Haugeland on Representation and Intentionality

In Philosophy of Mental Representation. H. Clapin, ed. Oxford: Oxford University Press, 2002, pp. 122–138.

Robert Cummins
University of Illinois
rcummins@uiuc.edu

Haugeland doesn't have what I would call a theory of mental representation. Indeed, it isn't clear that he believes there is such a thing. But he does have a theory of intentionality and a correlative theory of objectivity, and it is this material that I will be discussing in what follows.

It will facilitate the discussion that follows to have at hand some distinctions and accompanying terminology I introduced in Representations, Targets and Attitudes (Cummins, 1996. RTA hereafter.). Couching the discussion in these terms will, I hope, help to identify points of agreement and disagreement between Haugaland and myself. In RTA, I distinguished between the target a representation has on a given occasion of its application, and its content. RTA takes representation deployment to be the business of intenders: mechanisms whose business it is to represent some particular class of targets. Thus, on standard stories about speech perception, there is a mechanism (called a parser) whose business it is to represent the phrase structure of the linguistic input currently being processed. When this intender passes a representation R to the consumers of its products, those consumers will take R to be a representation of the phrase structure of the current input.

There is no explicit vocabulary to mark the target-content distinction in ordinary language. Expressions like "what I referred to," "what I meant," and the like, are ambiguous. Sometimes they mean targets, sometimes contents. Consider the following dialogue:

You: I used this map and got around the city with no problem.
Me: Which city do you mean? (or: Which city are you referring to?)

Here, I am asking for your target, the city against which the map's accuracy is to be measured. Now compare this exchange:

You: Here is the city-map I was telling you about.
Me: Which city do you mean? (or: Which city are you referring to?)

Here, I am asking for the content, i.e., for the city the map actually represents.[1] The potential for ambiguity derives from the fact that there are two senses in which representations are semantically related to the world, the content sense and the target sense, but the distinction is not marked in ordinary vocabulary. The distinction is there, as the example just rehearsed makes clear, but it is marked only by different uses of the same expressions in ordinary language.

Representational error, then, is a mismatch between the representation produced and its target on the occasion in question. The obvious way to think of representational error, as Jerry Fodor once said to me in conversation, is this: Error occurs when a representation is applied to something it is not true of, e.g., when one applies a representation of a horse to a cow. The distinction in this formulation between what a representation is applied to and what it is true of is precisely the distinction between a representation's target and its content. The crucial point is that what determines what a representation is true of must be independent of what determines what it is applied to, otherwise error is problematic. It follows from this that a theory of representational content - a theory that says what it is for R to be true of T - is only part of the story about representation. We require, in addition, a theory of target fixation, a theory that says what it is for R to be applied to T. Since the target of tokening a representation is, as it were, the thing the representation is intended to represent, I shall say that representations represent their contents, but that a use of a representation intends its target. Intentionality is thus different from semantic content. The former is part of the theory of targets, while the latter is part of the theory of representational content. The intentional content of r is therefore not the actual content of r at all, but, rather, the intended content of some use of r.[2]

I construe Haugeland's theory of intentionality and objectivity as a theory about targets. The theory has two internally related parts:

---

[1]The standard Twin-Earth cases simply trade on this ambiguity in semantic terms. See RTA, 126ff.

[2]"Intend" is a technical term here. I do not, of course, suppose that cognitive systems generally intend their targets consciously, i.e., that whenever t is the target of a use of r in ∑, ∑ forms the intention to use r to represent t. But I do think the technical sense is a natural extension of this usual sense. In the case of conscious, deliberate use, intended content is quite literally the content one intends to represent. As always, one may not succeed in doing what one intends, hence one may fail to represent what one intends to represent.

(1) Intentionality: how are targets fixed--i.e., what determines the target of a representation on a given occasion of its use?

(2) Objectivity: what targets it is possible for a given system to have?

