Phenomenology and the Cognitive Sciences
https://doi.org/10.1007/s11097-025-10068-5

Pluralistic Teleosemantics: Why we need both Bickhard-Representations and Millikan-Representations

Lucas Thorpe[1]

Accepted: 25 March 2025

Abstract
Ruth Millikan and Mark Bickhard both offer theories of representation that can be understood as broadly teleosemantic. Both agree that representations have an essentially normative character and that their normativity should be understood by appealing to some biological notion of function. Their fundamental difference has to do with their accounts of biological function. Millikan offers an etiological account of function, according to which the function of a thing is to be understood in terms of what is has been designed to do, or been selected for. According to such an account, whether or not a thing has a particular function depends on the history of the thing. Bickhard, in contrast, argues that we should understand function in non-etiological terms. Instead of appealing to design or selection, Bickhard appeals to the notion of dynamic presuppositions. Living beings are complex systems that involve multiple interconnected processes and there are dynamic presuppositions between the various processes. So according to such an account the reason why the function of the heart is to pump blood, is not because it was selected for this task, but that if it stopped pumping blood the whole organism would die, including the heart itself. I will argue that both Bickhard and Millikan are substantially correct in their positive theses. Rather than being incompatible theories that are attempting to explain the same thing, they offer distinct notions of function, biological normativity and representation that both have their place and provide for distinct, but equally valid and non-competing types of explanation.

Keywords Representational content · Functions · Interactivism · Singularism · Decriptivism

✉ Lucas Thorpe
lucas.thorpe@bogazici.edu.tr

[1] Philosophy Department and Cognitive Science Program, Boğaziçi University, Istanbul, Turkey

Published online: 11 April 2025

Ruth Millikan and Mark Bickhard both offer theories of representation that can be understood as broadly teleosemantic. Both agree that representations have an essentially normative character and that their normativity should be understood by appealing to some biological notion of function. Their fundamental difference has to do with their accounts of biological function. Millikan offers an etiological account of function, according to which the function of a thing is to be understood in terms of what it has been designed to do, or been selected for. According to such an account, whether or not a thing has a particular function depends on its history. Bickhard, in contrast, argues that we should understand function in non-etiological terms. Instead of appealing to design or selection, Bickhard appeals to the notion of dynamic presuppositions. Living beings are complex systems that involve multiple interconnected processes and there are dynamic presuppositions between the various processes. So according to such an account the reason why the function of the heart is to pump blood, is not because it was selected for this task, but that if it stopped pumping blood the whole organism would die, including the heart itself. I will argue that both Bickhard and Millikan are substantially correct in their positive theses. Rather than being incompatible theories that are attempting to explain the same thing, they offer distinct notions of function, biological normativity and representation that both have their place and provide for distinct, but equally valid and non-competing types of explanation.

If we think that representational content is to be explained in terms of function, and we are pluralists about biological functions, then we should be pluralists about representations, thinking that there are at least as many types of representation as there are types of biological function.[1] A pluralist about biological functions should think that there is space for both Bickhard-representations and Millikan-representations. One reason to be a pluralist is that you might think that the different types of representation are found in different types of organisms. So, for example, one might think that the behaviour of bacteria is best explained in terms of Bickhard-representations, whereas human representations of individuals and kinds are best explained in terms of Millikan-representations.[2] But one might also be a pluralist about individual phenomena, thinking that a particular phenomenon can be explained in distinct ways. So, for example, one may think that even in the case of a single phenomenon, such as a frog's capacity to detect and capture flies, we might want to explain this phenomenon using both Bickhard-representations and Millikan-representations,

[1] Since the 1990 s there has been a growing trend towards pluralism about biological functions in the philosophy of biology. In particular philosopher of biology have suggested that we can distinguish between a notion of function understood in etiological terms and a notion of function understood in terms of causal role, and that within biology both notions of function play an important role. For defenses of pluralism about functions in biology see, for example, Garson (2016) and Cusimano & Sterner (2019).

[2] I will explain this distinction between what I call Millikan-representations and Bickhard-representations later in this paper. Bickhard (2017, p.187) distinguishes between his own "pragmatically future-orientated" account of representation and compares this to what he takes to be the dominant "passively past-oriented... spectator model" of representation.

as these two distinct modes of explanation may have different and complimentary explanatory virtues. I will argue that we should be a pluralist in both of these senses.

1 Bickard-representations

Mark Bickhard's interactivist model of representations is based upon the idea that living beings are far from thermodynamic equilibrium recursively self-maintaining systems. A useful label for any such system is an *autonomous* system. This provides us with an understanding of biological normativity independent of any considerations having to do with evolutionary history or selection. For with any such system we have a basic notion of serving a function in terms of contributing to the stability, and hence the continued existence, of the system.[3]

Bickard's canonical example of a far from thermodynamic equilibrium self-maintaining process is a candle flame. A flame, among other things induces convection which sucks in a fresh supply of oxygen and removes waste, and melts wax in the candle which can then serve as fuel to keep the flame going. With a suitable supply of oxygen and wax, the flame will maintain itself and continue in its existence. What distinguishes living systems from candle flames is that living systems are not just *self-maintenant*, but are *recursively self-maintenant*; "they maintain their self-maintenance" (Bickhard, 2009, p.555). Such systems are able to maintain their self-maintenance as the local environment changes by modifying their activities. Bickhard's canonical example of a recursively self-maintaining system is a bacterium that swims so long as it is heading up a sugar gradient, but tumbles if it is heading down such a gradient. (Bickhard, 2009, p.555. See also Campbell, 1974, 1990).

