved simultaneously by recourse to the biological explanation of the origin of purposive systems.

Skepticism concerning naturalization derives from the concern that explanations of purpose will beg the question; assuming the higher-level purposes that they seek to explain in the explanans by ascribing purposes to pre-purposive nature. To avoid circularity, question-begging, and regress, the Darwinian must demonstrate that natural selection is the source of purposes without itself being purposive.

Arguing that proper functions are the source of the normative features of representation and are the basis upon which higher-level features of human cognition can emerge faces a basic challenge. The worry is that we interpreters

¹² Artifacts and biological systems will have different individuating characteristics by virtue of the kinds of functions that are proper to them. For an alternative approach to the individuation of biological entities in virtue of functions see Symons (2010)

ascribe functions to biological systems and, moreover, that an indeterminate variety of interpretations are compatible with the facts. As Mark Rowlands explains (1997) grounding content in terms of function is open to criticism insofar as any particular biological function is not, by itself, sufficient to fix a single interpretation of the semantic content it is understood to ground.

According to advocates of teleosemantics, when the coyote detects the roadrunner in its environment, his properly functioning neural system is in a state that can be understood as being associated with the content 'roadrunner'. This is because, on Millikan's view, the proper function of this system is to detect roadrunners. However, as Fodor (1996) and others have noted, there will always be alternative interpretations of the function of this system that are consistent with what we observe. The content of the system might be the pattern of the roadrunner's tailfeathers, roadrunner-like movement, food here now, etc. Millikan is not especially concerned with the logical possibility that there are interpretations of any representation's content that are coextensive with the one that we assume is in play. The fact that the 'roadrunner' is virtually always co-present with 'roadrunner-like movement' is not a problem for teleosemantics insofar as it assumes that the content should be grounded in the causal history of the representation. 'Roadrunners' played a causally relevant role in the appearance of the representation on this view while 'roadrunner-like movements' were caused by roadrunners and are, at best, causally secondary or irrelevant to the selection pressure that gave rise to the representation.

Neander (2012) points out another aspect of the indeterminacy problem when she notes that contributions to fitness that support a proper function can include a multitude of factors that we would ordinarily regard as irrelevant to the content of the representation. It is difficult for Millikan's account of content to screen off such contributions in a non-question begging manner. So, for example, Coyote's representation could consistently be interpreted as 'roadrunner in the presence of oxygen'. Following Pietroski (1992) Neander points out that teleosemantic theories are in trouble if they force us to say that an animal's representation is something that the animal could know nothing about. Coyote is in no position to have views about oxygen, but oxygen is certainly some part of the function of his representational system.

Pietroski objects to teleosemantics insofar as it requires us to deny the existence of the very thing that we had hoped to explain. Whereas we assumed that the animal's representation had some content, and that this was the target of our explanation, in fact we were wrong, the real content is the set of conditions that led to selection of the function of that representation. If, for example, the content of an animal's representation is fixed by some disposition that increased fitness, this disposition might have only an accidental relationship to what we would ordinarily say the representation is about. This means that teleosemantics is at least partially a revisionary account of mental life.

The next section examines the manner in which the indeterminacy challenge relate to the problem of determining what constitutes a good explanation of content.

5. Information and the explanation of content.

As we have seen, Millikan's view is shaped by reflection on the characteristics of the so-called consumers of representation. The content of representations, on this view, is to be understood in terms of the historical conditions under which the relevant function of the organism or system was selected. Those conditions are partly determined by the sources or the producers of the representation, but only insofar as the producers are relevant to the selective pressures on consumers. In some cases, matters are complicated by the producers and consumers having a co-adaptive relationship, but for the most part her view is that representational content is grounded in the biological functions of the consumer.

Critics have argued that thinking about content in this way simply fails to capture the fact that referring to the truth of the representation is often the most straightforward and natural way of explaining why the consumer's behavior succeeded rather than failed. Why did Roadrunner reject the booby-trapped birthday cake from Coyote? Because, his belief that the cake was a trap was true. Roadrunner's successful avoidance of danger is easiest to explain in terms of the truth of his beliefs.