A target of a representing is, by definition, normative for the accuracy, on that occasion, of the representation deployed. If you have targets that you cannot represent accurately, you get what I called forced error in RTA. Accuracy and error come in degrees, of course. A certain amount of error is tolerable, sometimes even desirable, as when accuracy is traded for tractability or speed. Moreover, a given representation may be an accurate representation of one target and an inaccurate representation of another. Accuracy, then, is not intrinsic to a given representation. Indeed, representational types are not accurate or inaccurate. Only particular deployments of a representational type can be accurate or inaccurate, and the degree of accuracy is just the degree of match between the representation deployed and the target at which it is aimed. To understand objectivity, then, we must understand how it is possible for a representing to be aimed at some particular target, and this is what makes intentionality and objectivity two sides of the same coin. To understand one is to understand the other.

Haugeland emphasizes that, in this inquiry, we must not take the objects for granted, and I agree. To see why, consider parsers once again. A parser is an intender whose function is to represent the phrase structure of the current linguistic input. Targets, then, are fixed by the representational function or functions of the intender-mechanisms that produce/deploy them. So, when we ask what targets a system can have, we are asking what representational functions it can have. Whatever your favorite account of functions, it is not going to be the function of any canid intender to represent chess positions. It follows that canids cannot have chess positions as representational targets, and hence that no chess position as such is normative for the accuracy of any canid representation deployment. In this sense, then, chess positions are not objects for canids. For exactly analogous reasons, it cannot be the function of any sub-cultural intender of which we have any experience to represent positrons. What representational functions something can have evidently depends in part on its conceptual sophistication and on its perceptual and inferential resources. This, then, is why the theory of intentionality cannot take the objects for granted. Nothing can be an object for me if I cannot have it as a target. I cannot have it as a target unless it can be the function of one of my intenders to represent it. But which intenders I can have depends, at least in part, on my sophistication

and resources. How I am designed--my functional analysis--puts constraints on what can be an object for me.

Unlike representational content, function (and hence target fixation) is holistic in the sense that functions arise out of the organization of some containing system.[3] So, when Haugeland asks what perceptions are perceptions of, or what our thoughts are thoughts about, he is, I think, thinking of what I call target fixation. And he is right that target fixation is holistic. You cannot have chess positions as such as targets unless you have the sort of functional organization that supports an intender whose job is to represent chess positions. About intentionality and its relation to objectivity, then, we are in substantial agreement about the basics.

What about the details? Here, I am less confident. I propose to look at three different areas where I think that I must have misunderstood his position, or one of us must have misunderstood something about the phenomena themselves.

Seeing

I am pretty sure that Sheila, our canny Australian Shepherd mix, can see can openers and tennis balls. I am also pretty sure that it cannot be the function of any canid intender to represent tennis balls or can openers. How is this dissonance to be harmonized?

It seems impossible to explain how Sheila could chase and catch a thrown tennis ball if she could not see it. But how can it make sense to say she sees it if it cannot be a target for her? Seeing is normative, an accomplishment. 'See', as we used to say, is a success verb. If tennis balls cannot be targets for her, then they cannot be norms against which the accuracy of her perceptual representations are measured.

The overwhelmingly tempting move here is to say something like this: It certainly is the function of one of Sheila's intenders to represent the currently foveated thing, understood here as a coherent physical lump, and this thing could well be a tennis ball or a can opener. So, though she sees the tennis ball, she doesn't see it as a tennis ball. Realistically, she doesn't just see it as a

---

[3]Indeed, it is my view that holism in the theory of meaning is a response to the holism of target fixation. Since reference fixing, in one sense, is target fixing, and since Davidson (1967) turned the theory of meaning into the theory of reference, target fixation can look like the whole story about meaning. If you squint and don't move.