Appealing to far from equilibrium systems allows us to offer a non-etiological account of biological function. The first step is to offer an account of what it is to *serve* a function, and then we can offer an account of what it is to *have* a function, in terms of being a functional presupposition of a process that *serves* a function. A process *serves* a function in an autonomous system in so far as it contributes to the maintenance of the system. That a particular process serves its function, however, presupposes that other processes succeed in serving their function. This allows us to offer an account of what it is for a part of an autonomous system to *have* a function: a part of an autonomous system *has* a function if serving its function is functionally presupposed in order for other processes to *serve* their functions. So, for example, your heart helps keep you alive by pumping blood around your body (and in so doing it helps maintain the conditions for its own continued existence),

[3] Bickhard himself is committed to process metaphysics. However, even if one does not buy into a view that is committed to processes all the way down, we have very good reasons to think that living beings should be thought of in process terms in the way that the Bickhard advocates. See Nicholson and Dupré (2018). So, the plausibility of such an account of biological normativity is independent of a general commitment to process metaphysics. One does not need to make this general metahpysical commitment to processes all the way down to buy into Bickhard's account of representational content. As long as one is sympathetic to a process view of biological entities then one can buy into interactivism and remain neutral about process metaphysics in general.

but in order for the heart to *serve* this function, the lungs need to be bringing fresh oxygen into the blood and be removing excess carbon-dioxide. The lungs, then, *have* the function of bringing fresh oxygen into the blood, because this is functionally presupposed by the heart (and the rest of the circulatory system) *serving* its function (Bickhard, 2004, p.78; See also: Bickhard, 2017, p.183). Appealing to such dynamic functional presuppositions offers us a way of explaining how a particular sub-process in an autonomous system can be mis-functioning that does not appeal to any notions of selection or design, and so is independent of any considerations concerning the history of the organism. To understand the function of a particular part of such a system, one just needs to know about its current ongoing dynamics. This seems a perfectly respectable notion of function.

This notion of functional presupposition underlies the interactivist account of representational content, for certain process don't just functionally presuppose other processes taking place within the organism, but also presuppose that the world is a certain way. The process of walking, for example, functionally presupposes that the ground carries on ahead of me.[4] This allows the interactivist to provide an account of what Bickhard calls minimal representing. Living autonomous systems need to continuously interact with their environments in order to survive. And, Bickhard explains,

> Engaging in such interactions is functionally anticipatory that they will in fact succeed in contributing to the maintenance of the autonomous system. That is, the success or lack thereof, and the conditions upon which that might depend, are in the future of the interaction once initiated. As such, there is an implicit predication that "this" is one of those environments in which the initiated interaction will proceed as anticipated. That predication, therefore, might itself be true or false: the environment might or might not be among the supportive kinds. Under some conditions, the activity will be functional, while under others it will not. Initiating the activity, therefore, presupposes that those supportive conditions hold. (Bickhard, 2009, p.569)

On the basis of this account, Bickhard builds up more complex representations in terms of what he calls webs of interactive potentialities. And he offers a Piaget-inspired account of how representations of small manipulatable objects can emerge.[5]

[4] As we shall see in section three this way of characterizing what is presupposed by walking turns out to be inadequate, which raises problems for Bickhard's account.

[5] "[I]nteraction indications can both be multiple they can branch and they can iterate. As such, they can form webs of interconnected conditionalized indications of interaction potentiality m perhaps vast and complex webs. Some subwebs of such a larger web may come to have special properties. In particular, they may be internally reachable, in the sense that any indicated interaction anywhere in the subweb is reachable as a direct interaction potentiality, perhaps via various intermediary conditional interactions, and that intemal reachability property may remain invariant under some relevant class of other kinds of interactions. For example, a child's toy block will afford multiple potentialities of visual scans and manipulations. Any one of these potentialities is available from any other... so the subweb of interactive potentialities for this block is internally reachable... This outlines the general manner in which the interactive model can scale from simple interaction possibility representations to representations of physical objects. It is a generally Piagetian, or pragmatic, model of object representation..." (Bickhard, 2004, pp. 79–80).

Now, I think that this model offers a plausible account of a certain type of representation, which can function within certain types of explanation. And Bickhard representations, understood as being the implicit presuppositions of anticipatory processes, allow for organism detectable error, which allows for a plausible and needed account of learning. For the organism recognizes the presupposition has not been met, when the anticipations are not fulfilled. Bickhard argues that the possibility of accounting for system detectable representational error is a central desideratum of any theory of representation (Bickhard, 2004, pp. 79–80). And he argues that etiological theories of representational content cannot account for system detectable error. However, while it is clear that some representational errors need to be system detectable to explain, for example, how certain types of learning are possible, it seems plausible to assume that there are going to be classes of representational errors that are not system detectable – but which may be detectable from the standpoint of an external observer. (I will return to this claim later). But there is no reason in principle why we cannot add etiologically specified representations on top of an interactivist framework.[6]

Now Bickhard sometimes argues that all representations are such that they must allow for organism detectable error. If this were the case, Millikan style representations would not really count as representations at all. Thus, he argues, in the context of discussing theories like Millikan's that appeal to evolutionary and learning histories, that,

> [i]t requires an *external observer*, and a fairly knowledgeable such observer, to assess such issues as asymmetric dependencies or evolutionary or learning histories. The organism cannot do it for itself. So the purported solutions work, if they do at all, only for observers, not for the intentional organisms or systems themselves.//This is unacceptable. It violates a basic constraint of naturalism: representation, and representational error cannot be dependent on external analysis and ascription on pain of making the analysis of intentionality circularly dependent on external intentional observers... Representational error must be possible, and it must be detectable - not necessarily infallibly - in the system, by the system, and for the system, independent of any observer. (Bickhard, 1998, p.188)

But this criticism misrepresents the role of external observation in Millikan's account. The representational content of Millikan representations depend upon the evolutionary or developmental history of an organism, and such history can perhaps be recognised by an external observer. But this does not make such representational content "dependent upon external analysis and ascription"; to claim that something

[6] The vehicles of such content could as Shea (2018, 40) suggests be particular dynamic parameters such as frequency, phase, impedance and gain. And we can think of algorithmic operations on such vehicles in terms of dependence relations between such parameters and offer historical accounts (either in terms of learning of evolution) as to how such parameters have the function of representing certain environmental features. Such explanations could be used not to explain a particular dynamic flow, but to explain certain regularities in patterns of activity: for example how the organism is robustly able to succeed in certain tasks but fail, in systematic ways, in others. But it is not clear that Shea's potential response is compatible with Bickhard's thought that representing is something that we do, not a structure that we use. See Erdin and Bickhard (2018).

is *observable* does not entail that its existence is dependent on being observed. The fact that the proper function of an evolved mechanism may perhaps be reliably attributable by an external observer (but not by the system in which the mechanism exists), does not entail that the proper function is observation dependent. So, the explanation of representational contents in terms of proper functions that can only be attributed by external observers does not make the representations observer dependent. Having a proper function is observer independent; it is only the attribution of proper functions that is observer dependent. Now, as we shall see, Millikan's account depends upon the coherence of adaptationist accounts of proper function, and such accounts have their own internal problems as famously argued by Gould and Lewontin (1979).[7] But this is a distinct issue from the question of whether *all* misrepresentation must be system detectable. And I can see no good reason to believe, as Bickhard seems to, that the fact that a misrepresentation is only (potentially) detectable by an external observer entails that the representation is observer dependent.

