

# BIO-SEMANTICS, ECOLOGY AND CONTENT ASCRIPTION

Alejandro Rosas  
Departamento de Filosofía  
Universidad Nacional de Colombia  
arosasl@unal.edu.co

**Resumen:** Los defensores de la bio-semántica proponen que el contenido de las representaciones sea fijado con referencia a su función biológica. Sin embargo, una legítima pluralidad de funciones etiológicas para una representación dada impide fijar el contenido. Los intentos de superar por medio de un argumento contrafáctico esta dificultad en la *atribución de contenido* han sido acusados de tergiversar la distinción *selección para/selección de*. Reconstruyo el argumento contrafáctico en el contexto de los constreñimientos al diseño evolutivo que subyacen a una teoría ecológica de las presiones selectivas. Seguidamente, argumento que estos constreñimientos asignan un rol modesto a la función etiológica en un método para la atribución de contenido. Sin embargo, esto requiere abandonar la tesis de que pensar en términos de una función etiológica resuelve *a priori* la atribución de contenido; en particular, no puede reemplazar la experimentación y la teorización cognitiva. Tal como lo han advertido los ecólogos del comportamiento, las funciones biológicas no deben ser tan fácilmente traducidas a mecanismos próximos, en particular perceptuales.

**Palabras clave:** bio-semántica, contenido, funciones etiológicas, ecología, Neander, representaciones.

**Abstract:** Advocates of bio-semantics propose that the content of representations be fixed with reference to their biological function. A legitimate plurality of etiological functions for a given representation, however, stands in the way of fixing content. Attempts to overcome this problem for *content ascription* with a counterfactual argument have been accused of misusing the *selection for/selection of* distinction. I reconstruct the counterfactual argument in the context of the constraints on evolutionary design that underlie an ecological theory of selection pressures. I then argue that these constraints assign a modest role to etiological function in a method for content ascription. But it requires abandoning the claim that *a priori* thinking in terms of etiological function will solve content ascription; in particular, it cannot replace cognitive theory and experimentation. As behavioural ecologists have warned, biological functions should not be too quickly translated into proximate, in particular perceptual, mechanisms.

**Keywords:** bio-semantics, content, etiological functions, ecology, Neander, representations.


## I. INDETERMINACY: THE CHALLENGE

Theories of content can be understood as answering two different concerns. The first one is a concern about *content ascription* and asks: *what* content (and what *type* of content) does a given representation have? How do we establish this when we attribute representations to other minds? A more fundamental concern is the following: How do we know whether any given mental item has content? In virtue of what does it have content? This paper is concerned about bio-semantics as an answer to the first question. It inquires whether bio-semantics solves the problem of *content-ascription*, not on whether it proposes a plausible (naturalistic) analysis of the



relation “— has the content that —”. Arguably, bio-semantics fails on this second task (Pietroski 1992), but so does any other theory that purports to offer a non-circular analysis, naturalistic or not. Bio-semantics can be construed modestly as an attempt to answer the first question. As such it will be treated in this paper.

Bio-semantics (also known as teleosemantics) was born as a plausible solution to the indeterminacy affecting causal theories of content. For it is a consequence of causal theories that a given representation represents every event that actually causes it, even if it is not what it is *supposed* to represent. Bio-semantics is meant to overcome this difficulty: it proposes to nail down the appropriate content by linking the representational state to its biological function. The classical trial for this approach is the content of a frog’s representational brain-state when it flicks its tongue to catch prey. If its function is to represent those worldly conditions in the presence of which the flick of the tongue will normally lead to the ingestion of food, these conditions are the content of its representation. However, it has turned out very difficult to describe these conditions univocally. Do frogs flick their tongues at flies (McGinn 1989 and Sterelny 1990)? Do they flick their tongues at food (Shapiro 1992 and Millikan 1989)? Do we have to know first their desires, if they have any (Papineau 1998), in order to know what they flick their tongues at? Or do they simply snap at small, black moving things (Neander 1995) or perhaps at black, moving, non-dangerous food (Agar 1993)? The proposals can be multiplied. Bio-semantics seems to end up with an indeterminacy problem of its own.