coherent lump, either, since she can distinguish tennis balls from balls of other sorts of the same size and color, and will pick one out of a pile of assorted balls when asked to get the tennis ball. So she sees it as a ball (sphere?) of a distinct sort, pretty much as do young children who haven't a clue about tennis, but recognize the characteristic size, texture, and characteristic seam pattern of two interlocking bones wrapped around a sphere. In just this way, non-chess players can recognize standardized chess pieces. Because of this, terms like 'tennis ball' and 'chess piece' are ambiguous. There is a functional sense in which a salt shaker can be the white queen (or any other piece) in a chess game provided it is played (used) like one, and any ball of roughly the right size and elasticity can be a tennis-ball in a game of tennis provided it is played (used) like one. But there is another sense in which these terms pick out objects with characteristic shapes and other intrinsic properties because these are the standard occupants of the functional roles in question. Sheila can see tennis balls in this latter sense, but not in the functional game-embedded sense.

The natural moral of this story seems to be that all seeing is seeing-as. Since seeing is an accomplishment, it must have an intentional object--a target--against which its accuracy is measured.[4] For a tennis ball, as such--i.e., taking the term functionally--to be a perceptual target, you have to be able to see it as a tennis ball. Hence, the only things against which the accuracy of your seeing can be measures are things that can be seen as something or other. Since all seeing is normative, if Sheila doesn't see the tennis ball as a tennis ball, she must see it as something else if she sees it at all. So all seeing is seeing-as. QED

So what? It depends on what is built into seeing-as. Or better, it all depends on what it takes to have tennis-balls and chess positions as targets. The capacity to see something as a tennis ball is just the capacity to have tennis balls as such as targets. So the question is what it takes to have an intender that has the function of representing tennis-balls as such. Here, I think,

---

[4]Representational accuracy should not be confused with perceptual or cognitive or behavioral success. Successful perception needn't be accurate perception. Indeed, to reiterate a point made earlier, accuracy can be expensive, hence undermine success. This means that no simple behavioral test is a test of representational accuracy. It also means that cognitive science has an extra degree of freedom to deploy in explaining successes and failures. In particular, it can distinguish between those failures that are due to the representations, and those that are due to an inability to exploit them. I will return to this point at length below.

Haugeland and I part company. I think his account over-intellectualizes target fixation and underestimates forced error. He writes:

> "Surely no creature or system can see a given configuration as a knight fork without having some sense of what a knight fork is. To put in a familiar but perhaps misleading terminology, nothing can apply a concept unless it has that concept." (Haugeland, 1996, p. 247)

The respect in which Haugeland thinks this remark might be misleading is that it might suggest to some readers that seeing-as requires linguistic abilities. I am not even slightly tempted to think that having a concept — even a concept like the concept of a knight fork — requires having linguistic abilities, so I am not worried about being misled in this respect. My worry is, rather, the claim that having a knight fork as a perceptual target (or a target of thought, for that matter) requires having a concept of a knight fork, where this is meant to imply having at least a minimal knowledge, though perhaps not explicit knowledge, of chess in general and knight forks in particular.

To get a sense of why this worries me, consider a case of learning. What Haugeland's account rules out is a scenario like this: Before learning, you have t as a target but do not know enough about it to successfully hit it. The errors you make, however, drive incremental improvements in aim, as back propagation does in connectionist learning. If this is possible, you must be able have targets you know next to nothing about. If the targets you could have were constrained by how much you knew about them, the size of the error would not decrease as you learned more. Indeed, your early shots would not be misses at all, since they wouldn't be aimed at the target you ultimately acquire.