2 Millikan and Teleo-semantics

In a recent paper, Ruth Millikan explains that the core commitment of teolosemantics is the idea that representations are essentially normative and that this normativity is to be explained by appealing to functions. Thus representations can be "correct or incorrect, true or false, fulfilled or unfulfilled, satisfied or unsatisfied" (Millikan, 2020, p.2) and the normative status of a representation depends upon whether or not it is doing what it is supposed to do. Thus, Millikan quotes her 1984 book *Language, Thought and other Biological Categories*:

> If language device tokens and mental intentional states (believing that, intending to, hoping that) are members of …function …categories, then they are language devices or intentional states not by virtue of their powers but by virtue of what they are supposed to be able to do yet perhaps cannot do. For example, just as hearts and kidneys are sometimes diseased or malformed, so sentences and beliefs are sometimes false, and words and concepts are sometimes ambiguous and sometimes vacuous. Such sentences, beliefs, words, and concepts are not able to perform their… functions. (Millikan, 1984, p17)

And she continues that "this is all there is to the claim that was later dubbed "teleosemantics". Teleosemantics says nothing about what the function of a representation is or about how a representation carries out that function. Despite widespread confusion on this matter, it says nothing about how the "content" of a representation is determined. It leaves almost all of the work undone. Also, it is, or was originally, a claim about representations generally, not merely mental representations." (Millikan, 2020, p.3).

[7] There is a huge literature on adaptationist explanations and a discussion of this issue is beyond the scope of this paper. But.

Understood in such broad terms, Bickhard should also be understood as advocating a version of teleosemantics because he also believes that the correctness conditions of representations need to be explained in functional terms. Now, Millikan (2020) suggests that any proponents of teleosemantics thus broadly understood need to answer three questions. And how one answers these questions will lead to a specific version of teleosemantics:

> (1) What is it to have a function in the appropriate sense?
> (2a) What is the specific function of a representation as such, the function required for it to be a representation? What is its job?
> (2b) What is the particular method, if any, by which a good undamaged representation performs its function? (p.3)

The fundamental break between Millikan and Bickhard has to do with how they answer the first question.[8] Millikan, like most contemporary defenders of teleosemantics offers an etiological account of what it is to have a function, with the function of a thing being what it has been designed for or selected for. Bickhard, in contrast, takes the fundamental notion of a process serving a function to be one of contributing to the stability of the far from thermodynamic equilibrium process of which it is a part. As such, this notion of serving a function is not dependent upon the past history of the system, but depends purely on its current (and future) dynamics.

There is no reason why a proponent of teleosemantics should think that there need be a single answer to the first question. And if one thinks that there are distinct ways of answering this first question, then one should think there will be distinct species of representation[9]; given that there are at least two ways in which something can have or serve a function, either of them can potentially underlie the normativity of particular types of representations. Bickhard offers a nice non-etiological dynamic account of how bacteria can represent and misrepresent their environments. Etiological accounts, such as that proposed by Millikan, offer plausible accounts of how maps and other artefacts represent or misrepresent that appeal to what they are designed to do. Thus, I suggest that a proponent of teleosemantics should be a

[8] In terms of her answer to the second question. Millikan (2020) argues that something counts as a representation only if it is used as a representation. And she argues that this suggests the necessity of some sort of interpreter (or consumer) that is able to use the representation appropriately. Thus Millikan (2020) argues that "[e]very kind of representation, including a neural representation, should there be such, is for guiding the behavior or operation of another mechanism or other mechanisms, "interpreters" or "consumers," having functions that can be performed successfully only if their activities are set or molded to fit or take account of the state of affairs represented. Representations are used by their interpreters for guidance in performing functions of the interpreter." (p.4) Now, although Bickhard would not use such language, his account, with his distinction between what he calls contact and content, is committed to making an analogous move, although expressed in dynamic rather than mechanistic terms. In terms of Millikan's third question, most contemporary accounts appeal to mechanistic and computational accounts of how a particular representation performs its function. Bickhard, in contrast, suggests that we should answer this how question by appealing to dynamic process-based explanations.

[9] The position I am defending here is similar in spirit to the pluralism defended by Nicholas Shea (2018).

pluralist about the function of representations. And such pluralism leaves us with the interesting question when it comes to any particular purported type of representation, whether the functions that determine the correctness conditions are best explained in etiological or dynamic terms, (or perhaps whether a particular representational system can be explained in both ways but for different explanatory purposes). That is for any purported explanation of a system that appeals to representations, we can ask whether it is best understood in terms of a Millikan-Representation or a Bickhard-Representation.

3 Why Bickhard-representations are not sufficient

In this section I will focus on four problem cases, where Bickard's model seems to be inadequately capture the situation. I name these cases: (1) Magoo-cases, (2) Inaccurate-frog cases, (3) Clubbed-frog cases, and (4) Not-getting-mauled cases. I am not claiming that Bickhard's explanation of these cases is wrong. But that explanation that appeal to Millikan-representations can capture something about these cases that cannot be captured by explanations in terms of Bickhard-representations.