Godfrey-Smith (1996) and Nicholas Shea (2007) both make the point that while the teleosemantic approach denies that the principal function of representation is to carry information it still seems to assume that a true representation is one that depicts the world in way that captures the conditions that have historically caused a certain kind of success. The notions of success and truth are thoroughly entangled according to these critics such that success over the course of natural history simply cannot explain truth in a non-question-begging manner.

Shea's criticism assumes that Godfrey Smith's slogan 'truth is a fuel for success', is intuitively obvious, at least in the explanation of intelligent behavior in adult humans. He also argues that it can be extended to the explanation of simple representational systems in biology, well beyond its home in folk psychological explanation. He is also committed to the idea that scientists will explain the success of a piece of behavior by reference to the behaviors having been caused by a true representation. (2007, 417) These assumptions are all weaker than they first appear. Stephen Stich for example has long argued against the view that psychological explanations should follow the standard belief-desire model. Alternatives to treating beliefs as simply true or false have been proposed by other philosophers for example by Chris Stephens in his 2001 paper. On Stephens's view, the correct approach is to understand the functional utility of belief and desire forming mechanisms. On his account natural selection is not always going to favor reliable inference methods. Allowing for the possibility that behavior can be successful in spite of the possession false beliefs runs counter to the intuitions underlying folk psychology, but it should not be surprising if the science of simple representational systems provides revisionary results.

Even in the case of adult human cognition, it seems intuitively obvious that certain behaviors might be motivated by false beliefs while still being evolutionarily advantageous. So for example positive self-deception (unrealistically optimistic beliefs, or overestimation of one's own virtues) is a commonly recognized feature of

adult psychology and is also regarded as conducive to reproductive success. (Taylor and Brown 1988)

While we might admit the possibility of a revisionary account of the relationship between truth and success, the evolutionary accounts provided by teleosemantics do seem to rely on something like reliable correlations between representations and states of the world. Nevertheless, carrying information is simply not the proper function of representation on their view. Neither, according to teleosemantics is it the purpose of the producer system to produce patterns or entities that carry correlation information. However, Shea argues that it is simply implausible to explain simple representation without reference to correlational information. Millikan responds by pointing out that the success of the behaviors of humans and animals may depend on any number of different kinds of representations for example maps, beliefs, perceptions, or no representations at all. There may be cases where an informational account is required, but the argument from success to truth does not entail that an explanation of the truthfulness of representations is universally required of all explanations of behavioral success. Recall our discussion in Section 3.1 of the ways that different aspects of folk psychological explanation can be decoupled from one another.

In fact, Millikan does present an account of what she calls “locally recurrent natural information” in her book *Varieties of Meaning*. On first glance, including information in this way might be regarded as a concession to criticisms from Shea and Godfrey-Smith. However, her view of information plays a different explanatory role in her theory from the role Shea sees for correlational information.

What she means by the notion of locally recurrent natural information (henceforth LRNI) is illustrated by examples of coinciding features of the natural world, like, for example the motion of dark clouds and the coming of rain: She writes: “that there is a black cloud at a certain time, t , moving towards a place, p , is a sign that it will rain at p shortly after t .” (2004, 47-8) On her view, these natural signs are useful to an organism insofar as the variables (in simple cases they often represent times and places) allow the organism to apply the understanding of the regular coinciding of black clouds and rain to novel circumstances. The coinciding of the two structures in the natural world is an example of the kind of LRNI that an animal can use to generate some successful behavior. In her response to Shea (Millikan 2007, 453) she clarifies the sense in LRNI are local. What she means by local is not a geographic region rather what she has in mind is something like the domain of a mathematical function. The generality LRNI is restricted to “the set of all actual instances falling under that locally recurring sign type.” On her view, an animal can use LRNI to successfully guide behavior in case the places that they live happen to be such that LRNI are statistically reliable.