Advocates of bio-semantics give a cogent explanation for this multiplicity. Traits are adaptive when they enhance the fitness (reproductive success) of an organism, but they do not usually enhance fitness as an immediate effect. They enhance fitness through a causal chain of events, involving several steps until the relevant fitness effect emerges (Neander 1995). Accordingly, the function of a trait receives a different description, depending on the event in the causal chain one chooses to focus on (Goode & Griffiths 1995). The proposals listed above are equally good candidates for the function of the frog’s brain-state because they mention links in a chain of events leading from the frog’s brain-state to its reproductive success. The frog’s brain-state causes a tongue-flick at things that are black, moving and small, thereby catching flies, thereby feeding well, thereby building an efficient and healthy body and thereby increasing its chances of reproductive success. That trait was selected because it did all these things. They are equally entitled to be called its functions (Goode & Griffiths 1995). This explains the plurality of proposals for the biological function of that trait as legitimate. It solves the indeterminacy problem affecting function ascriptions, or at least shows that the indeterminacy is not harmful. However, one is left with various possibilities of content ascription.

Neander suggests that it is possible to choose one among this plurality of genuine functions as specific to a given trait by focusing on the phenomenon of malfunction (Neander 1995). The specific function of a trait is the immediate effect whose absence univocally indicates that the trait is malfunctioning, because no other trait possibly contributes to its absence. This solution may be applied to detector devices as well. In order to find which content detector devices are designed to respond to, Neander’s strategy suggests directly probing, through cognitive experiments, into the device and its detecting states. But if cognitive experiments can solve the question of the content of representations, what is the use of a selectionist approach to the mind advocated by bio-semantics? If, on the other hand, we approach content from the viewpoint of selective



pressures, we will tend to direct attention, *a priori*, to functions located higher up in the causal chain leading to a fitness enhancement. Since these two methods result in different content ascriptions, the indeterminacy problem persists. However, it is my purpose to argue that there is no indeterminacy here. What is required is a methodological decision. Cognitive experiments inquiring into content and theoretical thinking in terms of selective histories must be assigned different, but methodologically complementary roles in the task of content ascription.

### 2. FALLACIOUS COUNTERFACTUALS?

Let us call 'D' the brain-state (type) of the perceptual system that signals the presence of whatever it is that it signals to the neural system that causes the flick of the tongue. Cognitive experiments (Lettvin *et al.* 1959) suggest that 'small, black, moving thing' is a correct description of the content registered by D. Let us assume that this is the case. This is the first link in a chain of events leading ultimately to a fitness enhancement: Detection of a small, black, moving thing triggers the flick of the tongue; the frog thereby catches a fly (in normal environments), thereby feeds, thereby builds or maintains its body, thereby contributes to its reproductive success. In Neander's terms, signalling to some other system in the frog's brain the presence of small, black, moving things is the specific effect for which the frog's brain state is directly responsible. The frog's detector device malfunctions if it fails to instantiate D when a small black moving thing passes through its visual field. Participants in the debate have argued against fixing the content of D through these experimental results. They use the following counterfactual, theoretical argument: if D caused the frog to snap at small, black, moving things, but did not cause it to take food in, D would not have been selected. Therefore, D must signal the presence of food rather than of small, black, moving things, even if 'small, black, moving thing' happens to be a correct description of a stimulus that *causes* brain-state D.