Magoo cases Explanations in terms of Bickhard-representations are going to be such that Mr Magoo is going to count as correctly representing his environment. Mr Magoo is a classic cartoon character, who first appeared in 1949, who is extremely short sighted. He wanders through the world, always on the brink of falling off a cliff edge or off the edge of a building site, but through luck there is some intervention, such as a girder swinging in front of him or his son throwing a banjo in front of him to serve as a bridge, that allows him to continue his activity successfully.[10] So, in terms of Bickhard-representation it looks like we must say that Magoo correctly represents his environment as Magoo successfully anticipates that he will be able to continue walking. And it looks like the content of his Bickhard-representation will be massively disjunctive as the presupposition of this anticipation is that *either* the ground continues ahead of him, *or* someone throws a banjo over the chasm *or* a girder swings into place or... etc. etc. Intuitively, however, we would like to say that Magoo is short-sighted, misrepresents his environment but gets lucky. We need some account of representational content that allows us to describe the situation in these terms, and this is precisely what etiological accounts of content determination can provide. I don't see how Bickhard has the resources to explain the possibility of misrepresentation combined with good luck, which allows for ones actions to proceed as anticipated. Now of course, Bickhard can offer an account of how Magoo has limited "situation knowledge", which is going to capture something about his short-sightedness. But I think he is going to have to say that when the banjo falls in front of him to provide a bridge, he has correctly represented his environment.

[10] See, for example https://www.youtube.com/watch?v=gKZcwsM3RH0, For the banjo over the bridge watch at time 1.43.

Inaccurate-Frog cases We can think of converse cases to Magoo cases – where we might think that an organism has correctly represented the environment, but has been unlucky in their actions. Once again an explanation appealing to Bickhard-representations will not allow us to make such a fine-grained distinction, whereas an explanation in terms of Millikan-representations will. Let's suppose that a frog has evolved such that in certain types of environments in which there are flies a certain neuron fires and this causes its tongue to flick which sometimes leads to the frog catching and eating a fly. Now, let's assume that tongue flicking is extremely inaccurate, and it only rarely catches a fly. We can distinguish then between three possible cases:

(1) **Good-case.** Flies are present, the neuron fires and anticipatory tongue flicking ensues, and the frog catches and eats a fly.
(2) **Bad-aim.** Flies are present, the neuron fires and anticipatory tongue flicking ensues, and the frog does **not** catch and eat a fly.
(3) **No-Flies.** Flies are **not** present, the neuron fires and anticipatory tongue flicking ensues, and the frog does **not** catch and eat a fly.

Now, in explaining certain aspects of frog behaviour is might be useful to be able to distinguish between cases of absence of flies and failure in action execution, and so claim that in both case (1) and case (2) the frog correctly represents the presence of flies, but in case (2) the frog is either unlucky or there is some form of motor failure. And appealing to Millikan-representations allows us to make the distinction in this way. Now, one of Bickhard's main criticisms of etiological accounts of representational content like Millikan's is that they cannot account for system detectable representational error. But I take it that this is feature of such explanations, not a bug. For when something goes wrong, the system may not be able to know if it is in a situation like case (2) or case (3). But this is something that may be quite clear to an outside observer. There is some fact of the matter that distinguishes these cases; but the frog may not be in a position to know this. Now, of course, if the frog flicks for many hours and catches nothing, it is likely to move to a different location. So, the frog may have a certain (behavioural) sensitivity to errors of misrepresentation which allows it to modify its behaviour in an adaptive way (e.g. by changing its location), although this sensitivity to failure is not really organism detectable (representational) error. Now, looking at the frog from the outside, this seems like a perfectly respectable explanation of what is going on.

Bickhard, cannot, I think, tell such a story. For, from his perspective, we cannot really distinguish between situations of type-(2) and situations of type-(3). Both will be situations of misrepresentation. In both cases the anticipation is equally unsuccessful. The frog flicking at a fly who misses with the tongue flick, and the frog who flicks at "flies" on an I-phone screen fail in exactly the same way; they both anticipate swallowing a fly and don't. So, the only case of successful representation for Bickhard would seem to be situations of type-(1). In the other two cases we have cases of misrepresentation. Now, although, I don't think Bickhard's account fits in with our intuitions here, I am not claiming that his account is wrong. But I am suggesting that he is offering a quite different notion

of representation. If the only correct case of correct representation are cases of type-(1) then what is being represented is not a fly-rich environment, but something like *the-possibility-of-successfully-catching-something-edible-with-this-sort-of-tongue-flick*.[11] So the frog never really Bickhard-represents the presence of flies. And Bichard is quite explicit about this. Thus in the course of discussing frogs catching flies, he argues that "[t]here is no necessity for any representational content *about flies* here, only a functional or control phenomena. However, there is representational content in the indication in the frog that it is in a tongue licking and eating situation." (Bickhard, 1998, p.191).

Clubbed-frog cases Similarly, Bickhard is unable to distinguish between the I-phone case and a case in which we have a catastrophic break-down of anticipations. So, for example:

(4) **Catastrophe.** Flies are present, the neuron fires and anticipatory tongue flicking ensues, and the frog catches a fly but before it can swallow it, it is hit over the head with a club and stunned.

Once again, this would seem to be a case of misrepresentation for Bickhard. The frog anticipates eating flies, and is instead stunned. Cases (2), (3) and (4) would seem, for Bickhard to all equally be cases of misrepresentation. For it is a functional presupposition of the anticipation that the frog is not hit over the head with a club. So, if our frog is hit over the head with a club mid-tongue flicking it has misrepresented its environment. Now, once again, this is not a bug, but a feature of Bickhards's position. Such an understanding of representational content is a perfectly good and usable notion that can play a role in certain types of explanation, especially explanations that require organism detectable representational error. But, there are also other perfectly good notions of representation that can also play a role in quite different but also perfectly good types of explanation. Bickhard offers a coarse-grained representational content that partitions environments in terms of whether or not they successfully support particular types of activity, but there is clearly a useful, more fine-grained, notion of representational content that distinguishes between cases 1 (success) and 2 (bad aim) on the one hand and cases 3 (no-flies) and 4 (catastrophe) on the other.

Now, I have argued that explanations in terms of Bickhard-representations cannot capture the useful distinction between failure due to misrepresentation and failure due to bad luck, clumsy execution of an action, or some catastrophe. Now perhaps an interactivist might try to respond to this claim by appealing to Bickhard's distinction between contact and content. As Bickhard explains:

[11] One might think that one can distinguish between situations of type-(2) and situations of type-(3) in counterfactual terms. In that had the process of tongue flicking in a situation like type-(2) been similar but slightly different, an ant would have been caught. Whereas a slight change in technique in situation-(3) would not have resulted in tongue flicking. But it is not clear that a story like this provides us with the sort of account of organism detectable error Bickhard wants to offer.