Sometimes organisms just happen to live and die within areas where the statistics on a certain recurrent sign are good or good enough. Other times they may develop crude or less crude ways of tracking locally recurrent signs in domains well enough to be useful – ways of tracking that work, at least, in the areas in which they live. (2007, 454)

Shea had argued that correlational information is needed in order to explain how simple representational systems are connected to reality. By contrast Millikan argues that LRNI suffices to address those cases where reference to the truth of a representation needs to be included in an explanation. Nevertheless, as argued above, the fact that an organism can make use of LRNI does not entail that the proper function of the representation is to carry information in the traditional sense advocated by Dretske.

On one level, Shea and Godfrey-Smith argue convincingly that teleosemantic accounts fail to provide a deeper explanation of the kind of truth supporting correlations that obtain between representations and states of the world. However, teleosemantic accounts are, in some sense, orthogonal to that particular explanatory project. Insofar as teleosemantics might seem question begging - assuming the truth of representations rather than explaining it - Millikan suggests in response that reference to history removes the question-begging feature of the explanation.

Shea argues that there is something Panglossian in the teleosemantic approach. Dr. Pangloss, we recall from *Candide* explained the power of sleeping pills to cause his patients to sleep by adverting to their possession of dormitive virtue. Similarly, tokens of representations coincide with their contents because they contributed to the successful behavior that led to their being adaptive. This strikes many readers as question-begging.

Millikan responds by pointing out that Dr. Pangloss' explanation can be somewhat more satisfying once we take into account the history of sleeping pills. If we assumed "that a sleeping pill is something that by definition has in its history that it was selected for manufacture owing to containing something capable of causing sleep." (2007, 437) then it would make perfect sense to say that the pills put John to sleep because they were sleeping pills. The issue between Millikan and Shea concerns the differing kinds of explanations that philosophers seek, or more specifically, the differing standards for what counts as an acceptable explanation of intentional content.

6. Conclusion

The claim that content can be grounded in historical events and that a theory of representation can be based solely on the properties of consumers is controversial insofar as it leaves so many questions unanswered. Some of those questions might seem less urgent or interesting given a commitment to teleosemantics. However, it is important to recognize that many of the criticisms of teleosemantics derive from explanatory goals that Millikan simply doesn't share. Millikan acknowledges that there is a range of distinct explanatory demands that her account will not address for reasons explained above.

Responses to Millikan tend to emphasize the distinction between explanations that are based on the properties of the organism and those that derive from properties of environmental information. While this is a natural reading of the debate over her work, her most significant theoretical contributions are targeted at

the normative aspect of representation and they leave ample space for information-based theories to play a role in explaining other aspects of the problem of representation. Furthermore, in reading Millikan and her critics one must be attentive to the kinds of explanations that are at stake. Sometimes, critics see her work as failing to satisfy explanatory demands that they incorrectly assume she is aiming to meet. It is common to see Millikan and her critics operating on what we might think of as different explanatory registers.

While there is disagreement over the explanatory role of information in theories of representation, there are some obvious commonalities in all viable contemporary theories. To begin with, philosophers agree that organisms have adapted, via natural selection, to respond to their environments in ways that are conducive to their reproductive success. Teleosemantic theorists like Millikan regard such adaptations as continuous with the features that we associate with adult human cognitive capacities. Explanations of the latter will arise out of our understanding of the former. As have seen, for Millikan the kinds of purposes we find in higher level human agency – conscious awareness of one’s purposes for example – are not fundamentally different in kind from the kinds of biological purposes that a reflex or some other automatic biological process might have. Those natural purposes are the foundation upon which her account of representation stands.

Acknowledgments

Thanks to Polo Camacho, Luciano Floridi, and Armin Schulz for insightful comments and criticism of an earlier draft of this paper.

Bibliography

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

Adams, F. and K. Aizawa, (2010) "Causal Theories of Mental Content", *The Stanford Encyclopedia of Philosophy* (Spring 2010 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2010/entries/content-causal/>>.

Bar-Hillel, Y. (1964), *Language and Information*, Reading, MA: Addison-Wesley.

Calvo, P., & Symons, J. (2014). *The Architecture of Cognition: Rethinking Fodor and Pylyshyn's Systematicity Challenge*. MIT Press.