The argument contains a plausible claim: that in absence of a correlation between frog food and small, black and moving things, a brain-state with the latter content would not have been selected. This claim depends on a plausible hypothesis about the biological function of the larger system where the detector device does its work. But as already noted in the literature, the argument uses wrongly the theory of biological functions against the results of cognitive experiments. Consider iterating the argument against the thesis that the brain-state signals food: If a brain-state signalling food caused the frog to feed well, but not to reproduce (because normal frog food in the frog's evolutionary environment happened to produce sterility), then a brain-state representing food would not have been selected. But this would take us all the way to the claim that 'reproduction enhancer' is the only legitimate function and the only legitimate content of such a brain-state; and this is nonsense. Counterfactuals of this sort cannot show that a state representing small, black, moving things could not have been selected. Even if they are true, they do not allow any one event in a stable casual chain to be chosen as more appropriate than any other one to qualify either as the function or the content. The stability of the casual chain leading up to fitness sanctions every link in the chain as a genuine function of the trait.

This use of counterfactuals could easily appear as an abuse of Sober's *selection for/selection of* distinction: it looks as if authors were arguing for the thesis that there has been *selection of* brain-states that detect small, black, moving things, but no *selection for* those states (Goode & Griffiths 1995). Sober's distinction is meant to remind us of the fact that the total set of effects





of adaptive traits is larger than the set of effects *for* which the trait was selected. Hearts pump blood and produce noise, so when hearts are selected, pumping blood and producing noise are both selected. But unlike pumping blood, producing noise is not something *for doing which* a heart has been selected. It is merely a side-effect produced by the pumping mechanism which hearts in fact use (muscle contraction). In contrast to noise production, muscle contraction is something *for doing which* tokens of hearts have been selected. Hearts pump blood *by* muscle contraction but they do not pump blood *by* producing noise.

Similarly, if frogs detect food *by* detecting small, black, moving things, then there has been selection both for detecting food and for detecting small, black, moving things. Detecting small, black, moving things is not a mere side-effect of detecting food, but a capacity through which frogs detect food. Here we have two effects *for doing which* the frog's brain-state has been selected to cause the flick of the tongue. In the environment where this state evolved, it happened that snapping at small, black, moving things was *often enough* snapping at frog food. In the same way, all content ascriptions used by participants in the current debate truly name effects *for doing which* the brain-state in question was selected. They describe in every case events in the chain leading up to fitness. Therefore, no appeal to Sober's distinction will be able to support an *a priori* argument for ascribing content on the basis of any one effect within this chain rather than any other one.

Goode & Griffiths (1995) think that Sterelny (1990), Shapiro (1992), and in a more subtle manner Agar [11] are guilty of supporting their favourite candidate for function and content with this misuse of Sober's distinction. Neander (1995) and Papineau (1998) think this applies to Millikan (1989) as well. I will not try to decide whether or not these authors are guilty of this fallacy. Rather, I will show that the counterfactual argument can be read in a different way, as having an altogether different point in mind.

### 3. BIO-SEMANTICS: THE ORIGINAL INSIGHT

Counterfactual reasoning of the sort employed in this debate makes sense when one takes a particular view of the original insight behind bio-semantics. This insight is not only that etiological functions provide a privileged standpoint for understanding mind, but that a selectionist approach should specifically target those functions which have been the focus of the behavioural ecologist's theory of animal behaviour; namely the famous four F's of biology: feeding, fighting con-specific rivals, fleeing from predators and mating. This way of reading the argument supports the privileged status of one point in the chain leading up to fitness, but not by way of misusing Sober's distinction. As will be explained below, the rationale behind this privileged status is a robust theory about adaptive design. This theory is able to offer guidance to direct experimental inquiry into representational content.