> Consider an interactive (sub-)system engaged in interaction with an environment. The internal course of the interaction will depend in part on the organization of the system, and in part on the environment being interacted with. The internal course and outcome of such an interaction, therefore, will serve to differentiate categories of environments from each other. For ease of discussion, consider such a sub-system that has only two possible internal outcomes, A and B. If the interaction ends in state A, then the organism is in an A-type environment— the type of environment that yields internal outcome A—and, if it ends in state B, then it is in a B-type of environment. At this point, there is no content involved in being in internal state A or B. There is just a differentiation of two kinds of environments, with no characterization of those environments. So, a detection of an A-type environment is *just* differentiation, not a representation. Arriving at internal state A, however, can nevertheless potentially be useful to the organism. It might be learned, or hard-wired, for example, that, if state A is encountered, then an indication of the possibility of tongue flicking and eating of a particular sort can be set up. Such an indication is future oriented, anticipatory, and, therefore, involves content: it is about the current environment, and it could be true or false. (Bickhard, 2009, 574)

Let's say an A-type environment is one containing flies and a B-type environment is one not containing flies. Now, let's supposed that an A-type brain process is the firing of a particular neuron (let's call this neuron "Bob") and a B-type brain process is one in which Bob does not fire. In such a case it would seem that Bob can be thought of as bug detecting neuron, that differentiates between fly containing environments and flyless environments. And let's suppose that in addition to firing only in the presence of bugs, the firing of Bob leads to the anticipation of tongue-flicking and bug-digesting, and this process sometimes leads to success (but sometimes the tongue misses). Now we might think that this story allows us to differentiate between type- 2 situations (bad aim) and type- 3 situations (no flies) for it looks like we can make a distinction between processes downstream of Bob firing and process up-stream of Bob firing. We could then use this distinction to explain the difference between misrepresentation and failure due to poor execution. But Bickhard cannot tell such a story, because the process of differentiation plays no role in explaining representational content at any particular time – all the work of determining content is done down-stream.[12] Let's suppose that our frog is hit by lightning, and as a result Bob, although it still differentiates between buggy and bugless environments no longer has any feed forward effect. When Bob fires, nothing else happens; Bob just shouts into the void. Instead, another neuron (let's call this neuron "Bill") sets

[12] I think Bickhard (2019) makes this clear when he explains that "[s]uch actual interactions with environments, whether perceptual or not, constitute relevant sensitivity to those environments – relevant contact with those environments. Such contact does not itself have truth value: it differentiates, but what is differentiated is not thereby represented. (For a simple example: a vending machine might differentiate between dimes and quarters without thereby representing either dimes or quarters. Nevertheless, the differentiation can influence further processes in the machine.) It is the anticipations based on such contact that are representational, that have truth value, that have aboutness concerning the environment." (p.326).

off the exact anticipatory chain that Bob would have set off in the past. Except that Bill just fires randomly with no relationship to how things stand in the environment. However, when Bill fires, because all the anticipatory processes that follow from Bill firing are the same as they would have been if Bob had fired, then the representational content will remain the same. After the frog has been hit by lightning and Bill rather than Bob calls the shots, it sometimes turns out that while in an ant-rich environment Bill randomly sets off a chain of anticipatory activity that leads to ant eating. In such a scenario, I guess we would have to say that the frog correctly represents the environment. The success conditions of the chain of anticipatory processes are exactly the same whether they occurred as a result of Bob or Bill firing. And so, contact doesn't seem to have any role whatsoever in fixing the content. Instead, appeals to the notion of contact will merely have a role in showing how an organism's anticipations can be more or less successful. Bob is in contact with bugs, Bill is not. And so, the frog's success in catching bugs will be more likely if Bob rather than Bill is in the driving seat. So, although the notion of contact does have an important explanatory role within the interactivist model, an appeal to the contact/content distinction can't play a role in offering an interactivist account of the distinction between misrepresentation and bad technique. And while the interactivist might be right that the frog is not in a position to differentiate between misrepresentation and bad technique, we, as third person observers, might be in such a position. And from this third-person external perspective in order to be able to distinguish between misrepresentation and bad technique we need to attribute something like Millikan-style representations to the frog.

Not-getting-mauled cases Finally, it is unclear what sort of account the interactivist can give of the representational content of successful avoidance behaviour. Suppose a gazelle spots a lion in the bushes and runs. What does it anticipate? Running? Being eaten? The interactivist faces a dilemma here. In so far as the gazelle anticipates running this anticipation is successful whether or not there is a lion present in the environment. And so the gazelle who spots a lion and runs and a gazelle who is spooked by a shadow and runs both successfully represent the environment. Both environments allow for running away, and in so far as the gazelle successfully runs, its anticipations are realized. But, if we take it that the gazelle anticipates being mauled, the gazelle who spots a lion and escapes untouched fails to represent the environment correctly; only the unlucky gazelle that gets mauls successfully represents the environment. Now, I used the example of complex organisms, but much simpler organisms are capable of avoiding particular types of threats. The etiological account offers a plausible story as to how an organism is able to detect certain types of threats in the environment and avoid them.[13] Cats have evolved to be afraid

[13] Now of course, one may worry, following Gould and Lewontin (1979), that adaptationist explanations are purely conjectural just-so stories. And while I agree that this is always a legitimate concern with adaptationist explanations, I think that there are cases where such explanations are extremely plausible and should be accepted based upon inference to the best explanation. And so I think that such worries should be addressed on a case by case basis rather than leading to a global rejection of adaptationist explanations. There is a huge literature on this issue and a general defense of adaptationist explanations is beyond the scope of this paper.

of snakes – that's probably why they freak out when you put a cucumber behind them.[14] It is easy to offer a plausible explanation here in terms of Millikan-representations: the cat misrepresents the cucumber as a snake. But it is not clear to me what sort of explanation the interactivist can provide here. When an organism successfully flees from a predator, how is the interactivist going to explain the content of the organism's representations? Interactivism involves an appeal to presuppositions of anticipated interaction with the environment, but the gazelle is doing all that it can do to *avoid* interaction with the lion. Now Bickhard's account of the emotions (and particularly fear) may partially explain how an organism can represent a part of the environment as not-to-be-interacted-with. But it is not clear how an appeal to fear can explain the difference in representational content between the case when the deer runs away from a lion and when it is just spooked by a shadow. And it is not obvious that the capacity to represent and avoid predators in the environment requires complex emotions.[15]

4 Singular content

In this final section I shall argue that Bickhard-representations are essentially descriptivist, and his model is not able to offer a plausible account of singularist, essentially object involving, *de re* representations.[16] This is important because for complex organisms like us, our capacity to represent and reidentify individuals over time is essential to our cognitive lives.[17] And such capacities can only be understood in etiological terms, because the referent of a *de re* representation (such as a name) depends upon upon what the name was designed to keep track of, and so the representational content is essentially related to past practices and so needs to be cashed out in etiological terms. And Millikan offers a plausible account of *de re* singular representations.

