

# TOWARD AN INFORMATIONAL TELEOSEMANTICS

KAREN NEANDER

One might think that a natural place to look for mental content is in the information carrying functions of mental representations. One might even think that the proponents of teleosemantics would especially like to look there. Fred Dretske (1986) does look there. However, Ruth Millikan and David Papineau, who have most extensively developed and defended teleosemantics, are critics of this approach. They deny that representations have the function to carry information or that representation producers have the function to produce representations that carry information. In my view, they are wrong, and once we see why they are wrong we can develop an alternative approach, an informational yet teleosemantic theory of content. This chapter makes a start with the contents of sensory representations.

## Introduction

Informational theories use a “natural” (as opposed to intentional) notion of information.<sup>1</sup> The starting idea is that sensory representations and perhaps other perceptual, conceptual, and belief-like states, carry information *about* things in the world insofar as they are caused by them, are correlated with them, or make their presence more probable. On this type of view, the “aboutness” of content is thought to derive from the “aboutness” of natural information. Teleological theories of mental content instead (or possibly, as well) appeal to functions that are given an etiological analysis. A function of something is what it was selected for.<sup>2</sup> Content is said to be “normative” because representations can

<sup>1</sup> See esp. Dretske (1981) and Fodor (1990).

<sup>2</sup> See esp. Millikan (LTOBC, 1989a); Papineau (1987).

misrepresent. And these functions are said to be “normative” too because items with such functions can malfunction. For teleosemantics, the core idea is that the “normativity” of content derives from the “normativity” of functions.

Each idea is further refined in different ways by different people but, at a first pass, teleosemantic theories, by invoking functions, are said to focus on the effects or results of representations, whereas informational theories, by invoking natural information, are said to focus on the causes of representations or on conditions that co-occur with their causes. The difference between the two approaches can thus seem to be one of output-based versus input-based views of the constitutive conditions for mental content. One might nonetheless think that the two approaches should be combined. A hybrid idea is that representation-producing systems have the function of producing states that carry natural information. Then the “aboutness” of content could be explained as originating with the “aboutness” of information *and* the “normativity” of content could be explained as originating with the “normativity” of functions. Here is the Holy Grail! A source of naturalistic normative aboutness! Or so one might think.

Yet Millikan scotches this idea. In her early criticism of informational theories,<sup>3</sup> she says that “... [i]f ‘detecting’ is a function of a representational state it must be something that the state effects or produces. For example, it cannot be the function of a state to have been produced in response to something” (Millikan 1989a, 283).<sup>4</sup> She then argues that a theory of content should ignore the production of representations altogether, in favor of their “consumption.” More recently, this emphasis on consumption versus production has been softened in her theory. Now Millikan speaks of co-adaptation between producing and consuming systems. However, her view is still that “... intentional signs often do carry natural information ... [but] ... it is not the purpose of an intentional sign to carry natural information” (VOM, 31). And, when discussing Nicholas Shea’s (2007b) argument for an input condition for teleosemantics, she approves his observation that, on her view, “... carrying information is not a function of a representation” and “... it is not a purpose of the producer system to produce items that carry correlational [or I would add [she adds] any other kind of natural] information” (Millikan 2007, 445).

David Papineau (1998, 3) agrees with Millikan about this and chides me for ignoring her point:

I would argue that Neander has taken insufficient note of Millikan’s point that representational content hinges on how the representation is used, not on what causes it. In her general discussion of teleology, Neander focuses, quite rightly, on the effects of biological traits. But as soon as she turns to representation she shifts to the question of what ... it is supposed to detect.... From the teleological point of view this is to start at the wrong end. The teleosemantic strategy requires us first to identify which result the state is supposed to produce and then use this to tell us what it is representing. (Papineau 1998, 7)

I am not about to repent. This chapter explains why, despite being in broad sympathy with Millikan and Papineau, I believe that asking what a system is supposed to detect is to start

<sup>3</sup> Including the teleonomically inclined, such as Stampe (1979), Dretske (1986), and Matthen (1988).

<sup>4</sup> See also Millikan (WQ, ch. 11, OCCI, VOM chs. 3–6, 2007).

in the right place.<sup>5</sup> My plan is first to explain why there are what I call “response functions,” which are functions to respond to something by doing something (in section 2). Sensory systems, I claim, have response functions. Then (in section 3) I argue that sensory systems have the function of producing representations that carry natural information, on a certain (singular, causal) understanding of information that I recommend. Section 4 briefly considers the implications of this for what we can say about the informational functions of the sensory representations themselves (the argument in section 4 is not crucial). Then I give a two-part proposal toward an informational teleosemantic (IT) theory, which is for the indicative content of sensory representations. Section 5 explains how, on this version of IT, the problem of error is handled and section 6 explains how the problem of distal content is handled.

## Response Functions

A system has a response function if it has the function to respond to something by doing something else. The term ‘response’ is used in a purely causal and non-intentional sense here; a system counts as producing a RED-state in response to red if some red (the instantiating of red) caused the system to produce a RED-state.

