

Content and action: The guidance theory of representation

MICHAEL L. ANDERSON

Institute for Advanced Computer Studies, University of Maryland, College Park, USA (E-mail: anderson@cs.umd.edu)

GREGG ROSENBERG

Artificial Intelligence Center, University of Georgia, USA (E-mail: ghrosenb@ai.uga.edu)

[Authors listed in alphabetical order]

Abstract. The current essay introduces the guidance theory of representation, according to which the content and intentionality of representations can be accounted for in terms of the way they provide guidance for action. The guidance theory offers a way of fixing representational content that gives the causal and evolutionary history of the subject only an indirect (non-necessary) role, and an account of representational error, based on failure of action, that does not rely on any such notions as proper functions, ideal conditions, or normal circumstances. Moreover, because the notion of error is defined in terms of failure of action, the guidance theory meets the “meta-epistemological requirement” that representational error should be potentially detectable by the representing system itself. In this essay, we offer a brief account of the biological origins of representation, a formal characterization of the guidance theory, some examples of its use, and show how the guidance theory handles some traditional problem cases for representation: the problems of error and of representation of fictional and abstract entities. Being both *representational* and *action-grounded*, the guidance theory may provide some common ground between embodied and cognitivist approaches to the study of the mind.

Keywords. action, intentionality, mental content, reference, representation

Content and action: The guidance theory of representation

Abstract. The current essay introduces the guidance theory of representation, according to which the content and intentionality of representations can be accounted for in terms of the way they provide guidance for action. The guidance theory offers a way of fixing representational content that gives the causal and evolutionary history of the subject only an indirect (non-necessary) role, and an account of representational error, based on failure of action, that does not rely on any such notions as proper functions, ideal conditions, or normal circumstances. Moreover, because the notion of error is defined in terms of failure of action, the guidance theory meets the “meta-epistemological requirement” that representational error should be potentially detectable by the representing system itself. In this essay, we offer a brief account of the biological origins of representation, a formal characterization of the guidance theory, some examples of its use, and show how the guidance theory handles some traditional problem cases for representation: the problems of error and of representation of fictional and abstract entities. Being both *representational* and *action-grounded*, the guidance theory may provide some common ground between embodied and cognitivist approaches to the study of the mind.

Keywords. action, intentionality, mental content, reference, representation

1. Introduction

Representations occupy a prominent place in—on some accounts the very center of—our mental lives. Representations succeed, not just sometimes but usually, and not by chance but by function, in being *about* specific things; a representation is something that stands in for something else.

How is one thing ever *about* another? To answer this question is usually to analyze this relation of aboutness—the intentionality of a representation—in terms of some other, presumably more basic relation. For instance, a simplistic causal theory of representation might hold that a given representation **R** is about **E** just in case it has a certain specified set of causal relations to **E**, for instance, that perceiving an instance of **E** will cause one to represent with **R**. Likewise an information-content approach might hold that a given representation is about that object from which the information it contains in fact derived. Conceptual role theories, on the other hand, try to analyze meaning in terms of the role played by the concept in inferential and other conceptual/cognitive processes: roughly

speaking, the representation **R** is about **E** just in case it is used to make warranted inferences about **E**. Naturally, there are also theories that try to combine these two approaches, producing the so-called “two-factor” accounts.¹ In this essay we will be introducing a somewhat different approach to the matter, and while we will not begin with detailed critiques of these forebears (for this see Anderson, forthcoming-a) let us say, by way of situating our own account of representational content, that we find the various causal approaches too *input focused*, meaning they give too much importance to the ways in which the environment affects the organism to endow its states with representational meaning. While the conceptual role theories seem to us a step in the right direction in that they draw attention to the importance of cognitive actions taken by the subject with its representations, *none* of the approaches outlined above give sufficient weight to the full range of what a subject *does* with its representations.² In contrast, we would like to suggest that there is a rather different, inside-out perspective on these matters available to the naturalist, one that is in line both with a broadly functionalist conception of representation and also a naturally ecological view of the organism.³ We ask first not what a representation *is*, but what it *does* for the representing agent, and what the agent does with it; what is a representation *for*? Our contention is essentially that representations are what representations do.

¹ For a well-known causal theory see (Fodor, 1981; 1987); for an information-content approach see (Dretske, 1981; 1986; 1988); a conceptual role theory is detailed in (Harman, 1982; 1987); and two-factor theories are developed in (Block, 1986; Loar, 1981; Lycan, 1984). The sketches above are meant only as a simple way to situate the account being developed here, and are in no way meant to be representative of the depth and sophistication of the work just listed.

² This is arguably the result of a lingering Cartesian influence on the cognitive sciences (Anderson, 2003; Thompson, 1996), and a misguided view that the primary function of the senses is to provide veridical, objective information about the state of the world (Akins, 1996).

³ Of course Dretske (1986; 1988) also adopts the functional perspective. But his theory remains an information-content approach to representation, which contrasts sharply with our approach of cashing out content entirely in terms of action-guidance.

What do representations do? We hold that what a representation *does* is provide guidance for action. Expressed at the same level of simplicity as those approaches mentioned above, on the guidance theory **R** is about **E** just in case **R** is standardly used by an agent to guide its actions with respect to **E**. Whatever the details of its instantiation or structure, whatever its physical or informational features (and these are quite various across different representing systems), what makes a given item *representational* is its role in providing guidance to the cognitive agent for taking actions with respect to the represented entity.⁴ In our view, those other special features a given representing token might possess—e.g. co-variance with, openness to the causal influence of, or resemblance to its object—are each elements in a range of strategies that our various representation-forming and representation-consuming systems have evolved to solve the biologically fundamental problem of providing autonomous organisms with guidance for action.

Thus, we will not (nor would we want to) deny the importance of all these various features of representations—although this can make our theory look like any number of alternatives at specific points in the exegesis. For instance, Bickhard’s theory of interactive representation (Bickhard, 1993; 1999) is very much part of the same pragmatic, naturalistic tradition as the guidance theory, and it is motivated by the same

⁴ Interestingly, although Kathleen Akins (Akins, 1996) makes much of the fact that the senses evolved primarily as motor-control systems, she draws from this what from our perspective is exactly the wrong moral: that the various sensory-motor systems are *not* representational, and that there is, therefore, a gap between the sensory systems, which guide our actions, and what she calls the “ontological systems”, which represent what is where (see esp. pp. 366-72). It is not clear whether the theory on offer in the current article should be understood as a way of *closing* this gap, or whether, in virtue of our starting point, we merely deny it.

fundamental insight regarding the epistemic importance of action and interaction.

Likewise, the guidance theory sides with Grush (1997; 2003; 2004), Lloyd (1989) and Millikan (1984; 1993)—against, for instance, Dretske, (1981)—in emphasizing the importance of “consumer-oriented” accounts of representation. What is new about the guidance theory is not that it is naturalistic, functionalist, and consumer-oriented, but rather that it insists that the fundamental ground of representational content is action guidance.

To see how this plays out in a specific case, consider a comparison with Rick Grush’s emulation theory (1997; 2003; 2004). In our view, Grush’s theory has two main attractions. First, Grush rightly emphasizes the centrality of the consumer in any account of representation—a representation is something that is *used as* a stand-in for something else. Second, Grush partially analyzes this relationship of standing-in in terms of an item **R** being a mental model employed to make predictions regarding the state of a system **E**, given some (hypothetical) set of interventions by the user—Grush calls such models *emulators*. For Grush, emulation is a paradigmatic case of representation. As will soon be clear, the guidance theory is fully compatible with these claims.

However, Grush leaves a crucial part of his account underdeveloped—namely, what, exactly, it is for **R** to be *used as an emulator of E*, that is, what ties it to **E** in particular? Is it some kind of functional isomorphism of **R** to **E**? The participation of **R** and **E** in an informational feedback loop? Our account answers this crucial question by saying that **R** is used as an emulator of **E**, and is, therefore, a representation of **E**, just in case (as with

all representations) **R** is used to guide the subject's actions with respect to **E** (and, as will be seen below, on-line guidance of overt action is only one of the many relevant sorts of guidance).

This particular line is not developed at all by Grush, and although it appears to us to be a natural choice, there are other options available, and it is not clear from Grush's writings that he would endorse grounding his account in the way the guidance theory advocates. Thus, one way of understanding the relationship between the guidance theory and the emulation theory is that we explore in detail what it is for an **R** to be used to guide actions with respect to an **E**, by way of grounding the intentional connection in the behavior of the representing system, thus providing to the emulation theory a more basic and general account of the crucial relation of *standing-in*.

Naturally, as should be clear, *compatibility* does not imply *reducibility*. In fact, the guidance theory is significantly broader than the emulation theory; a Grushian emulator is only one of many things that represent. Likewise, this comparison ignores the various differences in detail and emphasis between the theories; for instance, we make no strong distinction between *presentation* and *representation*. And finally, Grush's main interest is not the foundations of intentionality, but rather the usefulness and biological plausibility of emulators as control systems.