We might think of such singular representations in terms of what Millikan (2000) calls "substance concepts" or what Recanati (2012) calls "mental files". Etiological accounts of representation, which appeal to a robust notion of what is in involved in having an ability or skill can explain such representations in terms of the exercise of an ability or skill to keep track of and reidentify particular individuals (and perhaps other entities). Proponents of singularist *de re* representations are generally

[14] Cats being scared of cucumbers here: https://www.youtube.com/watch?v=_7vML9C3PZk

[15] Obviously there is a lot more to be said on this topic, and the interactivist has various strategies they can appeal to. But for reasons of space I am unable to engage with these responses.

[16] I suspect that Bickhard might argue that from a process perspective there is no place for singular content because everything is process and there are no "substances" for singular thought to latch on to. But I think such an answer would be unsatisfactory, for although everything flows, some stuff flows more slowly than other stuff. This more slowly flowing stuff can server the role of what Bickhard calls "infrastructure". And such infrastructure (such as this keyboard and my wife) can serve as the persisting objects of singular thought.

[17] And this capacity to reidentify individuals over time is found in our capacity to reidentify members of kinds, stuffs and even, I argue, event-types. See Author Reference 1.

explicitly semantic externalists, and so are not committed to the possibility of organism detectable error. Now, if one thought that *de re* representations were the only type of representation, or if one thought that *de re* representations were somehow basic and foundational, then the lack of organism detectable error may be a problem. But I take it that *de re* singular representations are not basic as I think that only quite sophisticated organisms are capable of singular thought. As a pluralist I see no reason to assume that *all* species of representation involve organism detectable error. As long as some representations are such that we can have organism detectable error, we can explain initial learning in terms of this type of representation. And I see no reason why Bickhard needs to reject the existence of *de re* Millikan representations, for the existence of functions understood in dynamic terms is also compatible with positing etiologically specified functions. As singular thought involves the capacity to reidentify individuals over time, which we may think of a form of tracking, such capacities are essentially past-involving. Whether our tracking is successful depends upon whether we are now tracking what we started to track. If I start tracking Moriarty in the morning, but half way through the day he is replaced by a look-alike called Bob, at the end of the day I am misrepresenting Bob. I am representing him as the person I started following in the morning, but he is not the person I started following in the morning. Such forms of misrepresentation are ubiquitous for creatures like us, and are essentially past-involving and so are best understood in etiological terms.

In contrast, Bickhard's account of content, as he recognizes, only provides us with *partitions* of environment types. The representational content of an anticipatory process is understood in terms of the way the world must be in order for the process to unfold as anticipated. So, the representation is true if the world is such as to allow for successful interaction, false if it does not allow for successful interaction.[18] If representational content is only based on partitions of the environment, this is not going to explain how it is possible to represent individuals, and Bickhard himself seems to accept this. Thus, he claims that,

> The structure of possible outcomes of a differentiating interaction imposes a partition of the differentiation outcome classes on possible environments. Such partitionings replace standard conceptions of epistemology as based on correspondence—generating a *partition epistemology*—... the partitions [are] not necessarily unique to individual entities (or to entities at all)... (Bickhard, 2009, 574)

This seems to be a pragmatist analogue of a descriptivist account of reference determination. Recanati (2012) defines descriptivism in the following terms:

> Descriptivism is the view that our mental relation to individual objects goes through properties of those objects. What are given to us are, first and foremost, properties whose worldly instantiation we are able to detect, and only indirectly objects. That is so because (according to the view) our knowledge of objects is mediated by our knowledge of their properties. (p.3)

[18] To engage in an activity "constitutites an implicit predication that this is one of these environments in which this kind of interaction is appropriate" (Bickhard, 2010, 218).

According to such a position, "we get at physical objects only by a semantic shot in the dark: we specify properties or relations and hope that they are uniquely exemplified". (Chastain 1975:254; Quoted by Recanati, 2012:4). Such accounts offer a satisfactional account of reference determination. The object picked out by a descriptivist representation is whatever satisfies the description. Now Bickhard, because his representations are action orientated, is not a descriptivist in the way some philosophers are who try to cash out descriptive content in terms of sensory primitives. Instead, he should be thought of as a pragmatist analogue of a descriptivist. The infant learns to recognize the red sphere by learning how it can be interacted with. But anything that can be interacted with in the same way will be recognised at "that again". So, I take it that the interactivist account of object representations is an action-based analogue of the satisfactional account of reference determination. We represent an object in terms of a web of action possibilities, and whatever we can interact with in the same way is what is represented.

Now, as a pluralist, I think Bickard is right that all organisms have such representations, which are best understood in descriptivist terms. And I also think that such representations are more basic than singular thought. Worms likely have Bickhard representations, but not singular thought. But I think we have good reason to think that human beings are capable of singular thought and reference, and probably other animals too. We are able to track and reidentify individuals. Such concepts can be thought of as what Ruth Millikan (1999) calls 'substances concepts'. For Millikan, the ontological category 'substances', as she uses the term, is roughly 'that extensive category consisting of items about which it is possible to learn from one encounter something about what to expect on other encounters' (p. 528). And the *concept* of a substance 'is the capacity to represent the substance in thought for the purpose of information gathering and storage, inference, and ultimately the guidance of action' (pp. 530–1).[19] The function of substance concepts is not to classify things or to describe them, but to identify and re-identify them in the service of informed action. To possess such a concept one must have a reliable capacity to re-identify the object in question, but this capacity of re-identification does not fix the reference of the concept. Rather the extension of the concept is determined by the function of the concept. If we think of such concepts as files, then we can think that the file has the function of collecting information about a particular thing, and it is this that determines the reference. The function of my representation of Zübeyde is to collect useful information about Zübeyde, so I know how to interact with her next time I meet her.[20]

[19] See also Millikan (2000, 13). I prefer Millikan's earlier term 'substance concepts' rather than her more recent (2017) term, 'unicepts'.