If we view bio-semantics as defending not only etiological functions, but more specifically the classification of behaviours and their evolutionary functions that is current in behavioural ecology, the internal quarrel among advocates of bio-semantics can be construed as a quarrel about the significance of the four F's of biology for the an evolutionary approach to mind and cognition. On one side, we find those authors who have defended some variant of the thesis that the content of the frog's brain-state is or includes food. I am considering the fly-thesis as





a variant of the food-thesis: flies or organic substances in general are just instances of *edible* prey. Save for Neander, practically every participant in the debate defends this claim. Neander (1995) called this camp 'high church' teleosemantics ('standard teleosemantics' in Neander 1996). Standard teleosemantics appeals, not to the immediate effect – which is the correct one for talk of malfunction to make sense – but to a higher-level effect in the chain of events leading up to fitness. I will label this camp 'ecological bio-semantics'. Authors in this camp assume, explicitly or not, that the functions which are going to help us get a grip on the nature of mind are precisely those which behavioural ecologists find illuminating when studying animal behaviour. According to behavioural ecology, feeding is one of the basic functions in terms of which animals are usually designed.

Neander's position represents the other camp (she has argued that Dretske belongs here too: see Neander 1996). Even though she thinks that selection pressures explain mental content, she seems committed to deny that a universal theory of selection pressures, as commonly defended by behavioural ecologists, will help us decide between the several possibilities of content ascription. Direct cognitive investigation into content will be more appropriate to the task.

What is at stake is the plausibility of an appeal to the biological function of larger systems embedding cognitive devices to guide content ascription. Ecological bio-semanticians believe that behavioural ecology and its theory of basic selection pressures should provide a privileged access to mind, its structure and its representational content. One consequence is that the frog's brain-state that causes the flick of the tongue is bound to represent food, since feeding is one of the basic functions in terms of which natural selection designs organisms. However, the behavioural ecologist's theory of functions will illuminate representational content only if particular conditions obtain. In the following, I will try to spell out these conditions. In any case, the upshot will be that ecological thinking cannot be used to build *aprioristic* and 'armchair' arguments, but as a methodological guide to be used in combination with cognitive experiments.



#### 4. ADAPTIVE CONSTRAINTS ON EVOLUTIONARY DESIGN

Under adaptive constraints I understand capacities that organisms *must* have if they are to be adaptively designed at all. Without intending to be exhaustive, I will mention some that appear to be non-controversial. Some people may view their non-controversial character to follow from the fact that they are true by definition. I will not take any sides on whether they are true by definition or rather because of the theory we currently hold. My argument will hinge on a difference in status between these *necessary* capacities and others that are clearly contingent. Among the capacities that are *necessary* in this sense, the following can be mentioned:

1. Organisms must be able to reproduce, or to help relatives reproduce.<sup>1</sup>
2. Organisms must be able to take in energy from the environment, that is, to feed.
3. Feeding includes, necessarily, the ability to detect relevant substances and to carry out activities for incorporating them into the organism.

<sup>1</sup> If we would count conceivable immortal beings as life forms, we would have to re-define biological fitness, but such theoretical possibilities are not relevant in our current biology.




Notice that all these capacities are general and may be realized by a variety of different mechanisms with a variety of specific operations. Behavioural ecology focuses initially at this general level of description. Therefore, this level of description can appropriately be called the ecological level of description of the capacities of an organism.

Biology is also interested in finding the specific mechanisms that realize those capacities in particular species:

4. Particular mechanisms for detecting and taking food in, even though they all respond to this general description, may be very variable in the specific operations by which they detect and take in food.<sup>2</sup>

Statements 1-3 express rough truths about living organisms; they represent capacities that organisms *must* have if they are to be adaptively designed. These fall under the four F's of biology. Assuming that they have been designed by natural selection, we can conveniently call these capacities the *ecological-level* effects of biological design. In contrast, capacities mentioned in statement 4 – bits of mechanisms or their immediate operations, by which a living being in a given environment detects and takes in food – allow of many different versions. Their reliability is highly sensitive to changes in environmental conditions. Assuming they have been designed by natural selection, we can conveniently call them the *mechanism-level* effects of biological design.