The guidance theory offers some different advantages over, for instance, Lloyd (1989) and Millikan (1984; 1993). In particular, fixing representational content with respect to

current/ongoing action-guidance gives evolutionary history only an indirect role, and does not require any particular history of causal interaction between the representing subject and the represented entity. Further, the guidance theory formulates an account of representational error, based on failure of action, that does not rely on any notion of proper function, or normal circumstances, or ideal conditions (all of which notions are extremely vexed). And finally, because the notion of error is defined in terms of failure of action, the guidance theory meets Mark Bickhard's "meta-epistemological requirement" (1993; 1999) that representational error should be potentially detectable by the representing system itself. Insofar as representing systems are not generally in a position to assess the causal circumstances, nor the causal or evolutionary histories, that define correct representation, they will not be in a position to determine when they are in error.

As has been noted already, above, the guidance theory is very close in spirit to Mark Bickhard's interactivist theory of representation (1993; 1999). However, we offer a significantly different development and formalization of our shared insights. For instance, Bickhard's analysis relies heavily on control theory, cashes out representational content in terms of "environmental interactive properties", and assumes some version of process ontology. The guidance theory, while compatible with these possibilities, does not require them, making our analysis somewhat more flexible, and more accommodative of mainstream analytic frameworks.

One way to think about the relation between Bickard's interactivism and the guidance theory is to say that the guidance theory singles out the action-guidance aspect of representation that is also a central part of Bickard's theory, and formalizes it as a stand-alone theory of content. The intent is both to bring out more clearly the power and implications of the claim, and also to make aspects of the theory more compatible with mainstream approaches to the study of the mind. One hope is that it might prove possible for the theory to be investigated and tested as part an empirical theory of representation, without having to pull apart its empirical implications from the normative and metaphysical aspects of Bickard's theory. Another is that expressing an action-grounded theory of representation in the formal language familiar to analytic philosophy will allow it to make more headway against the popular analytic theories of reference than has been possible for interactivism. Still, the relative advantages of these two analyses remain largely to be determined.

All these issues deserve further discussion, but neither this essay (nor especially the introduction to this essay) is the place to offer it (but see Anderson, forthcoming-a). Hopefully what has been said is nevertheless sufficient to situate the guidance theory with respect some well-known landmarks, illuminate its main distinctive feature(s), and thus prepare the way for the detailed exegesis to come.

In what follows we will offer a brief, hypothetical account of the biological origin of representation, which will serve both to explicate our basic claim regarding the function of representation, and to highlight some key features of representations in terms of which

we can understand that function. From there we will offer a formal characterization of the guidance theory of representation and some examples of its use. We will end with some suggestions as to how we can use the guidance theory to approach some traditional problem cases for a theory of representation: the problems of error and of representation of fictional and abstract entities.

2. The biological origin of representation

The legacy of evolutionary development is a set of capacities and structures—physical (kidney), perceptual (edge-detection module), psychological (memory)—each of which evolved to perform a certain function, or set of functions, which ultimately aid the reproductive success of the members of the species that possess them. Representation is an adaptation for making behavior more flexible and effective. In what follows we will present two brief “case studies” of behavior, the first a base-line study of turning behavior in the slime-mold, effected without representations, and the second a study of very similar, but representation-guided, behavior in the frog. The discussion is meant both to buttress the main claim of the paper, that the function of representation is to guide behavior, and also to suggest an evolutionary justification for the emergence of representations and their central features.

2.1 Case study 1: Phototaxis in D. discoideum

The slime mold *Dictyostelium discoideum* is a remarkably interesting creature with a surprising and strange life cycle. In one stage of this cycle, the slime mold exists as single-cell amoebae, feeding on bacteria in forest litter. Starvation in these amoebae,

however, induces them to aggregate into a large (up to 100,000 cell) multicellular mass, the “slug”, which looks and behaves like a single organism. Eventually this slug transforms itself into a fruiting body, sending up a stalk and spore head, facilitating the dispersion of spores over the surrounding area, and beginning the cycle anew (Bonner 1985).

It is the slug stage that is of interest to us here. Coordinated by waves of cyclic adenosine monophosphate (cAMP) originating in its tip, the slug migrates to a favorable location for fruiting, guided by sensitivity to light, pH, and temperature (Bonner, 1994). The particular behavior on which we will focus is phototactic turning: exposed to light, the slug will turn toward and head in the direction of the light. Although there is some functional specialization of cells in the slug stage—the slug is organized into prestalk cells at the anterior end, and prespore cells at the posterior; the coordinating waves of cAMP are generated only by the tip of the slug; and the slug is sensitive to light only in the anterior prestalk zone (Poff and Loomis, 1973; Häder and Burkart, 1983)—the slug has no sense organs or centralized signaling system, and thus no way to register⁵ or pass on information regarding the direction of a light source. Nevertheless, the phototactic turning is pronounced and accurate, and in the slug’s natural environment, light is an indication of an open area, favorable for the dispersal of spores. Phototactic turning and migration, then, are goal-directed *behaviors* of the slug, specifically selected by evolutionary forces.

⁵ In our usage, to register (or, more generally, a registration) has specific meaning. See footnote 8 and related text, below.

(Miura and Siegert, 2000) propose an interesting and plausible mechanism for this behavior. Their experiments show that light stimulates cAMP production, and thus increases the frequency of cAMP signaling waves. It is further known that the individual cells of *D. discoideum* are positively chemotactic with respect to cAMP (Bonner, 1985) and that anterior prestalk cells exhibit a stronger chemotactic response to cAMP than do anterior, prespore cells (Sternfeld and David, 1981; Traynor *et al.* 1992). These facts together suggest that when a slug is exposed to a unilateral light source, cAMP signaling is unilaterally increased, causing increased chemotaxis to that side of the slug tip, changing its geometry and orienting it toward the light source. As the tip becomes increasingly oriented to the light, one would expect the lateral gradient of cAMP production to diminish, leading to stimulation of production equally across the tip of the slug. The result, which is in fact observed, is increased movement in the direction of the light source.

We will take this to be a prototypical case of the evolutionary pre-conditions that allowed for the emergence of representation-driven behavior. Key to the slug's behavior is the fact that an advantageous action is driven by a distinct internal state—the cAMP gradient caused by unilateral illumination of the slug tip. This raises the possibility of internal mechanisms that can recognize and categorize these distinct states, and which can coordinate actions in response to these categorizations. In short, the slug's basic abilities and circumstances show how there could be a useful pre-representational selection environment in which representation consuming cognitive systems could evolve and provide advantages. A creature so endowed could use categorizations of its distinct

internal states to take more sophisticated advantage of those states' pre-existing capacities for providing guidance for action, allowing over generations for more and more indirect interaction with environmental influences. Thus distinct non-representational but action-guiding bodily states, like the slug's, by being categorized and consumed by a cognitive engine and exploited for self-directed behavioral control, can give rise to cognitively significant representational states.

2.2 Case study 2: Prey capture in frogs

Although of course the frog is a vastly more complicated animal, structurally and behaviorally, than the slime mold, the single behavior on which we will focus here is relatively simple and easy to describe: when a small, dark, moving dot comes within the frog's field of vision, it will orient its body to the dot, and snap at it.⁶

All three of these features are necessary to elicit this behavior; the frog does not respond to an object that is stationary with respect to, or is lighter than its background, and if the object is large it will move *away* from the object, and not attempt to capture it.⁷ To facilitate comparison with the phototactic turning of the slime mold slug, we can narrow our focus even further, and consider only the frog's orientation of its body in response to this stimulus. The general mechanisms whereby this behavior is elicited are also relatively straightforward: the appearance of a small, dark, moving object causes a typical

⁶ *R. pipiens* captures small dark things with its tongue, but other feeding behaviors are possible for other species of frog, and *R. pipiens* displays different behaviors in the presence of different prey (Deban *et al.*, 2001).

⁷ For some of the fascinating details of these mechanisms and responses see, e.g. (Lettvin *et al.*, 1959; Ingle, 1973).

pattern of firing in “class II” retinal ganglion cells, which project to the superficial laminae of the optic tectum (Ingle, 1991). Cells in the optic tectum, in turn, project to nuclei in the pons, medulla and spinal cord that control turning (Milner and Goodale, 1995). We needn’t suppose any complex encoding of spatial information is necessary to govern this response; rules for movement that bring the stimulus into a certain portion of the visual field will do.

There is a great deal of continuity between the phototactic turning of the slime mold slug and the prey-orientation of the frog. In each case, certain features or events in its environment stimulate an inner mechanism of the animal, causing a typical and predictable, but nevertheless useful and functional response. Yet, there are also some crucial, if (at this level of simplicity) subtle differences, in terms of which we can begin to understand the nature and role of representations in organisms (like the frog) that possess them.

First, and most obvious, it seems that in the slime mold, but not in the frog, the stimulus produces its effect through first-order changes to the slime mold. The bodily changes caused in the slime mold slug by a light source are themselves sufficient to drive its behavioral response. Conversely, in the frog, but not in the slime mold, the stimulus triggers a bodily change (at the frog’s retina) that is *registered* by a further, internal, self-driven process. The bodily state differences produced in the frog’s retina are not by themselves capable of driving or directing any behavior on the part of the frog. Rather, the stimulations of the retina generated by the small, dark, moving object are registered

by, and taken up into, a cognitive system that can consume the registration by exploiting its capacity to guide the frog's behavior in a sophisticated and coordinated way, in context with other registrations.⁸ In the slime mold slug there is no such intermediate registration of bodily changes in an integrated control system.