[20] Millikan herself focuses on three types of substance concepts: concepts of individuals (Mama), kinds (mouse) and stuffs (milk). She suggests, however, without elaborating on it, that event types can also be substances in this sense, for there are many types of events about which one can learn on a single encounter how to behave on future encounters with the same event type. In particular, it seems clear that young infants show a capacity for re-identifying event types that are parts of social routines, and the capacity to re-identify such events plays an important role in early learning. In particular, I take it, that being able to recognise words as event types is essential for the capacity to learn a language. So I take it that language learning requires the substance concepts, as we need to hear a sting of sound and be able to think "that again". For a defence of these claims about event concepts, see Thorpe (2022). For an account of the role that such concepts may play in perception, see Thorpe (2015) and Thorpe (2021).

Now one of Bickhards most frequent criticisms of Millikan is that her representations would be "epiphenomenal", in the sense of not really playing any causal role. The etiological notion of function is causally irrelevant. Supposed a lion springs into existence, that was molecule for molecule identical to a lion in a zoo. The dynamics of this lion would be identical to the dynamics of the lion in the zoo. But none of its organs would have functions understood etiologically. Thus, Bickhard concludes that on Millikan's account,

> Function, in this view, is dynamically—causally—epiphenomenal. It makes no difference to the causal or dynamic properties of an organism whether or not its organs have functions. Etiological models thus fail to naturalize function. Etiological history explains the etiology of something, but it does not constitute any of the causal or dynamic properties of that something. Etiology cannot constitute the dynamics of what it is the etiology of. (Bickhard, 2009, 557)

And because Millikan representations depend upon an etiological account of function, they are also causally epiphenomenal. Millikan-style representational content plays no causal role in the causal mechanisms, or unfolding dynamics, of the system. Now I think the proper response of a defender of an etiological account of function should be to just accept the fact that explanations in terms of Millikan representations are not causal in this sense. But I think that such accounts are not intended to be causal. All the internal causal work is done by the vehicles of representation. The ascription of representational content is not supposed to explain how the internal mechanism are supposed to be working. As Shea (2018) points out, according to defenders of a representational theory of mind "[a] complete causal account of the operation of the system will be available in non-contentful terms." (p.31).

Now Millikan does sometimes seem to suggest that representations understood in etiological terms are supposed to play a role in causal explanations. Thus, for example, she claims (2020) that, "[i]f the representation is correct, an interpreter can perform a certain task or tasks of its own by "using" or being adjusted by the representation to take account of the represented state of affairs. **That the representation is correct** will then be **part of a causal explanation of how the interpreter manages to perform that function**." (p.5–6, emphasis added). Now this passage really seems to open itself up to the sort of criticism that Bickhard offers, for it is really unclear how the correctness of the representational content can really play such a causal role. One could have a similar sub-system without the causal history that could play the necessary causal role and function in causal explanations. Here, instead of swamp-man, we could have a swamp-subsystem.[21] Suppose I possess a certain mental representation that plays a certain role in my

[21] While swamp-men are merely objects of thought experiments, it is quite likely there are "swamp-subsystems" indeed such systems are likely quite common: networks of neurones that have not been selected for a particular function but emerge through something like chance and end up serving a particular function.

mental life that is underwritten by certain neural circuitry. My twin has very similar neural circuitry but lacks this particular subsystem. However, he is struck by lightning which reorganizes his brain to be like mine. Before the accident, he possessed a certain brain subsystem (the subsystem in his brain corresponding to the "interpreter" in mine) that was unable to perform its function adequately, after the accident this brain subsystem (his "swamp interpreter") is able to function successfully because it is activated appropriately by the newly formed brain circuitry. Because this new circuitry has the wrong type of causal history, it does not have any proper function, and so cannot have a role in determining representational content. And in terms of offering a causal explanation of how my interpreter and his swamp-interpreter perform their function we can tell the same causal story. So, it is not clear how appealing to the correctness of the (Millikan-style) representational content is supposed to be a part of the causal explanation of how the interpreter manages to fulfil its function. So, this passage does seem to lay Millikan open to the worry that Bickhard raises.

Elsewhere, however, Millikan (2007) herself explicitly recognizes and responds to the epiphenomenalism objection, which is to argue that explanations in terms of etiological function are not meant to be (simple) causal explanations. Thus, she points out that functional explanations are explanations in terms of purposes, "[b]ut that a thing has a purpose does not give it ''causal powers.'' Being a smoke detector does not give a thing causal powers. Some smoke detectors don't work. Purposes are not, as such, causes." (p.440). Instead of being causal explanations, such explanations, she suggests, should be understood as (Russellian) definite descriptions of causal explanations. Thus she explains:

> Consider ''The alarm is ringing because you are smoking your pipe under a smoke detector.'' It implicitly offers a definite description of the cause of the alarm. The cause of the alarm is the operation, in accordance with its design, of a certain mechanism, situated over your head, that was designed to sound an alarm when it encounters smoke. Exactly what that correctly operating mechanism amounts to is not explained; it is merely described definitely as the mechanism that is inside the thing over your head by design. History, used in this way, is no part of the explanation proper. History (i.e., purpose) is used as a convenient way to give a definite description of the mechanism that is causally involved. (pp.440 - 41)[22]

[22] And she adds: "Maybe you will want to say that this kind of explanation shouldn't count as REAL explanation at all, or not, as Shea puts it, as ''substantive,'' because the causal mechanism involved is not directly described? All that is given is a reference to the existence of a definite historically and currently exemplified causal mechanism that might be found by examining the substance (the pill) or the mechanism (the smoke detector) itself or its history. That's all right with me. You can refuse that such an explanation is substantive. However, we should keep in mind that ordinary folk *do* count these forms as explanations and use them all the time. They seem to be capable of relieving puzzlement. The teleosemanticist claims that explaining success by reference to the truth of representations is this kind of explanation. It relieves puzzlement even if you prefer not to call it ''substantive.''" (2007, p.440).