The first question I want to consider is whether there can be response functions. Against this, Millikan (2007, 447) claims that “The function of an item, in the teleologist’s sense, is always something it effects.” But, before hastening to agree with this apparent truism, bear in mind that the issue is not whether functions must involve effects, which they must, but whether they must be effects to the exclusion of input causes. Might not a sensory system have the function, in part, to produce certain inner states in response to certain features of or occurrences in the environment?

A philosophical problem about biological functions is often posed as the problem of elucidating which effect(s) of an item count as its function(s). Why is it the function of the heart to pump blood and to circulate it through the body? Why are these effects its functions, not its making a whooshing noise? It is assumed at the outset that functions are certain effects of the items that have them, or anyway that functions are certain effects that they have in appropriate circumstances when they are functioning properly and not malfunctioning. But let’s look beyond the initial posing of the problem of functions to consult our best solution to it.

The proponents of teleosemantics invariably agree that the best solution is the etiological theory, and this includes Millikan and Papineau<sup>6</sup> and me. Leaving aside details that need not concern us now, the etiological theory says that functions are what things were selected for.

The details are variably elucidated, but note that phenotypic traits can be selected for the adaptive things that they *do*, and *doings* can involve input causes. Consider an analogous case first.

<sup>5</sup> And “[a natural sign] carries information if it is true and has been normally produced, that is all” (Millikan, 2007, 445–448).

<sup>6</sup> See Millikan (LTOBC, 1989b). Papineau (1987) approves of the etiological theory that I give in my (1983) PhD dissertation and later defend in Neander (1991).

If I am waiting for my friends to phone to tell me if they're coming to dinner, there are different things that I can do. I can pick up the phone at random or whenever I remember that my friends will call, or I can wait until it rings and pick it up only then. If I pick it up at random or whenever I remember that my friends will call, I'll waste my time and possibly frustrate my friends' efforts to get through. So doing one thing rather than another will serve my purposes better. And the crucial difference in these different things that I could do are differences in the input cause, or in other words in what I am responding to.

Now think about those biological functions that depend on natural selection. One type of mechanism might secrete melatonin in response to the dimming of light,<sup>7</sup> while another type of mechanism might secrete melatonin in response to light brightening and a third more or less randomly. Since melatonin makes us sleepy, the first mechanism might be more adaptive in creatures like us who have poor night vision. So the first type of mechanism might be selected in preference to the second and third because the first differs in its input cause, in what triggers its production of melatonin.

It is a function of our pineal gland to secrete melatonin in response to the dimming of light. It is also a function of B-lymphocytes to produce antibodies in response to antigens, a function of certain mechanisms to produce shivering in response to the early stages of hypothermia, and a function of the pancreas to secrete insulin in response to elevated levels of glucose in the blood.

A visual system can have the function, in part, to produce REDs in response to red being instantiated. Or, more generally, there can be a range of determinate values,  $e_1 \dots e_n$ , of a determinable (E), and a range of determinate values,  $r_1 \dots r_n$ , of another determinable (R), and a sensory system can have the function, in part, to produce certain values of R in response to certain values of E. It can have the function to produce  $r_1$  in response to  $e_1$ ,  $r_2$  in response to  $e_2$ ,  $r_3$  in response to  $e_3$ , and so on.<sup>8</sup>

To be sure, functions *must* also involve effects. For natural selection to occur, there must be downstream replication. But, to repeat, the issue is not whether functions involve effects but whether they can also involve input causes. And they can. Biological mechanisms are selected for their causal roles, which can include dispositions to respond to specific types of causes. Someone could *stipulate* that by the term 'function' she will refer to *effects* exclusively. But denying that there are response functions is meant to have significant consequences for how teleosemantics should be developed and mere stipulation is no motivation.

I am anyway not being revisionist. Let us take a look at some statements of the etiological theory. Larry Wright (1976, 81), for instance, says that "The function of X is Z iff: (i) Z is a consequence (result) of X's being there, and (ii) X is there because it does (results in) Z." A result of a sensory system "being there" can be its producing inner states in response to outer states and this can be, in part, why it is there (once we appropriately disentangle talk of types versus tokens, which we would anyway need to do here).

<sup>7</sup> More fully, in response to the norepinephrine produced in retinal photoreceptors in response to the dimming of light.

<sup>8</sup> It could be that the inner states are caused by the outer states in a way that involves second-order similarity or some other mapping, a complication that I don't explore but intend to allow here.

If we change the symbols for the sake of conformity, Millikan (1989b, 288) says that X has the proper function, Z, "... if X originated as a reproduction ... of some prior item or items that, due in part to possession of the properties reproduced, have actually performed Z in the past, and X exists because (causally, historically because) of this or these performances." Presumably, a performance could be a simple effecting. However, without an explicit ruling against this, it could also be an effecting of certain inner states in response to certain conditions obtaining.

In a paper titled "Functions as Selected Effects" I argue that functions are selected effects, thus apparently committing myself to the claim that I here deny. But my point there was that functions are effects of phenotypic traits for which the traits were selected rather than effects of traits that are currently adaptive. That is, the emphasis was on functions being *selected* effects, not selected *effects*. Of ordinary physiological functions, I say, "[i]t is the/a proper function of an item (X) of an organism (O) to do that which items of X's type did to contribute to the inclusive fitness of O's ancestors, and which caused the genotype, of which X is the phenotypic expression, to be selected by natural selection" (Neander 1991, 174). Again, response functions are not excluded. Something that a sensory system can do and be selected for doing is produce an inner state in response to a certain type of cause.