Both ecological-level and mechanism-level effects are selected effects, things for doing which selection rewards a given design. However, the difference is important. The frog's brain-state D (type, not token) that mediates between a distal food-stimulus and the flick of the tongue is a mechanism-level effect. There is a point in saying that unless D causes the frog to snap at food often enough, there will be no selection for D, but selection for some other state, or else frogs constructed with D will go. But there is no comparable point in saying that unless feeding leads to fitness, no selection for feeding will take place, but rather selection for something else. The difference between both statements results from the difference between ecological-level and mechanism-level effects. Natural selection has practically fixed the link among ecological-level effects: the link that leads from feeding well to reproduction and to fitness. Capacities belonging to this level are as good as fixed requirements for building adapted organisms. This is the reason why they have been included among the four F's of biology.

Natural selection does not in the same way embrace any capacity belonging to the mechanism-level. The effects of biological design at the level of mechanisms are highly sensitive to environmental change. Natural selection will start with rough designs that will do their job only across a limited range of environmental variation. The reliability with which effects at this level will enhance fitness depends strongly on environmental stability. It is not difficult to think of variations in environmental conditions that would make a particular way of identifying prey obsolete. In contrast, it is practically impossible to think of variations such that, leaving the reliability of capacities at the mechanism-level untouched, they would nevertheless disrupt the path that leads from level 3 through 1 and to fitness.

<sup>2</sup> In particular, as we will see, food detectors that mediate between a given stimulus and bodily movements may use representations, but they do not have to.



## 5. THE POINT IN COUNTERFACTUAL ARGUMENTS

Let us now see how these facts about biological design could motivate counterfactual thinking in the context of the indeterminacy of function and content. We want to find out what would be the point in saying that there would have been no selection for brain-state D (type) representing small, black, moving things if snapping at food had not been one of its effects. Why should this counterfactual be less trivial than one stating that unless feeding well enhanced the ability to reproduce, then a brain-state D representing food would not have been of any help either?

I think the clue lies in the biological facts stated in section 4. A counterfactual conditioning an effect at the mechanism level to an ecological-level effect has a point that is absent in a counterfactual conditioning one ecological-level effect to another ecological-level effect. In the causal process that controls the evolution of any given trait, natural selection is counting on a stable causal connection between the ecological-level effects and fitness: feeding well, fleeing from predators and winning fights against con-specific rivals in mating contests are abilities that contribute to successful reproduction. The chain leading up to fitness will rarely break up in the causal links connecting those abilities to reproductive success. In contrast, the connection between the mechanism-level effects and those abilities is fragile. The selection of effects at the mechanism level depends on how reliably they lead up to effects at the ecological level. Their reliability in this sense will depend on how well they perform on the face of environmental change. Since there is no guarantee that their performance will be constant under change, the chain leading up to fitness will often break at the mechanism level.

Natural selection will work to increase the fitness of traits by increasing the reliability of mechanism-level effects in achieving the ecological-level ones. How often natural selection re-designs mechanisms, will depend on how often the environment changes, and within which ranges of variation; on how challenging these changes are; and on whether natural selection finds solutions to them. In view of these historical contingencies, a counterfactual statement that conditions selection for a mechanism-level effect on the achievement of an ecological-level one, may be trying to make one of two points. Let us look at them in the particular example of the frog's detecting device.

### 5.1 *Food weakly shapes state D*

Suppose the ancestors of frogs initially detected food through brain-state D and that the content of this state was far from frog-food or from whatever organic substance occupied that role. It happened to be easy for natural selection to construct a state with this content. By happy chance, this content was reliable enough in the ancestral environment of frogs to make them tongue-flick at suitable food. Suppose that the frog's environment did not relevantly change over time to challenge the capability of this brain-state in accomplishing its function. Under these historical conditions, brain-state D was probably a quick and dirty way of designing the frog's food detecting abilities, and endured due to the convenient, but lucky and perhaps atypical, diachronic uniformity in the frog's environment. If one then says that no selection for D would have happened if D had not (luckily) caused the frog to snap at food, one may be suggesting that





it is not due to the special virtues of D, but to lucky correlation, that D performs its expected job of enabling the frog to feed well.