This difference is critical enough to introduce the notion of a *potential decoupling* of stimulus and response. This decoupling could occur at either of two substantial processing points. A first kind of potential decoupling could occur between the stimulus (and the bodily changes it typically produces) and its registration by a cognitive system. As an example of the this case, not just flies (the biologically relevant target), but bits of paper, dots on a screen, or an electrode applied to the right area(s) of the brain, could all generate the registrations (neural firings) characteristic of fly-detection, and thereby cause the frog to turn *as if* there were a fly. This possibility is a crucial marker of representations; the close coupling of the causal characteristics of the light source with the behavior of the slug, in contrast, does not appear to allow for the activation of the behavior-causing mechanism in the absence of the light.⁹

⁸ As we will use the term, a *registration* is a distinct and characteristic inner state, typically formed in response to a certain kind of bodily (sensory) change, and taken up into a behavioral control system. As we will argue below, registrations were the evolutionary forebears of fully-featured representations, and thus are one specific kind of (simple) *representation*.

⁹ This is a debatable point, interesting in its own right, but largely irrelevant to what follows, as it is the second type of de-coupling that is most important to the theory we offer. Nevertheless, for those interested in following out the point, we offer the following: It is tempting to object that the distinction cannot be maintained, since in both cases there must be a cause for any observed behavior, and in both cases the cause can be only those things actually capable of triggering the relevant internal process. The objection is true so far as it goes, but very misleading. For let us assume that there *were* another kind of stimulation capable of driving the mechanism responsible for phototactic turning. It wouldn't follow from this that phototaxis in fact utilizes representations, for the explanation for the behavior in such a case would necessarily make reference to the *causal features* of the stimulation that allow it to drive the same mechanisms that light normally drives. Not so in the case of the frog, for the wider range of possible objects—flies, bits of paper, dots on a screen, well-placed electrodes—with quite different causal features and physical natures, that could be responsible for stimulating the behavior in question strongly suggests

A second potential decoupling is of the registration from the behavior it might otherwise mechanically cause. Consider the classic experiments reported in (Ingle, 1973). Ingle unilaterally removed the frog's optic tectum, to which the optic ganglia implicated in fly-detection originally projected. Initially, the frog simply lost its ability to respond to prey-like stimuli in the region of the visual field contralateral to the ablation. However, over a period of months, the severed optic tract regenerated, and, finding no tectum to innervate, continued to grow until it had passed over the brain's mid-line and innervated the intact optic tectum. As this occurred, the frog regained its ability to respond to stimulus in the formerly "blind" area of its visual field, but its *response* was to turn to the equivalent location on the *other* side of its body. That is, it would turn away from the fly and snap at the air where a mirror image of the fly would be.

Clearly, the fly-registration delivered by the optic ganglia does not, in and of itself, determine the behavior of the animal; the same registration can cause quite different behavior depending on the action-guiding system to which it is fed. Note the implication that the registration does not carry *objectively specified information* about the location of

that it is *not* the causal features of objects that are relevant here. The only thing that unites all these possibilities is the very fact that they can trigger the characteristic patterns of neural firing associated with fly-detection, which is just to say that they all trigger the frog's fly-registration. The dogged objector may push further, and suggest that this common feature should be characterized *causally* in terms of the ability to generate the registration. But, first of all, this is a somewhat non-standard way to characterize an empirical cause (which perhaps should focus on such things as the kind, amount, and availability of energy). Second, it hardly accomplishes the purpose of the objector, if that is to suggest that the frog and slime mold are, in this regard, ultimately indistinguishable, for no such locutions are necessary in the case of the slime mold; straightforward empirical characteristics would be sufficient to identify the relevant class of causes. In the case of the frog, however, the characterization of the cause appears to require reference to registrations, and a cognitive system in which they are processed. Finally, even if the move is technically admissible, it misses the significance of representation-driven control of behavior, which is precisely that it allows for the liberation of the organism from the strict reliance on the causal characteristics or powers of the elements of its environment.

the fly-object; for if the optic tectum and other downstream structures were in the business of interpreting objectively specified information, then the contralateral projection of the optic ganglia ought to offer no barrier to correct behavior in this case. There are two separate but related points to be made in this connection. First, it makes little sense to think of the content of the fly-registration in isolation from the registration-consuming mechanisms it normally targets. Thus, it is perhaps better to think of the content of the registration by analogy with a gear of a certain shape (or a circuit board with a given set of connections) that will produce outcomes depending in large measure on the nature of the mechanism into which it is fitted. Second, to think of the frog as representing (or, in this case *mis*-representing) its environment, it is not necessary to think of it as *picturing* that environment. In the case of the frog, the registration is not processed to generate an objective picture of the world, but is used for motor-control, to determine the direction and value of the animal's turning.

In our view, it is a mistake for theories of representational content to focus overly much on the intrinsic characteristics of given environmental registrations, and to hold up picturing as the essential function of representation.¹⁰ Instead, representational content can be understood in terms of its ability to provide guidance for action, and the action-guiding content of a registration must be understood in terms of the function and structure of the guidance-taking mechanisms it normally targets. Note that this distributes

¹⁰ This is not to deny that there exist registration-consuming mechanisms the function of which *is* to generate a picture of the world. It is only to suggest that registration-consuming mechanisms that (help) generate behavior without picturing can still be considered representational. It is also worth considering the value to the organism of picture-generation. Might it not be the case that the justification for such mechanisms lies in their value for the generation and control of behavior?

responsibility for representational content between the registration and the guidance-taking mechanisms, a fact that is an explicit part of the formal theory we offer, below.¹¹

Still, as important as these implications may be for a theory of content, we should be careful not to lose sight of the central significance of the development of representation mediated behavioral control: it allows the behavioral response of the animal to be relatively unconstrained by any causal force transmitted to the animal by the registration-triggering object. Instead, behavior can be guided primarily by representations built upon these registrations.

When guiding behavior with representations, the central question is one of integrated control, not one of harnessing the specific causal forces of the environment for useful ends. Achieving integrated control requires transmuting different kinds of significant causal force in the environment into a generic causal basis whose form the integrated control system may process. Guidance then can come from differentiating features of tokens in the representation system, including both their representational form and their context of activity. These differentiating features turn out to be the most important items for directing the (self-driven) inner mechanisms of the animal. Vitamin D production, produced by sunlight hitting the skin, is one example of a (non behavioral) process in humans that relies on the causal characteristics of elements of the environment to produce

¹¹ Note further that, given that these guidance-taking mechanisms will involve such things as motor-control, muscle activations, and the like, and that representational content is generally appropriate to guide such systems, this introduces reason to doubt that representational content can always be cashed out in terms of conceptually-structured, language-like propositions.

its effect. We are claiming, essentially, that the phototactic turning of the slime mold slug is more like vitamin D production than like the prey orientation of the frog.

In comparing the slime mold to the frog, we can perhaps see this evolutionary progression illustrated: nature moves from externally caused bodily changes directly driving behavior in the slime mold, to the establishment of a behavior control system in the frog where distinct externally caused bodily effects are registered by the organism and built into representations used to guide the organism's behavior. This progression is immensely important, as it brings with it the potential for behavior driven by stimulus-registrations that occur without the stimulus being present, as well as for non-mechanical behavior with respect to the stimulus (by processing and contextualizing an organism's registrations within an integrated control system). For it's not just that an animal like the frog can evolve to react one way, or another way, to a given registration. A given registration can be processed in combination with other registrations of circumstance—for instance, responding to a fly-detection by suction-feeding while in water, and by tongue-prehension when on land (Deban *et al.*, 2001)—and furthermore, the same registration can be consumed by *more than one* action-guiding system, resulting in different (albeit coordinated) behaviors, as is the case with turning toward a prey-object and snapping at it, behaviors that are in fact controlled by separate systems (Ingle, 1991).

Finally, a behavioral control system based on action-guiding representations allows for the possibility of forming inner states that in fact guide action with respect to certain objects or properties, even though these objects and properties are not currently, nor ever

have been registered. Just as behavior need not be mechanically caused by the world, neither do representations need to rely for their formation on specific stimuli. A useful inner state can form by chance, by synthesis from extant states, or can be hypothesized or imagined into existence. But however a given action-guiding inner state came to be, what is important in determining its content is not its causal history, nor its relation to some set of stimuli, but only toward or with respect to *what* it guides the organism's behavior.

Putting this differently, the causal or informational *facts* about a given action-guiding inner state matter not a whit in themselves to its function; when it is utilized in guiding an action toward an object or property, the system is structured to utilize the representation *just as if* it carried information about the object or property in question. That is, a guidance control system is built to act on hypotheticals, in a certain sense: it makes *assumptions* that its representations are built upon registrations that provide it with information. But the ability of the representation consumer to make these implicit or explicit assumptions of information about objects or properties opens up the possibility that it might make them about things to which it is not, and never has been, causally or otherwise connected but about which it needs useful guidance.