In short, that functions are selected effects is true enough for most dialectical situations, but don't mistake the slogan for the theory. The etiological theory, as standardly formulated, and as it should be formulated in my view, does not rule out response functions.

That functions are effects and not causes is often taken to be too obvious to need arguing (Millikan's and Papineau's claims are often repeated by others). However, Shea (2007b) offers some motivation for the claim. He says:

Teleosemantics is sometimes mistakenly taken to be a refinement of informational semantics according to which items represent what they were designed by evolution to carry information about. The trouble with this gloss is that evolutionary functions are a matter of effects. Amongst the various effects that an evolved system can produce, those which have contributed systematically to the system's survival and reproduction in the past are its evolutionary functions. Evolution acts only on effects. It is blind to the mechanism by which those effects are produced. An effect produced by a system at random will be an evolutionary function if its production has contributed systematically to the survival and reproduction of the system. It is not part of any evolutionary function for the effect to have been caused in a particular way. (Shea 2007b, 409; italics added)

Shea goes on to argue that we could nonetheless still *supplement* teleosemantics with an input condition, and this is an important point. Teleosemantic theories can use other ingredients, besides functions, when they give a recipe for thought. But in defending the possibility of an input-based teleosemantics, Shea concedes more than he needs to.

It is true that natural selection cannot discriminate between two mechanisms that have the same impact on fitness. If one type of pineal gland with associated structures had randomly secreted melatonin *and* had contributed to survival and reproduction just as well as a system that secreted melatonin in response to the dimming of light, then natural

selection could not have discriminated between them. More generally, something that gets something done one way and something that gets exactly the same thing done another way are equal in the eyes of selection (all else being equal).<sup>9</sup>

However, there are cases where different things are done that have different impacts on fitness because they involve responses to different causes. It is safe to assume that a system that secreted melatonin randomly did *not* contribute to the survival and reproduction of our ancestors as effectively as one that secreted melatonin in response to the dimming of light. And it is safe to assume that a human visual system that produced REDs randomly did *not* contribute to the survival and reproduction of our ancestors as effectively as one that produced REDs in response to red. Given the complexity of the neurobiological systems involved, the inference to the best explanation is that natural selection selected the relevant pathways in our visual system, in part, for their disposition to respond to red by producing REDs (and to respond to green by producing GREENs and so on).

A difference in the triggering causes of a process cannot make a difference to fitness unless the difference in triggering causes *makes a difference* and impacts fitness. However, it does not follow that natural selection is blind to triggering causes (or mechanisms) and only sensitive to effects. Rather, natural selection is sensitive to triggering causes (and mechanisms) because it is sensitive to the different effects that different triggering causes (and different mechanisms) can bring about.

## Information and Singular Causation

How do response functions bear on information carrying functions? In a while, I'll argue for a different understanding of natural information, but it helps to appreciate a couple of points that Millikan makes about natural information construed as correlation first.

If a sensory system were selected, in part, for producing REDs in response to instantiations of red, this could increase the probability of a RED given red and the probability of some red given RED. But, as Millikan (2007) reminds us, a result of a system's normal functioning is not necessarily a result that it has the function to effect (not all that it normally does is something that it has the function to do). The normal functioning of the heart results in a whooshing sound, but the heart does not have the function to make a whooshing sound.

Further, correlations hold between *types* and whether two types of things are correlated and to what extent they are correlated depends on what is counted. As Millikan explains, all sorts of troubles follow on this score. An informational teleosemantics that interprets information as correlation will have to pay careful attention to specifying the relevant reference classes and it is hard to see how to appropriately circumscribe them.

In any event, Millikan argues, it cannot be the function of any one visual system to bring about such a correlation among REDs and red instantiations and thus ensure that

<sup>9</sup> See also Rosenberg and McShea (2008, 96–126).

its REDs carry information about red, no matter how the reference class is specified. My visual system does not have the function of ensuring that your visual system tends to produce REDs in (or only in) the presence of red.

Even if the correlation need only hold on an individual basis (i.e., even if the function of my visual system is to ensure that my REDs correlate with red in my vicinity), I would argue that there is a related problem. Johnny's visual system might function perfectly well until he is six, producing REDs when and only when he sees red. But his visual system could malfunction in the future, with the result that, for Johnny, his lifelong record of REDs in and only in the presence of red is very poor. By the time that he is celebrating his ninety-sixth birthday, his REDs might have correlated far more highly with green instantiations instead. If the function of his visual system is to ensure an individual lifelong correlation between REDs and red, it seems that in this scenario it never performed its function properly. The trouble with this is that, for his first six years, his visual system was functioning properly, and it was producing REDs that carry information about red, despite its later failure.

At this juncture we could, with Millikan, abandon the idea that sensory systems have the function of producing representations that carry natural information. Or we could explore other ways to understand natural information. Millikan explores other ways with no better outcome but there is one that she does not explore.