This seems a perfectly reasonable claim. And an analogous claim is made by Shea who appeals to the distinction made by Jackson and Pettit (1988) between *process explanations* and *program explanations*. Program explanations "tell us about the range of states that do or would produce the result without telling us which state in fact did the job" (Jackson & Pettit, 1988, 396). And Shea argues that if contents are fixed the way he suggests, this would show, "why semantic properties figure in program explanations, hence are explanatory of behaviour" (Shea, 2018, 208). So we might think that Bickhard-representations can figure in process explanations (the content is just the presupposition of the anticipatory activity, and so is internally related to the process) whereas explanations in terms of Millikan-representations are program explanations.

Whether we want to think of these non-causal explanations in terms of definite descriptions, or as program explanations, we have good reason to think of explanations involving singular *de re* representations in these terms. Indeed, with our representations of individuals it is clear that there is not a single mechanism involved in our capacity to keep track of individuals over time – even in the simple case of perceptual tracking of an individual. Instead, we keep track of individuals by using a complex set of quite distinct processes. And the only way of picking out or naming this set of processes is in terms of the individual they are involved in tracking. Just looking inside the organism there is no way of telling that these distinct processes and mechanisms are all involved in the exercise of a particular capacity to keep track of a particular individual. The unity of these processes can only be understood by looking outside the organism, and understanding what they are for.

For example, I may be looking at my wife when there is a power cut and our apartment is plunged into darkness. I hear her waling and then she calls out and I reach out and find her hand. Here I begin by tracking her visually; this visual tracking shifts to an auditory tracking until we finally make tactile contacts. What unifies the processes across distinct modalities is my wife. There is no single mechanism or process involved in this continued process of tracking. But it is my wife I am trying to keep track of (and if I reach out and grab the hand of another woman in the dark, I've made a big mistake.) The best we can do here is (something like) a definite description of an open-ended set of mechanisms (or processes) that have the same function: keeping track of my wife. So, explanations that appeal to *de re* representations of individuals, while perfectly good explanations, are not intended to be causal explanations, but they are perfectly good explanations for all that. So Bickhard's objection does not have any bite here.

Authors contribution Totally written by Dr. Lucas Thorpe.

Funding Boğaziçi University Research Fund (BAP) Grant Number 15681.

Data availability Not applicable.

Declarations

Ethical approval Not applicable.

Informed consent Not applicable.

Statement regarding research involving human participants and/or animals Not applicable

Competing interests None.

References

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

Bickhard, Mark H. (2004). "The Dynamic Emergence of Representation" in Hugh Clapin, Phillip Staines, Peter Slezak (eds.) *Perspectives on Cognitive Science, Representation in Mind,* Elsevier,Volume 1, 71–90.

Bickhard, M. H. (2009). The interactivist model. *Synthese, 166*, 547–591.

Bickhard, M. H. (2010). Interactive knowing: The metaphysics of intentionality. In R. Poli & J. Seibt (Eds.), *Theory and Applications of Ontology: Philosophical Perspectives* (pp. 207–229). Springer Verlag.

Bickhard, M. H. (2017). Information, representation, biology. *Biosemiotics, 10*(2), 179–193.

Bickhard, M. H. (2019). Anticipation and representation. In P. Roberto (Ed.), *Handbook of anticipation: Theoretical and applied aspects of the use of future in decision making* (pp. 323–338). Springer Verlag.

Campbell, D. T. (1974). Evolutionary Epistemology. In P. A. Schilpp (Ed.), *The philosophy of Karl Popper* (pp. 413–463). Open Court.

Campbell, D. T. (1990). Levels of organization, downward causation, and the selection-theory approach to evolutionary epistemology. In G. Greenberg & E. Tobach (Eds.), *Theories of the evolution of knowing* (pp. 1–17). Erlbaum.

Chastain, C. (1975). Reference and context. In K. Gunderson (Ed.), *Language, mind, and knowledge* (pp. 194–269). University of Minnesota Press.

Cusimano, S., & Sterner, B. (2019). Integrative pluralism for biological function. *Biology and Philosophy, 34*, 55.

Erdin, H. O., & Bickhard, M. H. (2018). Representing is something that we do, not a structure that we "use": Reply to Gładziejewski. *New Ideas in Psychology, 1*(49), 27–37.

Garson, J. (2016). *A Critical Overview of Biological Functions*. Springer.

Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society London, Series B, 205*, 581–598.

Jackson, F., & Pettit, P. (1988). "Functionalism and Broad Content" *Mind. XCVII, 387*, 381–400.

Millikan, R. G. (1984). *Language*. MIT Press.

Millikan, R. G. (1999). A common structure for concepts of individuals, stuffs, and real kinds: More mama, more milk, and more mouse. In E. Margolis, & S Laurence (Eds.), *Concepts: Core readings*. MIT Press.

Millikan, R. G. (2000). *On clear and confused ideas: An essay about substance concepts*. Cambridge University Press.

Millikan, R. G. (2007). An input condition for teleosemantics? Reply to Shea (and Godfrey-Smith). *Philosophy and Phenomenological Research, 75*(2), 436–455.

Millikan, Ruth Garrett. (2020). "Neuroscience and teleosemantics", *Synthese*.

Nicholson, D. J. spsampsps Dupré, J. (Eds.) (2018). *Everything Flows: Towards a Processual Philosophy of Biology.* Oxford University Press.

Recanati, F. (2012). *Mental Files*. Oxford University Press.

Shea, N. (2018). *Representation in cognitive science*. Oxford University Press.

Thorpe, L. (2015). Seeing white and wrong: Reid on the role of sensations in perception, with a focus on color perception. In T. Reid (Ed.), *Mind, knowledge, and value (Mind Association Occasional Series)* (pp. 100–123). Oxford University Press.

Thorpe, L. (2021). Thomas reid on the role of conception and belief in perception and memory. *History of Philosophy Quarterly, 38*(4), 357–374.

Thorpe, L. (2022). Atomic event concepts in perception, action and belief. *Journal of the American Philosophical Association, 8*(1), 110–127.