All this is to emphasize that, on the guidance theory, registrations—and the inner mechanisms that utilize them by building and consuming representations—emerge as a specific enhancement to the behavior control systems of the organism, and their content and function must be understood in this light. Although the development of

representation-producing and consuming systems was a giant evolutionary leap, its significance is not best elucidated in terms of information-containing, world-reflecting, or situation-modeling inner states. Functionally, these systems are instead best understood as continuous with the older, more world-driven behavioral systems they replaced: they are the things that provide guidance to the integrated systems for behavioral control. This aspect of the argument we offer is nicely summarized in (Milner and Goodale, 1995):

Vision in the frog, like vision in other organisms, did not evolve to provide perception of the world in any obvious sense, but rather to provide distal sensory control of the movements that the animal makes in order to survive and reproduce in that world. Natural selection operates at the level of overt behavior; it cares little about how well an animal ‘sees’ the world, but a great deal about how well the animal forages for food, avoids predators, finds mates, and moves efficiently from one part of the environment to another. To understand how the visuomotor systems controlling these behaviors are organized, it is necessary to study both the selectivity of their sensory inputs and the characteristics of the different motor outputs they control. (p.11)

We suggest that neither the primary function of registrations, nor the best way of specifying the representations they eventually came to support, radically changed as a result of any of their further evolutionary development. What we see instead are variations on and sophistications of this basic theme. For now, with this as a background, we turn to the formal account of the guidance theory.

3. A formal account of the guidance theory

Although we have introduced the guidance theory only in the case of mental representations, it is a *general* theory of representation. It applies with appropriate

modifications also to non-mental, natural and artifactual representation.¹² However, the development below is specific to mental representation.

In simplistic form (the final, technical definition of representation is offered in definition 11), the account we are offering argues that a given mental token **T** (of some given type) represents an entity **E** for subject **S**, just in case tokens of that type are standardly used by **S** to guide its actions with respect to **E**. Naturally, there are a great number of terms in this account that need explaining, foremost among them: what is it to “guide action”, what is it to be “standardly used” to do so; and what is it to take action “with respect to” an entity. We offer discussions and analyses of these terms, below.

However, let us acknowledge at the outset that all naturalistic theories of representation face a basic choice: to focus on the most primitive of cases, which invites the charge that the account cannot be scaled up to reflect the complexity of human representation; or, to try to include from the beginning the necessary machinery to capture human-level complexity, which choice invites the critical reader to search for circularities, those places where unanalyzed intentional terms have found their way into supposedly foundational definitions. We have embraced the latter approach, and are fully aware that in doing so we invite inspection of the sort described.

¹² Consider, for instance, the case of a wall calendar; it represents the days, weeks and months of the year just because it is standardly used to provide guidance to one’s actions with respect to these entities. Note that it represents the days, and provides guidance with respect to them, despite the fact that it is not causally connected to, nor does it co-vary with, the days (assuming that one does not mark off the days in using it). An item’s status as a representation depends not on such things as co-variance (although it could include such relations), but rather on its place in a system of practices and procedures within which it provides guidance.

The guidance theory is a functionalist account of representational content, stated in terms of a larger (and more basic) functionalist view of the intentionality of action. Like all functionalist views of intentionality, it does not seek to wholly eliminate intentional terms from its definitions, but embraces a certain amount of initial circularity. The challenge we accept is to appropriately circumscribe references to intentional notions so they form a recognizable pattern, and to point to how the pattern as a whole hooks onto parts of the organism and the world.

As an example, below we define representation partially in terms of a subject's circumstances and a subject's circumstances partially in terms of representation. We take the circularity invoked to be harmless because we do not believe assignment of content, in the relevant sense, to be necessary to identify that something is behaving like (and therefore is) a representation. Rather, the mere classification of something as a representation, as we define it, is possible based on simple functionalist grounds of understanding that it is providing guidance of the right sort to a creature's action-guiding mechanisms (see below).

We believe it is then a further question, once one recognizes that something is acting as a representation, to ask what its content is. Similar remarks apply for other seeming circular appeals to intentional states such as expectation or memory. We do not agree with some author's views that one must understand the content of mental states before one can even individuate them or classify them as representations, memories, values, preferences, and so forth, for it is possible to see at an abstract level what role a state is

playing in guiding a creature, to classify the state accordingly, and to treat assignments of content as a separate question.

The terms, definitions, analytical strategies and promissory notes that follow have been crafted according to this approach, and while we cannot, of course, be certain that no hidden circularities remain, we are confident that there are no *obvious* ones. Nevertheless, one cannot say everything all at once, and at any given point *in medias res*, there may well be ungrounded terms. From the reader, then, we ask the patience to let us reach the end of our formal account. Although we have hardly argued every detail nor explored every possibility, we do believe that we have laid down a consistent, non-circular foundation for a new theory of representation, and by that point it should be discernable, at least in outline.

First, then, let us say that a token *provides guidance* to a subject by making its features available to the subject's motor systems and rational control processes for use in making discriminating choices between possible actions or possible ways of executing actions. On the guidance theory action is fundamentally a directed engagement with the world. Representations come into existence and derive their content from their role supporting the basic directedness of action. It is important to note at the outset that, for the guidance theory, action-guidance is not limited to the on-line control of overt physical action (see definition 7). Thus, it is no part of our theory that the only items with content are those deployed in the immediate, on-line guidance of overt behavior. The guidance theory is an action-grounded theory of representational content, but is in no way a *behaviorist*

theory (really, it ought to go without saying that no theory of mental content is behaviorist).

Before moving on to the more involved discussions of “standardly used” and “taking action with respect to”, there are some basic terms that bear defining first: *token*, *type*, *entity* and *subject*. We intend these definitions to be relatively straightforward, and they are offered primarily in the interest of providing a clear and consistent formal account.

Definition 1: A *token* is any entity with a history and a location.

Definition 2: A *type* is any classical or fuzzy set resulting from a consistent way of categorizing tokens based on a natural similarity metric.

Definition 3: An *entity* is anything that can be represented: a property, a concrete particular, an aspect of a thing, a state of affairs, a number, etc.

Definition 4: A *subject* is any representation-consuming cognitive engine. To be a representation consumer, it must be capable of interacting in the world in a rational, goal directed way due at least partly to guidance it receives from tokens within its cognitive systems. Note that this definition ties representation to cognition, thereby excluding the slime mold’s states from being representations.

This puts us in a position to define what it means to be “standardly used” to guide action:

Definition 5: A subject *standardly uses* tokens (of a type) to provide guidance with respect to an entity **E** in a given type of circumstances **C** if, and only if, the subject has an enduring conscious preference or conditioned reflex to use the tokens (i.e., members of the type) to provide guidance with respect to **E** when in circumstances **C**.

There are a couple of things to notice about this definition. The first is that it introduces a new term, *circumstance*, which will be defined below. The second is that it already distances the guidance theory from theories requiring specific causal histories (e.g. of interaction with the represented entity) for tokens to qualify for content. This is clearest in the case of an enduring conscious preference; such a preference can, after all, be formed on the spot, or, as in the case of swamp man, inherited (along with some conditioned reflexes) as a result of the details of his construction. Our account of “standard use”, then, is a dispositional account, and not an historical one, although it is of course the case that such dispositions (preferences and reflexes) are generally fixed as a result of specific causal histories.

Definition 6: A *circumstance* is a circumstance of the subject. A circumstance consists in the subject’s internal states, including the subject’s bodily changes, registrations, representations, expectations, priorities, values, options for action, homeostatic self-evaluations, procedural knowledge, motor schemas and also the subject’s immediate environment.

The definition of “circumstance”, as mentioned, includes intentional states and tokens. It will be instructive to pause here to explain how the account avoids regress. Note that “circumstances”, as used in definition 5, do not play a content-grounding role in our theory of content. It is a fact that representing tokens are used in circumstances, and that these circumstances consist partly in the presence of other representing tokens. However, the importance of the circumstance is not in grounding the *content* of these tokens, but rather in illuminating and delimiting the conditions under which a subject would deploy them. A token has the content it has just in virtue of the guidance it provides to the subject with respect to an entity; whether and when such a token would be used depends on the circumstances of the subject. On our theory, the other representing tokens defining a circumstance can be independently assigned content according to the analysis of content offered here; or, if and when such tokens are necessary elements of the main action-guiding structure being assigned content, then our claim is that such tokens can be assigned content in virtue of their role in guiding the relevant parts of the overall action under consideration (for further discussion of this issue see section 3.4).

So far we have discussed all of the minor, and two of the major terms in our simple outline of the guidance theory. Italicizing those terms already covered, we have: a given mental *token T* (of some given *type*) represents an *entity E* for *subject S*, just in case *tokens of that type are standardly used by the subject to guide its actions with respect to E*). This leaves the last, and most important clause: what is it to guide action “with respect to E”? The fact that subjects take action with respect to things is what confers content on representations; it is how representations reach from the organism to the

things they stand in for. The guidance theory argues that the intentionality of representation can be grounded in the directedness of action. The central importance of the directedness of action means that it is vital to correctly understand—without regress—what it is for an action to be taken with respect to something. The following two definitions, the first of action, and the second of the clause “with respect to”, tackle this important task.