I suggest that  $r$  carries the natural indicative information that  $e$  if  $e$  is a cause of  $r$ , where this is singular causation and  $r$  and  $e$  are particulars.<sup>10</sup> Thus a particular RED carries the information that there is red if red being instantiated caused RED.<sup>11</sup> Sensory systems have the function to produce inner states in response to outer states (as I've argued). So they will have the function to produce inner states that carry information about outer states, on this proposal.

On this proposal, there is no difficulty of the kind that Millikan has raised for other informational proposals. On this proposal the information relation holds primarily between tokens and only secondarily between types.<sup>12</sup> In what follows, I will sometimes refer to this singular causal indicative information as 'information' (or 'natural information') without further qualification.

<sup>10</sup> In the end (elsewhere) I will say that  $r$  carries the information that  $e$  when  $e$  causes  $r$  or  $r$  causes  $e$ , which lets motor instructions carry information about the movements they cause. Thus I say that this is for *indicative* information.

<sup>11</sup> Fodor (1990) sometimes seems to have this in mind. Millikan (2001) thinks of information as "grounded" correlation. Millikan (VOM, 36) seems sympathetic to the idea that information is a matter of singular causation. But then she introduces a notion of a recurrent local sign (which reintroduces statistical considerations) and focuses on that. Of the latter, she says, a causal connection is not required (VOM, 44). Her concern there (in chapter 3 of VOM) is different from my concern here. Millikan is there concerned with one feature of the environment being a sign of another feature of the environment (e.g., magnetic north being a sign of the direction of oxygen-free water) and with a creature being able to use one as a sign of the other. Hence the connection between sign and signified, she says, needs to be at least locally recurrent. This epistemic issue is not my concern here.

<sup>12</sup> It holds between types only insofar as we aggregate what holds between tokens. For example, most REDs will carry information about red if most RED-tokens are caused by red instantiations.

It is often said that natural information must be “grounded.” On the present proposal, if  $e$  causes  $r$  and so  $r$  carries information about  $e$ , this will be in accord with natural laws and not a miracle. But singular causal information does not require that, if some REDs carry the information that there is red, there is a law to the effect that red and RED are correlated, or that there be reliable head-to-world indication, or that red instantiations always cause REDs or that REDs are only caused by red instantiations. A sensory system of a certain type can be selected for producing  $rs$  in response to  $es$  in one species, one part of the world, or one period of time.

Natural information, so understood, has a key feature that natural information is supposed to have. Crucially, misinformation is impossible. A RED carries the information that there is red at a certain location only if red at that location caused that RED tokening. Thus, if a RED carries the information that there is red, there is red.<sup>13</sup>

For what it is worth, this understanding of natural information is consistent with many of our intuitions about natural signs. Echoing, with modifications, Grice and Dretske, we take six rings in a cross-section of a tree to be a natural sign that the tree is six years old, but this is defeasible if the tree isn't six years old. If an unseasonable cold spell one spring added an extra dark ring and the tree is only five years old instead, the fact of there being six rings is not a natural sign that the tree is six years old. Six years of growth must cause the six rings, if they are to be a natural sign that the tree has been growing for six years. The red spots on Johnny's face signify that he has the measles only if he has the measles. If Johnny does not in fact have the measles, then the presence of such spots is not a natural sign that he does.

Even if Johnny has the measles, I would add, his spots do not carry the information that he has them *unless his measles cause his spots*. Imagine that Johnny has a mutation that prevents his developing spots in the course of the measles. If his spots were coincidentally caused during a bout of the measles by an allergic reaction to strawberries, so that his having the measles were a mere coincidence, his spots would not carry the information that he has the measles, no matter how similar his spots otherwise were to those usually caused by the measles.

Johnny's spots would not carry the information that he has the measles if his having the measles were a mere coincidence, even if the correlation (speaking in terms of mere frequency) between measles and the presence of such spots were perfect.

Perhaps Johnny's allergic reaction to strawberries is rare, or even also unique. Someone might balk at the implication of my proposal that Johnny's spots carry information about his allergy to strawberries. Such a rare causal relation will not be *informative* to us. But take care not to conflate an epistemic sense of information with the natural notion that we want here. Natural selection will bring it about that our sensory systems are adapted to produce inner states that carry information that our brains can exploit and it will bring it about that our brains are adapted to exploit the information that our sensory systems are adapted to provide. But that natural information be exploitable (let alone for adaptive

<sup>13</sup> Does this preserve the informational entailments of concern to Dretske (1981)? Perhaps. If a RED carries the information that there is red and a LINE <sub>$\theta$</sub>  carries the information that there is a line with a certain orientation, the conjunction of that RED and that LINE <sub>$\theta$</sub>  carries the information that there is red and a line so orientated.



behavioral responses or for the acquisition of knowledge) is not something that our notion of natural information needs to guarantee. A great deal could be useless, epistemically. Natural information *as such* need not be even locally recurrent. Natural selection can be left to sort out what is usable. Our analysis of natural information need not do that. Sensory systems will not be adapted to pick up on information unless it is at least locally recurrent, but this is not something that must be built into our analysis of natural information.

We could, if we want, skip talk of information and go straight to a causal theory of content. However, on this proposal, we can make good sense of the claim that sensory systems have the function of producing inner states that carry information about outer states. Their function is, in part, to causally mediate between the outer and inner states and hence to make it the case that certain inner states carry natural information about certain outer states. They have, in other words, information processing functions. Making sense of this way of speaking is salutary.