In connection with the above, if D registers the presence of a stimulus affecting one sense only and, on that basis, inflexibly triggers the flick of the tongue to catch prey, the counterfactual statement could be suggesting a difference between two capacities, namely *detecting* and *representing*. Though the frog's brain-state detects and registers a stimulus, it does not, strictly speaking, *represent*. The linguistic stipulation introduced here is not meant as an analysis of our concept of representation. It points to a plausible distinction between two capacities, however we call them. For one capacity, call it *representing*, it is a necessary condition, though not a sufficient one, that its tokens indicate the presence of environmental features by weighing inputs coming from different perceptual sources. When those features are essential to achieve a goal that is important for survival (like feeding), *representations* underwrite the ability to transfer a given behaviour (e.g. the tongue-flick) across different perceptual situations.<sup>3</sup> The counterfactual would then be used to suggest that D does not have the special virtues we attribute to *representations*.

## 5.2 Food strongly shapes state D

Alternatively, if it is empirically implausible to assume a diachronic uniformity in the environment of a given organism, the same counterfactual could be used to express a rather different point. If historical and/or comparative evidence suggests that the environment of ancestral frogs, or of the ancestor-species of modern frogs, changed often and within a wide range of conditions, the suggestion would be that selection for a brain-state D as a *quick and dirty* way of registering through a simple proximal stimulus (sense quality) the presence of a distal one (food) is unlikely. The expectation is that environmental changes will have forced D to adjust to these changes to keep a good record at detecting food across a relatively wide range of conditions. Adjustment to changes could have forced D to track food across different sense data and their combinations. Suppose that at the beginning of their natural history, frogs (or their ancestors) tracked food based on smell only, and did not need to take account of size because everything that smelled like prey had a suitable size. A simple environmental change, i.e. bigger-than-frogs prey-smelling animals, would force frogs to track their food based on smell and size. Such a challenge would force brain-state D to evolve towards a *representational* state, one which indicates the presence of that at which to tongue-flick on the basis of data coming from different perceptual sources.

## 6. FUNCTIONS AND NATURAL HISTORY

The preceding considerations allow for useful counterfactual reasoning in the context of ecological bio-semantics. Counterfactual reasoning gives ecological-level effects a special role in the causal process that decides the shape and fitness value of a trait. This can be justified by a plausible set of assumptions concerning the causal connections that natural selection counts on

<sup>3</sup> I follow a suggestion by Kim Sterelny (2000) who attributes representations to animals when they multitrack an environmental feature in a way that enables them to transfer a behavioural response across different perceptual situations to achieve a goal. However, I am not committed to the view that these *representations* support the context-free recombinant skills of full-fledged conceptual abilities. Cf. Susan Hurley's analysis and interpretation of experimental work to probe the conceptual skills of animals (Hurley 2003).



when designing organisms. However, the privileged status of facts concerning e.g. the frog's nutritional state does not automatically imply that they have robustly shaped brain-state D. It is still possible that D simply registers a proximal stimulus that picks out food in virtue of a superficial but stable correlation between food and some sense quality. Frogs may have been around for quite a long time now, but then no one can be sure *a priori* that this is partly because the frog's food-detector adjusted to repeated environmental challenges concerning food-detection. Perhaps those challenges never took place and perhaps frogs would have disappeared if they had. If the environment of ancestral frogs had in fact changed in the relevant ways to pose ever new challenges, then *that* would be one argument in favour of the thesis that the food-detector must have moved closer and closer to *representing*, rather than simply detecting, frog-food.