Definition 7: An *action* can be a motor process or a cognitive process. This yields two clauses in the definition of action:

Definition 7.1: In the case of a motor process, a motor process is an *action* if, and only if, it is activated under control of perceptual/cognitive feedback processes capable of effectively modulating or bringing about changes in the organism or in the world

Definition 7.2: In the case of cognitive processes, a cognitive process is an *action* if, and only if, it is a mental process under intentional control whose results contribute to circumstances (as defined above) used to direct motor processes. A cognitive process is under intentional control if the working of that cognitive process is subject to modification by processes of attention, short-term memory, valuation, assent and dissent, practiced learning, and consciously administered self-criticism and praise.

Definition 8: An action is taken *with respect to an entity E* if, and only if,

- (i) The action is a motor program, and **E** is the focus of the expected change or efforts at control in the world; or
- (ii) The action is a motor program and an assumption of information about **E** is a motivating reason that the given action, rather than some alternative non-**E** involving action, was undertaken; or
- (iii) The action is a cognitive process undertaken to discover or confirm facts, to modify values, or to decide between alternative actions, and an assumption of information about **E** is necessary if the process as a whole is to provide guidance for the subject's actions.

Although definition 7 is meant to be relatively simple and uncontroversial, we can make no such claim for definition 8. The key elements of that definition are the three further terms it introduces—*motivating reason*, *focus*, and *assumption of information*. These terms present the potential for regress and therefore require further discussion, offered in the three following sections.

3.1 Motivating reason

With respect to giving an account of *motivating reason*, we hold only that any analysis of “motivating reason” must be such that it would be applicable to goal-directed behavior of entities that do not have representations at all. For example, it must be of a piece with how we would identify the motivating reasons for why a plant turns toward the sunlight just as we identified the motivating reasons the slime mold moved towards the sunlight.

The plant's behavior is goal-directed behavior even if it is not action in the sense defined above, and the motivating reason for the behavior is to maximize the amount of sunlight available for photosynthesis. Because the plant does not have representations, a correct account of motivating reason cannot appeal to representational content. Any account of motivating reasons must allow for motivating reasons that are non-representational facts and entities, even for entities that possess representations.

We also distinguish motivating reasons from applications of causal force. A child may go to bed early on Christmas Eve to encourage Santa Claus to bring presents, and this may be the child's motivating reason, even though Santa Claus is not capable of applying causal force on the child's mind. A hungry wolf may look for prey and its motivating reason may be a future state of satiety, even if the cause of its behavior is a present internal state.

At its heart, the concept of a motivating reason is deeply tied to concepts of rational interpretation like the one found in Daniel Dennett's description of the intentional stance (Dennett, 1987). We will not describe specific standards of rational interpretation here and wish to purposely leave open from where the standards may come. Standards of rational interpretation may be wholly constituted by a Darwinian selection process or may instead be a kind of non-natural Platonic standard or may be a social construction, and so we take no position on naturalism versus non-naturalism for standards of rational interpretation. Keep in mind, however, that any account of motivating reasons must be sensitive to the systematicity of the modal properties that help to differentiate between an

action and an accidental or insufficiently controlled behavior. For example, the clown who falls accidentally because he was nervous about performing a pratfall he intended to perform, was not driven to perform an action by the definition of action given in this account, and there is therefore no motivating reason for his behavior, even though a non-ideal observer, knowing the clown's intention, might interpret the event as a rational action. That is, we are declining to specify specific standards of rational interpretation here, not because our position allows for or welcomes indeterminacy or multiplicity in the assignment of motivating reasons, but only because this is a complex issue in its own right that deserves separate treatment. What we do say is that these standards, whatever they are and wherever they come from, are measures with broad power and applicability; they can be applied equally to evolved and non-evolved objects, to creations of the imagination, and to creatures of design or creations of randomness.

3.2 Focus

As it is used above, the idea of an action's *focus* is intended to express a functionalist concept. When a subject is performing an action it places itself into a potential feedback loop with its environment. Its purpose is to monitor the result of the action and to plan adjustments to its course of action. Of special importance is that when a subject initiates an action it attempts to establish a causal connection to the focus, or, if the focus is not present in its circumstances, to something that is in its circumstances and that varies systematically and reliably with the focus of the action. As part of initiating the action the subject also primes certain mechanisms for potentially receiving feedback about the

result of the action through this contact. This priming is the setting of an expectation, and it is in such terms that the content of an expectation is to be cashed out.

The existence of an expectation has rich counterfactual consequences. Expectations are attuned to the feedback channels through which the subject expects to receive indicators of the action's results. Through the way its expectations are tuned upon initiating the action, the subject will respond differentially to the feedback received through these channels. Indications matching positive expectations will set off hedonic and other re-enforcement mechanisms, as well as possibly priming further motor programs consistent with the type of positive indication, given the subject's motivating reasons. Indications matching negative expectations will set off complementary kinds of re-enforcement mechanisms. In many cases, even *failure* to receive feedback will initiate appropriate responses, such as subsidiary action intended to connect it to some source of feedback. Given this picture we can define *focus* as follows.

Definition 8.1: The *focus* of an action is the ultimate entity being monitored through the feedback channels taken to provide indications of its status.

As mentioned, a subject may monitor the focus indirectly by monitoring the status of some entity being used as an indicator of facts about the focus. Because indicators are made part of an extended guidance control system, indications about the focus will cause in the subject beliefs, decisions or equivalent states about further appropriate actions or perhaps that action may cease. When the focus is monitored through an indicator the

subject may have an indirect causal connection to the focus or even no causal connection at all.

An example of an indirect causal connection to a focus would be an engineer monitoring a gauge that is itself monitoring engine pressure. In cases like this the focus of the action is distal while the directly monitored indicator is present in the circumstances of the subject. Examples of foci to which there is no causal connection are things like the time of day or a mathematical operation on numbers. To monitor the first we might monitor an indicator like a clock face and to monitor the second we might monitor a progression of numerals manipulated according to established rules. In both of these cases the focus of the action is something that is not present and to which the subject is not even indirectly causally connected, but which can be monitored nevertheless, despite the lack of causal connection, by establishing a connection to something else that can be manipulated to vary systematically with facts about the focus.

We now see three cases: A subject may be monitoring a focus directly and causally because it is part of its circumstances, indirectly and causally through a mediated causal connection with something in its circumstances, or indirectly and non-causally by monitoring an element of its circumstances which is believed to maintain a picturing, co-variation, or other tracking (but not necessarily causal) relationship to the focus.

Identifying the focus of an action in a given case requires establishing the facts about what the subject is monitoring in its circumstances by choosing between these three

cases, a choice which itself requires understanding these facts in terms of the subject's motivating reasons.

3.3 Assumption of information

An *assumption of information* is to be cashed out in terms of facts about the actual operation of the representing subject with respect to its operating environment. Beginning with an example will make the concept easier to grasp. Imagine a computer that is processing the command to print a document. To do this, the computer must determine to which printer it should send its own commands. The focus of the action is the printer. To guide its action, the computer reads several character strings contained on its hard disk, one identifying the printer and others with other information about the printer. These strings guide it regarding where it should send its print commands and what protocol it should use to communicate with the printer.

From the perspective of the guidance theory, here is the key fact: these character strings represent what they do both because of the circumstances in which the computer is reading them and also because of the *assumptions* built into those circumstances. The computer processes the strings *as if* they conveyed information about the printer to which it should send its commands and which communication protocol it should use. There is no regress involved in claiming it makes this assumption because the assumption itself is not a matter of having representational content. There is no representation inside the computer with the content: *I assume that this string has information about the printer.*

Even more strongly, its ability to make an *assumption* of information does not require that the computer actually *possesses* information, nor that it ever did.¹³ The character string it accesses could have been placed on the disk via the output of a random number generator and by coincidence be effective in directing it to the proper printer. Even were that to be true, the string still would be providing guidance and the computer would still be making an assumption that the string contained information about the correct printer.

Note that this is one of the significant points of divergence between the guidance theory and Millikan's biosemantics. For the guidance theory, content is determined by current function, but for Millikan content always depends on *proper* function; and proper function is defined historically. "It is not the facts about how the system *does* operate that make it a representing system and determine what it represents. Rather, it is the facts about what it would be doing if it were operating according to biological norms."

(Millikan, 1993:10-11) Specifically, according to biosemantics, the content of a token is determined by the environmental correspondence that must obtain if the representation consumer is to have its proper function, and proper function is canonically defined in terms of the selection history of the consumer. In contrast, the guidance theory allows for the possibility that a system serving some function other than that for which it was selected, or mal-functioning in some very lucky way, or with some other aberrant causal history, could, in its use of mental tokens, be *representing* just in case (roughly speaking)

¹³ The claim rests on the perhaps somewhat illegitimate assumption that to successfully possess information requires a certain kind of causal history. On weaker requirements, it might perhaps be said that the computer has information about the printer, since the character string in question is useful in guiding its action with respect to the printer. Such a possibility does not affect the main argument of this section.

the mental tokens in question were being used to guide the agent's actions with respect to the indicated entities.