Singular causal information is strict enough without being overly strict. To use one of Dretske's well-known examples, suppose that his neighbors begin making their doorbells out of nuts and squirrels start chewing on them and making them ring. Doorbell systems could still be useful in his neighborhood and people could still install them and use them if *often enough* when a doorbell rings someone is causing it to ring. A parallel point can be made for biological functions and natural selection. If our ancestors experienced color contrast illusions, or the light wasn't always white light, with the result that REDs were sometimes produced in the absence of red, certain pathways in the human visual system could nonetheless have been adaptive and could have been selected because *often enough* the REDs they produced carried information about red.<sup>14</sup>

## The Functions of Sensory Representations

Until now, I have discussed the functions of sensory *systems* and not, or not directly, the functions of the *states* that it is the function of these systems to produce. There are two terminological issues to be decided before we conclude (if we do) that the sensory representations themselves have response functions, or functions to carry information.

One is whether only systems, structures, mechanisms, and so on can count as having functions or whether their states have functions too. Someone might allow that the left ventricle of the heart has the function of relaxing after it contracts in order to let in blood, and yet deny that the relaxation of the left ventricle, after its contraction, has the function of letting in blood. I see no good reason why we should stick at this.

A second matter is this. Even if we grant that a sensory system has the function of producing REDs in response to red, and that the RED states can count as having a function of effecting something, one might still protest at the idea that these RED-states have the function of being produced in response to instantiations of red. As Millikan (2007, 447)

<sup>14</sup> Thus this proposal, though it is on Dretske's side of the teleosemantic family feud, avoids the worries that Millikan (OCCI, VOM) and others raise for Dretske's (1981) strict definition of 'indication'.

says, "... a thing cannot effect its own history." Nothing can, so to speak, have the responsibility to bring itself into existence in a certain way. Functions are things that were done, that were adaptive, and that caused things of a type (the type of thing that did them) to be selected. Coming into existence in a certain way isn't something that anything can do. Or at any rate it isn't something that anything can be "responsible" for doing.

However, it is of little consequence what we say about this for present purposes, because even if it is infelicitous to ascribe a response function to a sensory *representation*, we can always translate such an infelicitous ascription into an ascription of a response function to the sensory *system* that produces the representation. Crucially, to say that a sensory system has a response function is not to say that *the system* is responsible for bringing itself about. It is to say that it is responsible for bringing about a state change in itself in response to a certain type of triggering cause.

We may also say that REDs are "supposed" to be caused by red instantiations (unless it is objectionable to speak in this colloquial style) if we read this as saying that the RED-producing system has the function of producing REDs in response to instantiations of red.

### **The Contents of Sensory Representations: The Problem of Error**

Finally, we come to the content of sensory representations. How does any of this bear on their content? The now all too obvious suggestion, or so I hope, is that a sensory representation, RED, has the indicative content that there is red if the RED-producing system has the function of producing REDs in response to red, or (synonymously) if REDs are supposed to carry the information that there is red.<sup>15</sup> Similarly, a sensory representation, LINE<sub>*i*</sub>, will have the indicative content that there is a line with a certain orientation if the LINE<sub>*i*</sub>-producing system has the function of producing LINE<sub>*i*</sub>s in response to lines with that orientation, or (synonymously) if LINE<sub>*i*</sub>s are supposed to carry the information that there is a line with that orientation.

Let's look at a third example.<sup>16</sup> Simplifying somewhat, this involves representation by certain tectal (T) cells in toads. These cells are referred to as the T(5)2 cells, and they normally fire strongly in response to any stimulus that is within certain size parameters and elongated and moving in a direction that parallels its longest axis.<sup>17</sup> Neuroethologists refer to this configuration of features, for short, as "worm-like motion." Let '*r*' stand for the relevant sensory representation in the toad. Often enough, the stimulus to which *r* is a response is something nutritious, such as a worm, a cricket, a millipede, or some other nutritious tidbit (even a bird, frog, or a smaller toad, for toads are not fussy eaters). However, a toad cannot distinguish something that is nutritious from something that is not nutritious.

<sup>15</sup> I'm not implying that sensory representations have a syntactic structure isomorphic to the grammatical structure of the sentence, "There is red."

<sup>16</sup> See Neander (2006) for more detailed discussion of the example and the points made in this section.

<sup>17</sup> The content isn't necessarily the property that triggers the strongest reaction. Super stimuli might never have occurred in the past, let alone been systematically involved in past selection.

When  $r$  is tokened, the toad orients toward the stimulus, moves around obstacles and approaches if necessary, and then it tries to catch and swallow what is, if it is lucky, some nutritious prey. The tokening of  $r$  initiates the orienting and prey-capture behavior, but more information processing concerning localization and motor coordination is involved before the prey is caught.

The information processing that produces  $r$  is not mere transduction, as has sometimes been suggested in philosophical papers. It is not completed in the retina or even in the retinal ganglion cells. The processing that produces  $r$  involves mid-brain structures and it is complicated enough to remain not yet fully understood even after intensive neuroethological study.