Notice that I am not complaining about circularity here. It isn't that you need to know <u>about</u> t to have t as a target. I find this suspicious, all right, but I am sure some boot-strapping story, or co-constitution story, could leave me speechless, if not convinced. My complaint is rather that the account makes it impossible to articulate what it is that drives the learning process. NetTalk, for example, begins by knowing nothing about the phonetic values of letters in context (Sejnowski and Rosenberg, 1987). It learns the mapping by adjusting its weights as a function of the difference between its "guesses" and the correct values. But this means that the correct values are targets right from the start, since it is the correct values against which the accuracy of the guesses are measured. Whether or not you think NetTalk is a good model of how letter-to-phoneme mappings are learned is not relevant here. If you think <u>any</u> perceptual or recognitional learning proceeds by adjusting one's knowledge to reduce the mismatch between one's targets and one's representational attempts, you have to have an account of target fixation that largely decouples it from what you

know about the targets.[5]

Functions

How can systems have or acquire targets about which they know next to nothing? This is a question about representational functions, and hence depends to some extent on the theory of functions. I will discuss just two: the teleological theory, and my own.

On the teleological account, an intender has the function of representing a T (e.g., the phrase-structure of the current linguistic input) if it was selected for, historically, because, on various (perhaps rare) occasions, it did represent Ts accurately enough. How accurate is accurate enough? Enough to get selected. The point is that the intender got selected because, often enough (which might not be very often) it got however close it did (which needn't have been very close) to an accurate representation of Ts. Ts are the targets, on this account, because it is Ts as such, not something else, that are the standard you need to look at to see why the relevant intender was selected.

On this account it is plain that little if any knowledge of Ts is required to have them as targets. What is required is just that Ts were a good thing to get represented now and again, that the capacity to do that be heritable, and that competing lines didn't do it, or something even more important, better. Knight forks are bad candidates for canid targets because chess didn't figure in their evolutionary history, not because they don't have a clue about chess.

On my account of functions, an intender has the function of representing Ts if its capacity to represent Ts figures in a functional analysis of some capacity of a containing system. Consider, then, whatever complex capacity C

---

[5]A similar point can be made about representational contents. If you think, as all use theories (causal theories, conceptual role theories, selectionist/adaptationist theories) must, that there can be no more content in your representations than you can exploit, then you will be at a loss to even articulate the process whereby a perceptual system learns to exploit the information in its proximal representations to identify distal stimuli. Think now of the input end of NetTalk. Its input vectors carry information about the phonetic value of the targeted letter. Its job is to learn how to exploit this information. Since it cannot learn to exploit information that isn't there, it must be there prior to any ability to exploit it. Hence, input content must be independent of the system's sensitivity to it, contrary to all use theories of content. (See Cummins, 2006)

of the whole organism made it more fit than the competition in the teleological scenario. The capacity to represent Ts (to some degree of accuracy, with some reliability) figures in an analysis of C. And since teleological account didn't require knowledge of Ts, neither does mine.

Commitment

Let's rehearse the plot. Haugeland and I agree that what targets a system can have — what it can intend — and, hence, what can be objects for it, depends, in part, and in a more or less holistic way, on the sophistication and resources of the system. So an account of intentionality and an account of objectivity are just two sides of the same coin. So far, we are both Kantians. In my view, however, Haugeland's account of intentionality over-intellectualizes intentionality by assuming that you cannot have targets you know next to nothing about, and under-estimates the power of biological functions to make room for systematic error and therefore to ground targets and hence objectivity.

Given that Haugeland doesn't think biological function can ground targets, and hence objectivity, what does he think intentionality is grounded in? In a word, commitment; vigilant, resilient commitment. To see how this is supposed to work, we need to return to the chess example.

Nothing is a knight-fork except in the context of a chess game. Nothing is a chess-game unless the rules are in force. What enforces the rules? Nothing but a commitment on the part of the players not to tolerate deviations. You get a chess game, not when the rules are satisfied, but when the players take responsibility for their satisfaction.[6] Thus, there are no knight-forks unless there

---

[6]Compare this passage from William Alston's Philosophy of Language.