Insistence on data coming from natural history reflects our present lack of theory concerning the evolution of cognition. For example, does the bare existence of different sense organs in an organism imply that it is already beyond brain-states that merely register one-dimensional sense stimuli and are able to *represent* things crucial to survival, i.e., that they multi-track those things by weighing evidences coming from different sources? This is possible, but not certain on our present knowledge. Consider for example an early frog-food detector that started with a brain-state D that tracked visual sense-data coming typically from flies. The data registered by the detector were such that  $\Pr(D \mid \text{flies present})$  was high, but the  $\Pr(\text{flies present} \mid D)$  was low, by giving many false positives. The detector was not able to discriminate between flies and fly-looking creatures, given the stimuli that it picked out. Suppose false positives later turned out to be noxious. The detector responded by changing its data registers to increase  $\Pr(\text{flies present} \mid D)$ , i.e. to increase the probability that the content indicated flies, and flies only, by tracking, for example, flight patterns with certain characteristics. But alternatively, the detector could have coped with those changes by switching to other type of prey, for example worms, by changing its data-registers so that they would register some other movement pattern. Since, after any of these changes, the detector still tracks food or flies on the basis of one sense only (vision) one should not say that D *represents* flies or food. As a necessary condition for *representing* flies or food, D should have gone beyond registering sense data on one dimension, by at least combining data coming from different senses.

Alternatively, the frog's detector could have responded to environmental change by using *representations*, combining inputs from different sense-organs. For example, the frog's detector could acquire something like a digestive-system-directed *representation* for certain types of nutrients that it would identify following different types of data (visual, gustatory, olfactory and feedback from sensations in the digestive system). Since it is likely that sensations of hunger originating in the digestive system are responsible for putting the food detector in alert mode, it is also plausible that the food detector will receive feedback from positive or negative digestive sensations after ingestion. Since they plausibly shape the frog's detector, they are relevant to solve the problem of content ascription.

Now, considering the fact that these and other possibilities are open for the historical evolution of the frog's detector and its representational capacities, it seems to me that some participants in the debate are mistaken when they proceed as if *a priori* thinking in terms of function will solve the problem by leading to one univocal result.





## 7. NEANDER AND SELECTION PRESSURES

The selection pressures invoked by behavioural ecologists suggest that a cognitive mechanism linked to feeding behaviour should have food as its content. But as argued above, this will depend on the type of history undergone by the cognitive device. Contrary to what ecological bio-semantics assumes, historical details could have favoured a device based on one-dimensional data, luckily correlated with frog-food. I believe this possibility is accounted for in Neander's position and in her conviction that a bio-semantic approach will not necessarily tie psychological science to an ecological classification of behaviours and their functions (1995). This is a deviation from standard bio-semantics, but this deviation is not a sign of mistaken reasoning. It is a possibility within selectionist views on mind: it follows from the fact, explained above, that ecological-level effects may only weakly shape the content of representational devices.

The crucial difference between Neander and ecological bio-semantics is the stance she is adopting towards the consequences of admitting a selectionist approach to mind. On her view, ecological selection pressures on the foraging system of an organism give no specific clues regarding the cognitive devices in the system. Even if selection pressures explain the existence of particular representations, we do not need, or cannot hope, to determine what they represent by thinking about what those pressures have been. We can and must do it independently as cognitive psychologists, by looking into the mechanism with the appropriate experiments.

However, Neander's position seems to decree *a priori* that a theory of ecological functions will not be helpful in determining content. Her hypothesis about the content of the frog's detecting state – 'small, black moving thing' – relies too heavily on the experiments reported by Lettvin *et al.* (1959). This classic paper reports experimental evidence regarding the visual cues that the frog's brain-state responds to when it flicks its tongue. But the experiments they report were not designed to prove that the frog tongue-flicks in response to visual cues exclusively. Later experimentation has shown that this is not the case. Alcock (2001) reports experiments showing that toads (and tree frogs) learn to discriminate inedible insects after one single experience: they will not later flick their tongue at those insects. This means that the detector associates visual with gustatory cues and does not respond to visual cues alone. This is what one would expect when allowing for the possibility that ecological-level effects (catching edible prey) have contributed robustly to shape food-detectors. Neander seems to ignore this possibility altogether. Surely, more complexity would be required to say that frogs represent food, but their detecting devices are in fact closer to what we would call representing food than if they responded to nothing else than to visual cues.