The relevant pathways in the toad's brain normally produce this representation in response to the configuration of features characterized as worm-like motion, and it is safe to assume that they were adapted in part for doing so and hence have the function of doing so. Thus, on the present proposal, the  $r$ -states that it is their function to produce are supposed to carry information about worm-like motion.

A complication is that the activity in the toad's tectal cells normally carries both *what* and *where* information and they were adapted for doing so. So they have both *what* and *where* content on IT. There are different T(5)2 cells, and the different T(5)2 cells respond to moving stimuli in different parts of the toad's visual field (by which I mean different parts of the outside world, as it comes within the toad's field of view). Neighboring regions of the toad's retinas feed information forward to (causally feed forward to) neighboring regions in the tectum. So which T(5)2 cells fire tells the toad's brain the (rough) location of the moving stimuli. They tell it that *there* is worm-like motion *there*, and where *there* is depends on which T(5)2 cells fire. Thus, there is a range of relevant representations ( $r_1 \dots r_n$ ) with different localization contents.

Elsewhere (Neander 2006) I argue that this is the right content to attribute to  $r_1 \dots r_n$  if they are to play a role in the information-processing explanation of the toad's capacities that the neuroethologist seeks to provide. My argument there leans on the fact that the perceptual capacity to be explained is a capacity to distinguish between items in worm-like motion and items that are not in worm-like motion. This is the capacity to be explained because it is the capacity that the normal toad has. In contrast, the normal toad has *no* capacity to distinguish nutritional from non-nutritional substances and so this is not the capacity to be explained.

If we want to explain why the toad's capacity to detect worm-like motion *evolved*, then the relationship between worm-like motion and nutritional substances will be paramount. But, if we want to understand the toad's perceptual capacity and its prey-capture behavior and the role of its representations in this capacity we cannot afford to ignore the information processing that produces the representations. It is a mistake to think that talk of proper or normal functions has as its main or anyway its sole aim the explaining of the evolutionary origins of things and their capacities. Talk of normal or proper functions also has an important role in physiology and in neurophysiology, in explaining how the capacities are realized, or in other words in explaining how the *normal* system works. Informational teleosemantics (mine, anyway) has its eye on explanations like these and the role of representations in them.

The normal toad's brain processes information about the configuration of visible features that constitutes what we're calling "worm-like motion." That is, its perceptual processing in the detection of prey concerns the size, shape, motion, and motion relative to shape of the stimulus. There is no justification from a neurobiological point of view for the assumption that there is any inference on the part of the toad from this configuration of visible features to the presence of invisible nutrients. Such an inference surely does not occur. And even if there were a further inference to the presence of nutrients, it is not the content of the conclusion of any further inference that is under discussion. What is under discussion is the content of a sensory representation. To ignore the underlying information processing that produces this representation is to risk making the content that we ascribe to the sensory representation irrelevant to the explanation of the toad's perceptual capacity.

The pathways that produce *r* have the function to do so in response to a certain type of cause (worm-like motion) but they also have the function to effect certain things thereby. Most directly they have the function to initiate certain behavioral responses, starting with the orientation of the toad toward the worm-like motion that has been detected. Notice that when we focus on the further effects of the representation (on its "consumption") we are not anyway forced to start talking about nutritional substances right away. The components in the toad's brain that go on to more precisely locate the stimulus, and that go on to control the orienting, approaching, and catching behavior of the toad, have functions that can be described in terms of the item in worm-like motion that has been detected, rather than in terms of any nutritive potential it has.<sup>18</sup>

Millikan (2001, 105; cf. OCC1, 217) worries that informational semantics (including informational *teleosemantics*) implies a kind of verificationism but, to be clear, I am only making a claim about sensory representations here. I agree that we have what Millikan (OCCI) calls "substance concepts." I agree that we can have a concept of X but not Y, even if we cannot distinguish between Xs and Ys (extensionally speaking, we can have a concept of H<sub>2</sub>O but not of X<sub>γ</sub>Z, even if we cannot distinguish between them). But a toad has no substance concepts. Nor does a sensory system. Substance concepts will be involved in sophisticated acts of perceptual recognition or post-perceptual recognition, such as perceiving or judging that something is gold or a lion or recognizing Ruth when she walks into the room. The contents of sensory representations, in contrast, are constrained by what can be sensed and so by the distinctions to which our sensory systems are causally sensitive. Though some of our concepts must be able to refer to things with hidden structural or historical "essences," sensory representations do not do so. Sensory

<sup>18</sup> In Millikan's theory of content, it is not the focus on consuming systems *as such* that dictates that the indicative content of *r* is something nutritious (though her theory does, she says, entail this). What principally does that job in her theory is the notion of a Normal condition. As I understand it, the Normal condition is an environmental condition that corresponds (more or less, but well enough) with past uses of a representation and that crucially explains why its use was beneficial. The latter is cashed out in terms of the condition allowing the consuming system to perform its function. The Normal condition need not be the condition to which a sensory system is causally sensitive. It need not be the condition to which it is responding. Hence Pietroski's (1992) worry about the kimu who cannot tell a snorf from a non-snorf though, on Millikan's theory, the kimu love snorf-free spaces and seek them out.

representations represent the visible properties of things. How we get from them to substance concepts is a separate question.