The central point is that an assumption of information (e.g. that the computer has information about the printer) is not a matter of pre-existing representational content, nor is it grounded out in terms of any specific causal history, but is rather a matter of *know-how* that is identifiable by rational interpretation of the performance of the computer: that it takes actions, that these actions are guided, how it accesses representations for that guidance, in what circumstances it accesses them, how it reads and interprets their structure, what actions it initiates and monitors upon accessing them, how those actions cause it to interact with the world, and so forth. Knowledge of the computer's design history is merely useful but not necessary to make judgments about these things. As with the accounts of both *motivating reason* and *focus*, the account of assumption of information is specifically cashed out in terms of the actions of the subject.

3.4 Know-how

We can provide a candidate analysis of this *know-how*. To do this, we first need to define, for any given token, the class of actions it supports. The class of actions a token **T** supports is relative to the kinds of circumstances **C** where the system is prepared to use the token for guidance. It consists of all the actions the system can initiate or modulate in **C** due to its processing of **T**. Let us label this class of supported actions **A_{supp}**.

Definition 9: An action **A** is a member of the class of actions, \mathbf{A}_{supp} , supported by a token **T** used by a subject **S** in circumstances **C** if, and only if, **S** in **C** would use **T** for guidance regarding the initiation or manner of execution of **A**.

We should think of the actions in \mathbf{A}_{supp} as focus-neutral descriptions of an action in need of association with a focus in particular initiations. So, for example, if in some circumstances a system is prepared to use a token for guidance in running, the action *running* is the focus neutral description. If the specific initiation of this action occurs when the focus of the action is a bear, the focus-neutral action “running” is initiated as the focus-specific action “running away from a bear.” Actions obtain a focus in the way discussed above, by determining the ultimate entity being monitored by the subject in support of its action.

Furthermore, since subjects do not initiate actions at random, for each action in \mathbf{A}_{supp} , there will be a (possibly very large but) finite set of circumstances capable of triggering the initiation of the action. We can call this set of triggering circumstances \mathbf{A}_{circ} . The number of triples $\langle \mathbf{A} \in \mathbf{A}_{\text{supp}}, \mathbf{C} \in \mathbf{A}_{\text{circ}}, \mathbf{Focus} \rangle$ representing supported actions **A** initiated in circumstances **C** with focus **Focus** provides a class of counterfactual *action scenarios*, $\mathbf{A}_{\text{scene}}$, in which the token **T** provides guidance for a subject. These are the action scenarios in which **T** *participates*.

Most actions are complex, both in the sense that they have many different specific features that must be managed (e.g., the trajectory and velocity of a running motion), and

in the sense that they almost always require initiating smaller or tangential actions involving entities besides its focus if they are to succeed in affecting their intended change or control. For instance, while running away (main action) from the bear (main focus), one may have to jump over (sub-action) a branch (sub-focus). Because of the complexity of action, subjects needing to execute an action will almost always use representations other than the tokens representing the (main) focus of the (main) action. In fact, activation of these further tokens is necessary to fill out the circumstances in which all the tokens are used.

These other active representations will fall into several categories: conscious representations with foci of their own serving the larger action program; unconscious but potentially conscious representations supporting the interpretation of the circumstances and manner in which the action is executed; and sub-conscious representations that can never be conscious but that provide support for basic perception, adjusting bodily movement, and triggering emotion. In order to assign content to these representations in a consistent and non-circular manner, we construe the entities towards which the supporting tokens provide guidance as sub-foci in sub-actions lying under the umbrella of the main action. Therefore, these further tokens, the ones that support the guidance for the main action within a given $\mathbf{A}_{\text{scene}}$, have functional roles determined by their potential relationships to their own foci within the circumstances \mathbf{C} of $\mathbf{A}_{\text{scene}}$.

Relative to these action scenarios, the guidance theory supposes that in each A_{scene} where an active token succeeds in having reference¹⁴ the token can be mapped to an entity through its functional role with respect to its own focus, identified under the rational constraints associated with assigning motivating reasons to its actions or sub-actions. This supposition is justified because, in providing guidance, a token will make features of itself available to the subject, which the subject can use to differentially control its actions with respect to an entity which is a focus or sub-focus of a given action.

Thus, to summarize the analysis: know-how, as it is involved in an assumption of information, is a question of the way that the subject's decoders and action mechanisms process and/or respond to representations. Like the earlier example of the computer printing a document, it is a matter of how it accesses representations, in what circumstances it accesses them, how it reads and interprets their structure, what actions it initiates and monitors upon accessing them, how those actions cause it to interact with the world, and so forth, understood under the constraints associated with assigning motivating reasons. The general idea is that assumptions of information consist in non-representational facts about how the subject *works*, not in further representational facts about the subject. Although this account is clearly preliminary, it does at least show how the idea of an assumption of information can be interpreted, and used as part of the machinery involved in determining the content of a representation, without initiating a vicious regress or involving circular appeals to representational content. A more exact and certain analysis of this know-how requires discussion and debate by the interested

¹⁴ The concept of error will be defined formally at the end of the current section.

portion of the philosophical community. To stay within this framework, refinements or competing analyses simply must not violate the prohibition against a regress.

We have now defined and discussed all of the central terms necessary to the guidance theory, allowing us to offer the official, formal definition of representation:

Definition 10: A token **T** *tracks* an entity **E** for a subject **S** if, and only if, **T** is standardly used to provide guidance to **S** for taking action with respect to **E**.

On the guidance theory, *representation* is simply tracking in the sense defined above.

Definition 11: A token **T** *represents* an entity **E** for a subject **S** if, and only if, **T** tracks **E** for **S**.

By linking content to guidance in this way, the guidance theory distributes responsibility for the existence of representational content across a representational token (the representation) and an interpretative decoding mechanism (the decoder) integrated with a subject's action-determining processes. The effect of distributing responsibility is to introduce new degrees of freedom regarding the exact physical or informational requirements for something to be a representation, as the requirements on the representation will depend on the capabilities of the decoder and the circumstances in which it is used. In general, the demands on each part of the coupled system vary inversely with the demands on the other. A representation that is highly structured and

closely coupled with what it represents needs a less sophisticated decoding mechanism, while a very sophisticated (or very rigid and simple) decoding mechanism may embody (or presume) so much implicit domain knowledge that it can get by with very sparse representations.

The requirements on a representation can also vary greatly depending on the features of what is represented. For example, represented items that do not change, or change very slowly relative to the lifetime of the representation, can be represented by tokens that are not causally connected to their content because the decoding mechanisms can extract guidance by implicitly relying on the stability of the content.

3.5 Error

One of the most important problems that any theory of representation must solve is the problem of normativity: representations are assessable for accuracy, and therefore they can be in error. To be complete, the guidance theory must account for this feature of representations. Because the guidance theory is an action-based theory of representation, the natural thing to do is to base error on the failure of action and the way that a representation's guidance contributes to that failure.¹⁵ The intuitive idea, then, is that a representation is in error if it provided guidance to an action that failed in its intent, and it

¹⁵ It is because the guidance theory defines content in terms of action guidance, and thus can define representational error in terms of failure of action, that it can dispense with the difficult notions of *proper* function (Millikan, 1984; 1993), or *normal* circumstances (Dretske, 1981), generally required to account for the possibility of representational error. One advantage of an account of error based on failure of action is that it meets Bickhard's meta-epistemological requirement that representational error be potentially detectable by the representing system itself (Bickhard 1993, 1999).

failed partly or wholly because of the guidance provided by that representation. This intuitive idea can be formalized as follows.

Definition 12: An action **A** *fails in its intent* if, and only if,

- (i) **A** is a motor action and the intended change is not achieved or the intended process is not brought under control (as identified by functionally specified states of expectation satisfaction); or
- (ii) **A** is a cognitive process and it
 - a. confirms a representation that is in error¹⁶; or
 - b. disconfirms a representation that is not in error; or
 - c. modifies a value in a way that the subject later regrets; or
 - d. recommends a course of action that fails.

Definition 13: An action **A** fails in its intent *because of R* if, and only if,

- (i) **A** failed; and
- (ii) **A** was taken with respect to an entity **E**; and
- (iii) **R** provided guidance for **A** with respect to **E**; and
- (iv) **R** has feature **F**; and
- (v) **R** with **F** represents that **E** has property **P**; and
- (vi) **A** failed because **E** was not **P**.

Definition 14.1: A token representation **R** *is in error* for subject **S** in token circumstances **C** for action **A** if, and only if **A** would fail because of **R** if taken by **S** in **C**.

¹⁶ This clause in the definition is an embedded recursion, not a circularity.

Definition 14.2: A token representation **R** is in error for **S** *simpliciter* if, and only if, the class of actions for which **R** provides guidance in **S**'s circumstances **C** is dominated by actions that would fail because of **R**.

The definitions and explanation of terms above give the content of the guidance theory of representation. Applying the theory to a few standard examples of mental representations in action will help to make clear how the definitions work in concert to identify representations and their content.

4. Illustrations of the guidance theory

4.1 Stopping at a red light – Janice is driving her car. As she approaches an intersection she notices the light changing from green to yellow to red. She stops her car. In this case she has a perceptual representation of the light turning red. According to the guidance theory, what is it that makes her percept a representation of the state of the traffic light? It is a representation because such percepts standardly provide guidance for her driving actions. The percept is a representation of the traffic light's being red because, in this case,

- (i) *Circumstances* – Her token representation is active in circumstances that include her driving, her knowledge of the relevant practices, her motor skills that make her able to drive competently, and her wish to avoid tickets and remain safe.