> " If we set out to analyze the concept of a serve in tennis, the problems we encounter will be very similar to those we have just discussed [in connection with the analysis of the illocutionary act of requesting someone to close the door]. To serve is not just to make certain physical movements, even given certain external circumstances. (I can be standing at the baseline of a tennis court, swinging a racket so that it makes contact with a ball in such a way as to propel it into the diagonally opposite forecourt, without it being the case that I am serving. I may just be practicing.) Nor are any specific effects required. A shot can have widely varying effects — it can inspire one's opponent with fear, despair, exultation, contempt, or boredom; these variations, however do not keep it from being true that one was serving in all these cases. Then what does

are players taking responsibility for the satisfaction of the rules and conditions constitutive of chess. When they do take such responsibility, and when the world cooperates well enough to make satisfaction possible — e.g., by not having the pieces multiply and spread to adjoining squares––chess, and all that goes with it, is constituted. Haugeland thinks nothing less will do for objectivity generally, and consequently that mundane and scientific objects are on a par with chess positions when it comes to the *grounds* of objectivity. Hence genuine intentionality is possible only where there is not only the possibility, but the reality, of this sort of commitment. Since something comparable to human intelligence seems to be required for this sort of responsibility taking, it seems to follow that animal and machine intentionality is ersatz.

> Once ersatz intentionality is distinguished and characterized, however, it becomes apparent that there are candidates for it besides GOFAI robots. In particular, I want to suggest that (as far as we know) the intentionality of animals is entirely ersatz (except for purely tropistic creatures, whose intentionality is at best "as-if"). That is, we can understand animals as having intentional states, but only relative to standards that <u>we</u> establish for them. (Haugeland, 1992, p. 303).[7]

It is tempting to agree with Haugeland that nothing less than this will do for the objectivity of knight-forks. Knight-forks, after all, are, well, pretty conventional objects. So it is not too surprising that they can exist only on the supposition that the conventions are somehow kept in force. And what could keep them in force other than the commitments of the players of the game? But surely you have to have illegal concentrations of transcendental philosophy in your veins to think nothing less will do for sticks and stones, conspecifics,

---

> change when, after a few practice shots, I call to my opponent, "All right, this is it," and then proceed to serve? The new element in the situation, I suggest, is my readiness to countenance certain sorts of complaints, for example, that I stepped on the baseline, hit the ball when my opponent was not ready, or was standing on the wrong side of the court. I take responsibility for the holding of certain conditions, for example, that neither foot touches the ground in front of the baseline before the racket touches the ball." (Alston, 1964, p. 41-2)

[7]I think Haugeland underestimates the capacities of social animals, especially those living in a dominance hierarchy, to take responsibility for the satisfaction of social norms. See D. D. Cummins, 1999, and the papers referenced there.

smiles, and threats.

But this is unfair to Haugeland's position. Chess has to be playable, and this requires, to repeat, that the pieces do not multiply, that they do not move (in the chess sense) unless moved, etc., and these are not matters of convention. To give a very crude summary of a subtle and nuanced position, physics and chess differ in the extent to which the constitutive standards are matters of convention, but not at all in the grounds of objectivity, which have to do with a commitment to seeing that the standards are satisfied, not with their conventionality or lack of it.

And yet, surely this position vastly over-intellectualizes what it takes to have targets. And targets are all that objectivity requires, hence all that a theory of intentionality needs to accommodate. Targets are grounded in representational functions, and representational functions do not in general require either conceptual machinery or standard-grounding commitments. From this point of view, it seems it ought to be possible to see knight forks as such without having a clue about chess, and hence without being in a position to take responsibility for the satisfaction of its constitutive rules and conditions. It is tempting to give Haugeland knight-forks, and hold out for stones, but I don't think this will do. If target fixation is what I think it is, then, like Haugeland, I am committed to knight-forks and stones being in the same boat. It is just that my boat is little more than a fleet of floating logs, while his is, well, the Titanic.

Functions Again

To make this contention stick, I need to show how an organism could have knight-forks as targets on the cheap, as it were. To do that, I need to explain how it is possible for a relatively cheap intender to have as one of its functions the representation of knight-forks.

Haugeland points out that perceiving knight forks is a pretty abstract capacity.