In my view, the relevant tectal firings in the toad represent the visible properties of a toad's prey, not its invisible properties (such as its ability to nourish the toad). The version of IT that I am offering gives that result. IT is inconsistent with the possibility that  $r_s$  refer to toad food (or at least with the possibility that this is their indicative content) because the relevant pathways in the toad were not selected for causal sensitivity to toad food as such.

In some neurologically damaged toads, "anti-worm motion" by something that is elongated but moving perpendicular to its longest axis can trigger a spike in T(5)2 firing. In some neurologically damaged toads, large predator-like objects (e.g., the laboratory technician's hand) can trigger it too. IT counts the resulting  $r$ -states as erroneous because the relevant pathways in the toad were not selected for responding to such stimuli as these by producing  $r$ -states. These inappropriate responses on the part of a neurologically damaged toad also count as *abnormal* for the neuroethologist. That is, strong firing in T(5)2 cells in response to a worm that is hung by its tail and moved sideways (in anti-worm motion) is considered abnormal, no matter how nutritious the worm is. So is strong firing in response to a large (predator-like) chunk of processed toad food, even if a good bite out of that chunk would feed a hungry toad. The notion of normality invoked here is teleonomic and not statistical. It is the notion of normal invoked in talk of normal functioning that is synonymous with talk of proper functioning.

IT does not equate correct representation with representations produced as a result of normal functioning, even in the case of sensory systems, but it brings the two into closer alignment, as just noted. In neurologically normal toads, an electrode in the right place could cause  $r_s$  to be produced and that would also be erroneous although no malfunctioning would be involved. But to explain this further, we need the second part of IT, given in the next section.

### **The Contents of Sensory Representation: The Distality Problem**

In her "What has Natural Information to do with Intentional Representation" (2001, 118; cf. OCCI, 231), Millikan says that Neander (1995) "seriously claims that all representations must only be of proximal stimuli." As it happens, Millikan and I agree that there is distal content. Toward the end of my "Misrepresenting and Malfunctioning" (1995, hereafter M&M) I raised the problem of distal content as an, as yet, unsolved problem for my proposal in that paper. I had argued that several content indeterminacy problems had been conflated and that we needed to adopt a divide and conquer approach to them. And, while I did not tackle the problem of distal content, I did not deny that it needed a solution (on the contrary, see esp. Neander 1995, 136). I did not try to solve it then because I did not know how. But let me try to solve it now.

Suppose that you agree with me so far. (Just suppose.) So you agree, if only for the sake of the argument, that  $r_1 \dots r_n$  are, so to speak, supposed to carry information about worm-like motion in various locations and that this is therefore what they represent. Now

consider a specific representation in the range ( $r_i$ ) and the singular causal chain represented on the following line:

Wormy motion at L  $\Rightarrow$  light  $\Rightarrow$  *retinal firings*  $\Rightarrow$  *other stuff*  $\Rightarrow$   $r_i$

The problem of distal content in this case is to say why  $r_i$  represents worm-like motion at location L rather than the more proximal things that carry information about the worm-like motion to  $r_i$ . The more proximal things include the light reflected from the item in worm-like motion and the retinal firings that result when this light hits the retina and various other neurophysiological events that occur in brain structures *en route* to the tectum and the tokening of  $r_i$ .

It is an interesting question just how the problem of distal content relates to the contents of representations of depth. Just briefly (as explained in Neander 2006) the toad does have some depth perception. But  $r_i$  itself is not thought to carry information about depth. In part, this is because it is only when the toad subsequently orients and fixates on the worm-like motion that it can see it with both eyes. Anyway, as is usual, I'll treat the problem of distal content as in need of a solution that does not depend on the representation of depth relations, although these might not be wholly separate matters. 'Location L' refers to a region within the toad's field of view and this region will include both near and far. Our question is not how to account for the more precise localization content on the near-far dimension that later representations possess. Our question is how we can count  $r_i$  itself as representing the distal worm-like motion rather than the light waves reflected from the item whose motion it is, or the retinal firings that result, and so on.

The intuition behind my idea for a solution is that a sensory system is only adapted to respond to the more proximal items because doing so is the means by which it responds to the more distal ones, and not vice versa. This is a principle I wish to apply only to the determination of content on the distal-proximal dimension. Some of the more proximal items are states of the system itself, and the sensory system was not selected for carrying information about itself. So those states are ruled out anyway. But there are also proximal items outside of the system, namely the patterns of light. So, here goes.

Let there be, as before, a range of determinate values,  $e_1 \dots e_n$ , of a determinable (E), and a range of determinate values,  $r_1 \dots r_n$ , of a determinable (R), and a sensory system (S), which has the function of producing instances of certain determinates of R in response to instances of certain determinates of E. Thus, S will have the function to produce  $r_1$  in response to  $e_1$ ,  $r_2$  in response to  $e_2$ ,  $r_3$  in response to  $e_3$ , and so on. Now, in addition, let there be more proximal candidate contents (prox- $e_1 \dots$  prox- $e_n$ ), which are values of the determinable Prox-E. And let ' $r_i$ ', ' $e_i$ ' and 'prox- $e_i$ ' stand for arbitrary values in their respective ranges. Then the addition to IT is this:  $r_i$  refers to  $e_i$ , and not prox- $e_i$ , if S was selected for producing determinates of R in response to determinates of E *by* producing determinates of R in response to determinates of Prox-E, and not vice versa (i.e., S was not selected for producing determinates of R in response to determinates of Prox-E *by* producing determinates of R in response to determinates of E).