- (ii) *Action* – In these circumstances, her stopping is an action according to the motor control clause of the definition of action;
- (iii) *With Respect To* – Her action is with respect to the traffic light by clause (ii) of the definition of “being with respect to” because an assumption of information about the traffic light was a motivating reason for stopping the car; and
- (iv) *Providing Guidance* – The token provided guidance with respect to the traffic light in those circumstances because features of the token were used to discriminate between the possible actions with respect to the traffic light. For example, one of the features of the token was whether it contained a token representation of red or not.

4.2 Solving a math problem – Young Anna is learning her arithmetic. She is told that Ted has two apples and that Billy gives him three more apples. She is asked how many apples Ted has after he has received the apples from Billy. She writes down the problem as $2+3=?$ and then works out the solution on her fingers. According to the guidance theory, what is it that makes Anna’s fingers representations of the numbers 2, 3 and 5? They are representations because they standardly provide guidance to her for her arithmetic reasoning. Her fingers are representations of numbers because, in this case,

- (i) *Circumstances* – Her token representation is active in circumstances that include, among other things, her knowledge that she has been assigned an arithmetic problem, her desire to find an answer to the problem, her

knowledge of how to count on her fingers, and her previously used ability to manipulate her fingers to help with small addition problems.

- (ii) *Action* – In these circumstances, her working out of the arithmetic problem is an action according to the cognitive process clause of the definition of action;
- (iii) *With Respect To* – Her action is with respect to the numbers 2, 3, and 5 by clause (ii) of the definition of “being with respect to” because an assumption that her fingers gave her information about the numbers 2, 3, and 5 was a motivating reason for her counting on her fingers; and
- (iv) *Providing Guidance* – The token provided guidance with respect to the numbers 2, 3, and 5 in those circumstances because features of the token were used to discriminate between the possible solutions to the arithmetic problem, each of which represents an alternative course of action as defined in the cognitive process clause in the definition of *action*. For example, features of the token include the fingers she has pointing up-vs-those pointing towards her palm.

4.3 Non-existent and fictional entities – One way to attach a traditional frame of understanding to the guidance theory is to use centered possible worlds to understand representational content as defined so far. A centered possible world is a possible world with a specific individual at a specific place and time marked as its center (Chalmers, 2003). We can use centered possible worlds to define a kind of representational content that will help us to make sense of representations of non-existent and fictional entities.

First, we can say that the *potential* representational content of a representation **R** for a subject **S** in circumstances **C** is the set of entities it *would* represent within centered possible worlds *were* an entity with **R** in **C** present at the center of the world (where **C** is a type of circumstance if the representation is a type and **C** is a token circumstance if the representation is a token). A representation's potential content clearly includes entities that only exist in merely possible worlds. The *actual* content of the representation is the content it has in a set of centered possible worlds relevantly similar to the centered actual world where the representation is giving guidance.

The actual content of a representation is therefore a subset of its potential content: It is just that part of the representation's potential content that is left once we apply a similarity relation as our access relation between the actual world centered on the subject possessing **R** and other centered possible worlds where **R**'s guidance could be evaluated. Because similarity is more relaxed than identity, a similarity relation of this sort is clearly liberal with respect to whether the other centered possible worlds can contain entities that do not exist in the actual world. Given this,

Definition 14: A representation **R** in circumstances **C** has *non-empty content* if, and only if, there is some center within some possible world where **R** in **C** would not be in error.

Definition 15: A representation **R** in circumstances **C** represents a non-existent entity **E** if, and only if,

- (i) The representation **R** in **C** has non-empty content; and
- (ii) In every centered world relevantly similar to the actual world where **R** in **C** has content, **R** represents **E**; and
- (iii) **R** in **C** would be in error with respect to the presence of the entity **E** at every relevantly similar center within the actual world.

To illustrate: Imagine a subject **S** in the actual world has a representation **R** in **C** of an alleged planet between Mercury and the sun. We will say this is a representation of Vulcan and the location of the subject provides a center for the actual world in which **R** is providing guidance.

Does **R** represent a non-existent entity? Consider clause (i): **R** does not have empty content, as intuitively there is some centered world where **R** in **C** would succeed in providing guidance (after all, **R** does not seem to be self-contradictory). Consider clause (ii): Because **R** in **C** provides guidance appropriate only with respect to a planet between Mercury and the sun, it also seems intuitive that in every centered possible world relevantly similar to the centered actual world and where **R** is not in error, **R**'s guidance succeeds because Vulcan exists in that centered possible world. Therefore **R** represents Vulcan. Consider clause (iii): **R** in **C** would be in error with respect to the presence of Vulcan at every possible center within the *actual* world (since Vulcan does not exist). Clauses (i)-(iii) are all satisfied, therefore **R** represents a non-existent entity.

Human beings enjoy novels, plays, movies and stories of all types. These stories include fictional characters, places, events and times. The way we understand fiction certainly involves having representations of fictional things even though we know those things do not exist and even though those representations do not seem to provide guidance for action with respect to those things. How can these be representations if they do not provide guidance for action with respect to the entities they represent?

Representations of fictional entities are ordinary representations that are in error, and they are capable of providing guidance for our actions. They differ from ordinary representations only in the circumstances within which they are used: our action systems are reset to not respond to their guidance. In other words, the representations of fictional entities are capable of providing guidance and would do so *if the subject were to treat the entities as non-fictional*. However, the fictional character of the representations is flagged somehow and handled through special processing filters on the subject's action systems. These filters dampen our motor control systems with respect to the guidance provided by those representations.

4.4 Abstract entities and causal interaction – Abstract entities like the number two or time or the proposition that all men are created equal or the possibility that pigs could fly provide challenges for some theories of representation. The main challenge is that these entities are not, or at least do not seem to be, in the causal network. Because they can have no causal effects on our brain states or mind states, it can seem like magic that we are able to have representations of them (much less accurate representations of them).

How does the guidance theory account for the fact that we can have representations of these acausal entities?

The definitions in the guidance theory do not contain any essential condition of causal interaction with the represented entity. The theory only requires that the representation provide guidance to a subject for actions with respect to what it represents, and remains agnostic about how a representation may come to have its guidance providing features.

For instance, in the case where a representation is providing guidance with respect to an entity, and the actions participated in by the representation have the represented entity as a focus, all that is required is that the subject is tracking the represented entity through a feedback channel. These feedback channels can be mediated by proxies that track the represented domain (e.g., an abacus or number system in the case of mathematics) and do not have to causally connect to the represented entity, so there is no barrier in principle to an acausal entity being the focus of an action. In practice many different methods could be effective for different kinds of representations under different kinds of circumstances. Candidates other than causation include evolution and the exploitation of found isomorphisms.

In considering the case for evolution, we can presume that entities, such as the number two, and its domain, such as mathematics, present necessary conditions on the structure of the world and our experience of it. Presume also that the mathematical domain itself has structure worth knowing, and it is a small leap to believe that there are obvious

evolutionary advantages for creatures able to track those entities and the structure of their domain. If a possible mechanism for tracking abstract entities arises during a species' evolutionary history, perhaps as something providing supplementary guidance to causally based representations already in place, it is not hard to imagine evolution exploiting and selecting for those mechanisms. As previously mentioned, integrated control systems are kinds of hypothesis machines, built to make assumptions of information, and they work the same whether or not there is a true causal connection to the things about which they assume information. The only presumption required is that the naturally arising representational vehicles give proper guidance in the proper circumstances often enough to be useful, and therefore become selected for. One way they can do this is just by having features isomorphic to the entities and domains they track, and by making those features available to the proper action controlling systems for exploitation in the proper circumstances (e.g., when responding to differences in number is important). Random variation and selection can assure these kinds of features without a causal relation between representation and represented entity.

Notice that this view does not make a subject's evolutionary or design history definitionally responsible for representational content itself. The responsible features exist locally to a creature, so selection is not what makes a representation

representational.¹⁷ Rather, representations are just like other engineering features of a

¹⁷ This highlights an important distinction between the guidance theory and Millikan's biosemantics. For Millikan, content can be determined by identifying the environmental correspondence exploited by natural selection in the functional development of a given representation consuming system. For the guidance theory, it is a matter of identifying with respect to which entity a subject's actions are being guided. These two methods of fixing content will sometimes, but not always, give the same result. See (Anderson, forthcoming-a) for further discussion. Note the implication, however, that whereas Millikan advocates a direct and prominent role for evolutionary history in determining content, on the guidance theory

creature, such as its eyes or wings or lungs, and selection plays a role only in ensuring the persistence of those features if they arise, and perhaps their further articulation and growing sophistication through later generations.