> These complementary considerations, that chess can be played in widely different media, and that widely different games can be played in the same media, together with the fact that knight forks can occur (and be perceived) in all the former but none of the latter, show that the ability to perceive knight forks presupposes some grasp or understanding of the game of chess — at least enough to tell when it's being played, regardless of medium. (Haugeland, 1996, p. 248)

What I find interesting about this passage, and the argumentation that turns on it, is the idea that what you do doesn't count as perceiving knight forks unless you can do it all, i.e., recognize them in various media, and distinguish them

from schnight forks (configurations that look like knight forks but are part of schness, not chess) in the usual media. But this conflates two very distinct issues:

(a) how good is the system at recognizing knight forks?

(b) are knight forks among its targets, i.e., are some of its representational efforts to be assessed for accuracy against knight-forks?

I can, for example, be arbitrarily bad at recognizing colors (I am color blind), and yet have the color of the vegetable in my hand as a representational target. I won't hit it reliably, but it does not follow from this that it isn't my target. I can have the shape of the currently foveated thing as a perceptual target even though I always get it wrong to some degree (astigmatism). The mere fact that a dog, or a pigeon, or a novice human, will be reliably fooled by schnight forks, and fail to see knight forks in esoteric media, is beside the point. You cannot, in general, argue that having T as a perceptual target is expensive because accurate perception of Ts is expensive. Predator recognition systems tolerate a lot of false positives in the interest of speed. But when a prairie dog dives into its hole in response to a shadow made by a child's kite, this is a <u>false</u> positive, a mistaken predator identification, not an accruate shadow detection. This is not a point about the content of the animal's representations, which may not distinguish between shadows and predators, but a point about their targets.To understand what is going on, it is crucial to grasp that the targets are predators. Without this, <u>you cannot even articulate</u> the point about trading accuracy for speed.

In general, you are going to vastly under-estimate the representational (and other) functions a system can have if you do not maintain a principled distinction between representational accuracy on the one hand, and perceptual, cognitive or behavioral success on the other. This distinction is precisely what grounds explanations in terms of trade-offs between accuracy on the one hand, and speed and tractability on the other. I think Haugeland misses this point when he writes:

> But there is another important distinction that biological norms do not enable. That is the distinction between functioning properly (under the proper conditions) as an information carrier and getting things right (objective correctness or truth), or, equivalently, between malfunctioning and getting things wrong (mistaking them). Since there is no <u>other</u> determinant or constraint on the information carried than whatever properly functioning carriers carry, when there is no malfunction, it's as "right" as it can be. In other words, there can be no biological basis for understanding a system ad functioning properly, but nevertheless

> misinforming — functionally right but factually wrong, so to speak. (Haugeland, 1998b, 309-10)

But biological norms do enable this distinction, because it can be crucial to understanding the biological role of an intender to see that it is accuracy with respect to Ts that must be measured to judge its adaptiveness, even though it seldom or never gets Ts completely right.

If you miss this point, or just don't agree with it, then, like Haugeland, you will be forced to look elsewhere to find the resources to ground objectivity. You will be tempted to deal with my astigmatism and color blindness, either by claiming that I don't see—have as visual targets—colors or shapes (implausible), or by pointing out that I will take responsibility for corrections aimed at bringing my percepts (or, if that is not possible, my beliefs) into line with independently accessed color and shape information. Since prairie dogs don't take responsibility for their false predator recognitions, you will conclude that they don't see predators.