Back to the toad. The pathways in the toad's brain were selected for responding to both the distal worm-like motion and the more proximal patterns of light that carry information

about the distal worm-like motion to the toad. But there is an important asymmetry here.<sup>19</sup> These pathways in the toad's visual system were selected for responding to the light by producing certain tectal firings because *by that means* they responded to the distal worm-like motion, and not vice versa. That is, they were not selected for responding to the distal worm-like motion by producing certain tectal firings because *by that means* they responded to the more proximal patterns of light. That just isn't how the means-end analysis pans out. I believe that this solves the problem of distal content.

Millikan's (2001, 119 n7; cf. OCC1, 231 n7) also attributes to Neander (1995) the view that "...things really can't have distal functions," and so perhaps I should say something about this too. Actually, I did not mean to make this claim either. M&M is mostly about how, because things are selected for complex causal roles, there is a functional indeterminacy problem that needs to be distinguished from other indeterminacy problems for teleosemantics,<sup>20</sup> and about how this particular functional indeterminacy problem might be met. Things are selected for complex causal roles, I said, and so we can describe functions in complementary but different ways by focusing on different aspects of these roles. But we can, in a principled way, focus on different aspects of these roles in certain contexts, I said. For instance, it is the function of ovaries not only to produce ova, but also to contribute to conception and to live births. *All* of this is their function. However, ovaries do not *malfunction* if they cannot contribute to conception or a live birth because the fallopian tubes are blocked. Ovaries only *malfunction* if they cannot make their own individual contribution, I said. It was my attempt to characterize this "individual contribution" that led to so much misunderstanding.<sup>21</sup>

Of course, producing ova is not something that ovaries can do all by themselves, without assistance. The ovaries must be fed oxygen by the circulatory system, for instance.

I also spoke of the more "peculiar" function of something as that which it does "more immediately," and this is easily conflated with talk of the "more proximal," in discussions of distal content. I meant these contrasts to be orthogonal, but that wasn't made completely clear. In fact, I assumed that certain pathways in the frog's (or toad's) brain had the more "immediate" function, so to speak, of detecting something small, dark, and moving (or something in worm-like motion) that was in fact distal (though I had no solution to the problem of distal content then to offer). The less "immediate" function of these pathways, as I saw it, was to help the creature catch and eat prey and be better nourished and so on. That is in part to say that the relevant *sensory* pathways do not *malfunction* if the frog (or toad) cannot catch, eat, and benefit from its prey because, say, it cannot swallow, or its tongue-snapping mechanism misfires, or it is missing necessary digestive enzymes. The relevant *sensory* pathways count as *malfunctioning* if they cannot detect the type of stimuli to which it is their function to respond. This was not (and is not) by itself meant to be conclusive on its own for how we should approach content, but it seemed suggestive

<sup>19</sup> Shades of Fodor's asymmetric dependency theory here, but this is not at all like his attempt to solve the problem of distal content.

<sup>20</sup> Such as Fodor's (1990) problem, which he understands to be due, as he puts it, to natural selection being "extensional."

<sup>21</sup> Others besides Millikan were also misled (see Price 1998). Happily, Papineau (1998) understood what I said.

to me at the time. The idea was that the aspect of the complex causal role which is something's function that is most relevant to judgments of malfunction might also, in the case of representational systems, be most relevant to judgments about misrepresentation too.

To sum up. One phenotypic characteristic can be more adaptive than another and can be selected in preference to another because it brings about a certain effect in response to a certain cause. Sensory systems, more specifically, have the function of producing certain inner states in response to certain environmental conditions obtaining. They therefore have the function of producing states that carry natural indicative information if we understand natural indicative information in terms of singular causation, such that one thing carries natural indicative information about another if it is caused by it. We can then entertain the rather obvious suggestion that sensory representations represent that which they are supposed to carry information about. It is wrong to reject such a suggestion on the grounds that functions must be effects. I have tried to show that this idea can be formulated in a respectable way, and that as formulated here it generates suitably determinate content for at least sensory representations.

Hard questions remain. The hardest is how to get from here to where we want to go in the end, which is to much more sophisticated representational contents. Millikan and Papineau believe that their output-oriented approach will succeed, and that this more input-oriented approach will fail. They seem to think that the input-oriented approach is stillborn. That it is not. And I am betting that they are wrong about its future prospects too. It seems to me that they have begun in the wrong place, assuming (as I do) that we want a treatment of sensory representations that is sensitive to the actual capacities of sensory systems and to the actual information processing that is responsible for them. In Millikan's view, and in Papineau's view too, I turn teleosemantics on its head. But I think I've turned it the right way up and now, I hope, IT can get up and walk. I am in sympathy with much of what Millikan (OCCI) says in her important discussion of substance concepts. Much of it is, I think, compatible with the informational teleosemantics supported here. However, this is where I take my departure with the customary parting comment that a careful discussion of this must wait for another time.