Consider also the role that exploiting found isomorphisms may play. For example, the calendar for 2003 on the office wall is used to represent the days, weeks, and months of the year. The year 2003 is an abstract entity. Because the calendar was produced prior to the start of the year, it is difficult to argue that the year 2003 is causally responsible for the calendar. Also, because both the year and the calendar are essentially static objects, it is hard to argue that they co-vary. The feature the calendar has that makes it effective as a representation (i.e., the thing in virtue of which it can provide guidance to me) is that it has an isomorphic structure to the year along the dimensions of days, weeks, and months. Similar truths hold about our external representations of mathematical systems and our representations of things like space or time. In these cases the sufficient conditions for something to be a representation seem simply to be, (i) that there be a useful isomorphism between it and the represented thing; and (ii) that the subject possess decoders that are in a position to exploit this isomorphism to derive guidance for action with respect to the represented thing. Causal relations are entirely inessential to these conditions, as useful isomorphisms, and mechanisms for exploiting them, can arise through evolution, by chance, or even by divine intervention, and still provide guidance. Finally, although an

evolutionary history exerts only an indirect effect on representational content. It is not evolutionary history *per se* that determines content, but the function of a representation in guiding action. Since it is in virtue of their role in guiding action that the elements of an organism's cognitive systems are primarily exposed to selection pressures, this seems the proper place to locate the influence of evolutionary history on their structure and content.

isomorphism between representation and represented thing is required, no *specific* kind of isomorphism is required. The nature of the isomorphism can vary from case to case, and depends entirely on the abilities and needs of the representation using mechanisms for decoding and exploiting it for guidance within the circumstances where it is used.

5. Conclusion

This essay introduced the guidance theory of representation. The guidance theory is an action-focused theory of representational content. According to the guidance theory, representational content is derived from the role a representational vehicle plays in guiding a subject's actions with respect to other things. What qualifies an element of experience as a representation is, strictly speaking, only that the element of experience be capable of providing a subject with guidance for its actions with respect to entities. To be capable of providing guidance the further conditions an element of experience must meet are flexible and very generic. The element of experience only needs to have features useful for exploitation by the subject's action-producing mechanisms. As we have seen, the key features enabling exploitation are that it can be usefully interpreted, so there is required to be an exploitable feature of the representation relative to the entity it represents (and the decoder), but the constraints on the form of the feature are determined by the abilities and needs of the subject. A representation must also be usable in conjunction with other representations that might be required to extract guidance from it, but that too is largely a function of the abilities of the subject.

We have shown that the guidance theory can account for various problem cases for representational content such as abstract, fictional and non-existent objects. A more detailed comparison of the guidance theory to some alternatives, including Millikan's biosemantics, can be found in (Anderson, forthcoming-a). Future work will consider the implications of the guidance theory for the correspondence theory of truth and the differing cognitive significance of co-referring terms, to reference and self-reference (Anderson and Perlis, in press), to scientific realism (Anderson, forthcoming-b; O'Donovan-Anderson, 1996), and to consciousness and phenomenal content (Rosenberg, 2004). We will also consider the evolutionary development of representation in more detail, and the specific question of the teleological requirements on representations.

Bibliography

Akins, Kathleen (1996) Of sensory systems and the "aboutness" of mental states. *J. Phil.* 93: 337-72.

Anderson, M. L. (2003) Embodied cognition: A field guide. *Artificial Intelligence* 149(1): 91-130.

Anderson, M.L. (forthcoming-a) Representation, evolution and embodiment. In: D. Smith, ed. *Evolutionary Biology and the Central Problems of Cognitive Science*.

Anderson, M. L. (forthcoming-b) Cognitive science and epistemic openness. *Phenomenology and the Cognitive Sciences*.

Anderson, M. L. and Perlis, D. (in press) The roots of self-awareness. *Phenomenology and the Cognitive Sciences*, 4.

Barlow, H. (1999) Feature detectors. In: R. A. Wilson and F. C. Keil, eds. *The MIT Encyclopedia of the Cognitive Sciences*. pp. 311-13. Cambridge, MA: MIT Press.

Bickhard, M. H. (1993) Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5, 285-333.

Bickhard, M.H. (1999) Interaction and representation. *Theory & Psychology*, 9(4), 435-458.

- Block, N. (1986) Advertisement for a semantics for psychology. In: P. French, T. Uehling and H. Wettstein, eds. *Midwest Studies in Philosophy X*: 615-678. Minneapolis: University of Minnesota Press.
- Bonner, J. T. (1985) *The Cellular Slime Molds*, 2d ed.
- Bonner, J. T. (1994) The migration stage of *Dictyostelium*: behavior without muscles or nerves. *FEMS Microbiol. Lett.* 120: 1-8.
- Chalmers, D. (2003) The Nature of Narrow Content. *Philosophical Issues*, 13:46-66.
- Deban, S. M., O'Reilly, J. C. and Nishikawa, K. C. (2001) The evolution of the motor control of feeding in amphibians. *Amer. Zool.* 41: 1280-1298.
- Dennett, D.C. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.
- Dretske, F. (1981) *Knowledge and the flow of information*. Cambridge, MA, MIT Press.
- Dretske, Fred (1986) Misrepresentation. In: R. Bogdan, ed. *Belief: Form, Content, and Function*. New York: Oxford University Press.
- Dretske, F. (1988) *Explaining behavior*. Cambridge, MA: MIT Press.
- Ewert, J. -P. (1987) Neuroethology of releasing mechanisms: prey-catching in toads. *Behav. Brain Sci.* 10:337-405.
- Finke, R., Ward, T. and Smith S. (1992) *Creative Cognition*. Cambridge, MA: MIT Press.
- Fodor, Jerry A. (1981) *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge, MA: MIT Press.
- Fodor, J. (1987) *Psychosemantics*. Cambridge, MA: MIT Press.
- Gallistel, C. R. (1999) The replacement of general-purpose learning models with adaptively specialized learning modules. In M.S. Gazzaniga, Ed., The Cognitive Neurosciences 2d ed. (1179-1191). Cambridge, MA: MIT Press.
- Grush, R. (1997) The architecture of representation. *Philosophical Psychology* 10(1): 5-25.
- Grush, R. (2003) In defense of some 'Cartesian' assumptions concerning the brain and its operation. *Biology and Philosophy* 18(1): 53-93.
- Grush, R. (2004) The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences* 27(3): 377-396.

- Häder, D. P. and Burkart, U. (1983) Optical properties of *Dictyostelium discoideum* pseudoplasmodia responsible for phototactic orientation. *Exp. Mycol.* 7: 1-8.
- Harman, G. (1982) Conceptual role semantics. *Notre Dame Journal of Formal Logic*, 23, 242-56.
- Harman, G. (1987) (Nonsolopsistic) conceptual role semantics. In E. LePore (Ed.), *Semantics of Natural Language*. New York, Academic Press.
- Ingle, D. J. (1973) Two visual systems in the frog. *Science* 181: 1053-55.
- Ingle, D. J. (1980) Some effects of pretectum lesions on the frog's detection of stationary objects. *Behavioural Brain Research* 1:139-63
- Ingle, D. J. (1982) Organization of visuomotor behaviors in vertebrates. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield, eds., *Analysis of visual behavior*, MIT Press.
- Ingle, D. J. (1991) Functions of subcortical visual systems in vertebrates and the evolution of higher visual mechanisms. In: R.L. Gregory, J. Cronly-Dillon (Eds.), *Vision and Visual Dysfunction, Vol. 2, Evolution of the Eye and Visual System*. London. Macmillan, Basingstoke, pp. 152-164.
- Lettvin, J., Maturana, H., McCulloch, W., and W. Pitts. (1959) What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers*, 47: 1940-1951.
- Loar, B. (1981) *Mind and meaning*. London: Cambridge University Press.
- Lycan, W. (1984) *Logical form in natural language*. Cambridge, MA: MIT Press.
- Millikan, Ruth G. (1984) *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, Ruth G. (1993) *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.
- Milner, A. D. and Goodale, M. A. (1995) *The Visual Brain in Action*. Oxford: Oxford University Press.
- Miura, K. and Siegert, F. (2000) Light affects cAMP signaling and cell movement activity in *Dictyostelium discoideum*. *Proc. Natl. Acad. Sci. USA* 97: 2111-2116.
- O'Donovan-Anderson, M. (1996) Science and things: On scientific method as embodied access to the world. In: M. O'Donovan-Anderson, ed. *The Incorporated Self: Interdisciplinary Perspectives on Embodiment*. Lanham, MD: Rowman and Littlefield Publishers.

Poff, K. L. and Loomis, W. F. (1973) Control of phototactic migration in *Dictyostelium discoideum*. *Exp. Cell Res.* 82: 236-40.

Rosenberg, G. (2004) *A Place For Consciousness: The Theory of Natural Individuals*. Oxford: Oxford University Press.

Sternfeld, J. and David, C. N. (1981) Cell sorting during pattern formation in *Dictyostelium*. *Differentiation* 20: 2-21.

Thompson, E. (1996) The mindful body: Embodiment and cognitive science. In: M. O'Donovan-Anderson, ed. *The Incorporated Self: Interdisciplinary Perspectives on Embodiment*. Lanham, MD: Rowman and Littlefield Publishers.

Traynor, D., Kessin, R. H., and Williams, J.G. (1992) Chemotactic sorting to cAMP in the multicellular stages of *Dictyostelium* development. *Proc. Natl. Acad. Sci. USA* 89: 8303-8307.