Carnap, Rudolf, and Yehoshua Bar-Hillel. (1952) *An outline of a theory of semantic information*. Research Laboratory of Electronics, MIT.

Dennett, D. (1987) *The Intentional Stance*. MIT Press, Cambridge, MA

Dennett 1988 'Evolution, Error and Intentionality in Y. Wilks and D. Partridge, eds, *Sourcebook on the Foundations of Artificial Intelligence*, New Mexico University Press 1988.

Dretske, F., 1981, *Knowledge and the Flow of Information*, Cambridge, MA: MIT/Bradford Press.

Floridi, L. (forthcoming) "Perception and Testimony as Data Providers" *Logique et Analyse*

Fodor, J. (1996) "Deconstructing Dennett's Darwin", in *Mind and Language* 11: 246–262.

Godfrey-Smith, P. (1996) *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press

Meinong, A. (1899) "Über Gegenstände höherer Ordnung und deren Verhältniss zur inneren Wahrnehmung" in *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, 21, pp. 187–272.

Millikan, R., (1984), *Language, Thought and Other Biological Categories*, Cambridge, MA: MIT Press.

Millikan, R., (1989) "In Defense of Proper Functions", *Philosophy of Science*, 56, no. 2: 288–302,

Millikan, R., (1993), *White Queen Psychology and Other Essays for Alice*, Cambridge, MA: MIT Press.

Millikan, R., (2004), *Varieties of Meaning*, Cambridge, Mass: MIT Press.

Millikan, R., (2007) An Input Condition for Teleosemantics? Reply to Shea (and Godfrey-Smith) *Philosophy and Phenomenological Research*, 75(2), 436-455.

Neander, K. (2012) "Teleological Theories of Mental Content", *The Stanford Encyclopedia of Philosophy* (Spring 2012 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2012/entries/content-teleological/>>.

Papineau, D. (1984), "Representation and Explanation", in *Philosophy of Science*, 51: 550–72.

Papineau, D. (1987), *Reality and Representation*, Oxford: Basil Blackwell.

Papineau, D. (1997), "Teleosemantics and Indeterminacy", in *Australasian Journal of Philosophy*, 76: 1-14.

Pietroski, P., 1992, "Intentional and Teleological Error", in *Pacific Philosophical Quarterly*, 73: 267-81.

Plantinga, A. (1993). *Warrant and proper function*. Oxford: Oxford University Press

Rowlands M. (1997) 'Teleological Semantics', in *Mind*, Vol. 106, no. 422, April 1997, 279-303.

Shea, N. (2007). Consumers need information: supplementing teleosemantics with an input condition. *Philosophy and Phenomenological Research*, 75(2), 404-435.

Stampe, D., (1977) "Toward a Causal Theory of Linguistic Representation," in P. French, H. K. Wettstein, and T. E. Uehling (eds.), *Midwest Studies in Philosophy*, vol. 2, Minneapolis: University of Minnesota Press, pp. 42-63.

Stegmann, U. (2013). *Animal communication theory: information and influence*. Cambridge University Press.

Stephens, C. (2001) "When is it Selectively Advantageous to Have True Beliefs? Sandwiching the Better-Safe-than-Sorry Argument", *Philosophical Studies*, vol. 105, number 2, (2001) p. 161-189

Stich, S. 1983, *From Folk Psychology to Cognitive Science: The Case Against Belief*, MIT Press.

Stich, S. 1985, "Could Man Be An Irrational Animal?" *Synthese*, Vol. 64, No. 1, pp. 115-135.

Stich, S. 1990, *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*, MIT Press.

Symons, J. (2001). Explanation, representation and the dynamical hypothesis. *Minds and Machines*, 11(4), 521-541.

Symons, J. (2002) *On Dennett* Belmont: Wadsworth

Symons, John. "The Individuality of Artifacts and Organisms." *History and philosophy of the life sciences* (2010): 233-246.

Taylor, S.E.; Brown, J. (1988). "Illusion and well-being: A social psychological perspective on mental health". *Psychological Bulletin* 103 (2): 193-210