## 8. CONTENT AND SELECTION PRESSURES: CONCLUSION

The relation between cognitive investigation into representational content and the selectional approach to content is a particular case of the more general problem of relating biological functions (those handled in behavioural ecology) to proximate mechanisms. Ethologists and behavioural ecologists usually stress that biological functions should not be too quickly translated into proximate, in particular perceptual, mechanisms (Krebs and Davies 1997). Biological functions can and should serve as a guide to the investigation of proximate mechanisms, but both should





not be confused. For ecological bio-semantics, this amounts to a warning not to infer content too quickly from ecological function. One cannot infer a priori that detector devices represent precisely what the behavioural theorists view as their ecological-level effects. Inferences of that sort should be carefully supported with empirical evidence, either historical-comparative or cognitive-psychological. On the other hand, since ecological-level effects are ultimately responsible for shaping mechanism-level effects, cognitive scientists are well advised to accept guidance from ecological thinking. A fruitful interaction between an evolutionary and a cognitive approach should allow theory about ecological selection pressures to guide hypotheses concerning such matters as representational content or psychological architecture; and inversely, ecological theory should resort to cognitive experimentation on representational capacities in order to test those hypotheses. For the specific case of the frog's food-detector, thinking in ecological terms suggests that visual data will not be the only sense-data contained in the representational states linked to its foraging behaviour. And this suggestion has turned out to be true.

## REFERENCES

AGAR, N.

(1993) What do frogs really believe? *Australasian Journal of Philosophy* 71. 1-12.

ALCOCK, J.

(1975, 2001) *Animal Behavior. An Evolutionary Approach*. Sunderland: Sinauer.

FODOR, J.

(1990) *A Theory of Content and Other Essays*. Cambridge: MIT Press.

GOODE, R. AND P. E. GRIFFITHS.

(1995) The misuse of Sober's selection for/selection of distinction. *Biology & Philosophy* 10. 99-108.

HINDE, R.

(1982) *Ethology Its Nature and Relation to Other Sciences*. Oxford: Oxford University Press.

HURLEY, S.

(2003) Animal action in the space of reasons. *Mind and Language* 18 (3). 231-257.

KREBS, J. R. AND N. B. DAVIES (EDS.).

(1997) Introduction. In *Behavioural Ecology And Evolutionary Perspective*, 4th ed. Oxford: Blackwell Science.

LETTVIN, J. Y. ET AL.

(1959) What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers* 47. 1940-1951.

MCGINN, C.

(1989) *Mental Content*. Oxford: Blackwell.

MILLIKAN, R.

(1989) Bio-semantics. *The Journal Philosophy* LXXXVI, 6. 281-297





NEANDER, K.

(1995) Misrepresenting & Malfunctioning. *Philosophical Studies* 79. 109-141.

(1996) Dretske's Innate Modesty. *Australasian Journal of Philosophy* 74. 258-274.

PAPINEAU, D.

(1998) Teleosemantics and Indeterminacy. *Australasian Journal of Philosophy* 76. 1-14.

PIETROSKI, P.

(1992) Intentionality and Teleological Error. *Pacific Philosophical Quarterly* 73. 267-82

SHAPIRO, L.

(1992) Darwin and Disjunction: Foraging Theory and Univocal Assignments of Content. *Philosophy of Science Association* 1. 469-480.

STERELNY, K.

(1990) *The Representational Theory of Mind. An Introduction*. Oxford: Basil Blackwell.

(2000) "Primate Worlds". In *The Evolution of Cognition*. Heyes, C. and L. Huber (eds.). Cambridge: MIT Press.

WRIGHT, L.

(1976) Functions. *The Philosophical Review* 82. 139-168.