On the other hand, if you are on a budget and in market for esoteric perceptual targets, here is how to get weak king-side defenses cheap. Imagine someone who is trying to learn to recognize them. They start with some crude criteria that don't work very well. Still, the function of those criteria is to enable recognition of a weak king-side defense. This is their function because the trajectory of the learning is tracked by the difference between the student's guesses and the actual weakness of the king-side defenses in the training set. The student needn't even be conscious of this. Indeed, it may well be that teachers who treat this as like a chicken-sexer case will have better success than their more cognitively explicit competitors. Whether or not such learning could be successful, leading to reasonably reliable discrimination from foils in a variety of media is beside the point. What matters is that the relevant error is the gap between student guesses and weak king-side defenses. Nor need the students be committed to keeping the rules in force. For one thing, they may be spectators, not participants. More importantly, however, it isn't their commitments that count. They may be committed to anything you like, or nothing at all, provided only that it is differences between their guesses and the weakness of king-side defenses that must be compared to track the learning.[8]

---

[8]I don't mean to suggest that tracking learning trajectories is always the issue. I doubt that prairie dog predator recognition is learned at all. But thinking about learning highlights the difficulty because it is in the nature of the case that the target one is trying to learn to see (recognize, identify) is, in the early stages, beyond one's abilities. The novice learning about knight forks, and even some

If we are not careful, this talk of what we must do to track learning will make it look as if I have tried to pass off some ersatz intentionality for the real thing. Not so. It is the gap between weak king-side defenses and the student's guesses that drives the learning, whether or not anyone realizes this. It isn't that we make the learning intelligible by establishing this standard. It is rather that the learning becomes intelligible when we realize what the operative standard has been all along. The only reason to resist this point is the mistake about functions scouted above, the mistake rooted in the failure to distinguish representational accuracy from perceptual (cognitive, behavioral, reproductive) success.

Perhaps Haugeland thinks that that distinction, if viable at all, is a matter of us establishing standards for them, so that the functions in question are themselves ersatz functions. But I don't think it is. There is a subtle relationship between the accuracy of a representation and the success of its consumers. Once we see that accuracy can be sacrificed for speed or tractability in order to improve consumer performance, we will no longer be tempted to say, with Haugeland, that biological norms cannot enforce a distinction between success and accuracy.

Representational Content

I have been arguing for cheaper targets, and hence objects, than Haugeland is prepared to allow. But there is no free lunch. The distinction I have been urging between representational accuracy and consumer success will collapse if you have a use-theory of representational content. I said above that use theories of target fixation are not only acceptable, but inevitable. I also said that, to the extent that representational content is not distinguished from intentional content, use theories of representational content will seem inevitable. Use theories of representational content seem inevitable to almost everyone anyway. What could possibly endow a representation with a content that is completely independent of how that representation is used? And why should we care about such contents? These are large questions, and I will not

---

experts some of the time, will be stymied by non-standard media and fooled by schnight forks, as will a dog or a pigeon. All three may do quite well with the standard cases. And they may never get any better. The training may be discontinued, or the student inapt. But the fact that acquiring the "full ability" may be beyond certain non-human animals, and beyond certain humans, for that matter, is beside the point. Having a target doesn't require being able to hit it.

address them here. But I do want to explain why use theories of representational content threaten to collapse the distinction between representational accuracy and consumer success. If I am right about this, then everyone who assumes a use-theory of representational content should follow Haugeland in thinking that genuine intentional content is rare and expensive, rather than following me in thinking it is as common as dirt.

> Use theories can be understood as applications of the following strategy:
> 1. Begin by noting that the content of r is whatever it is applied to in cases of correct application.
> 2. Provide a naturalistic — i.e., non-question begging — specification N of a class of applications that can be presumed to be accurate.
> 3. Identify the content of R with whatever it is applied to (its target) when N is satisfied.

The trick is in supplying the non-question begging specification of a class of applications of r that can be presumed to be accurate. Use theories require that there be a sort of non-semantic natural kind of cases in which r is bound to be accurately applied to t. After the fact — with the definition in place — these will be cases in which application of r to t is accurate by definition.[9]

But: there <u>are</u> no non-question begging conditions that guarantee accuracy. The best one can do is to require success or effectiveness — normative notions, but not semantic ones. So, for example, you get the suggestion that the content-fixing cases are those in which the producer of the representation and its consumers behave in the ways that account, historically, for their being selected (Millikan, 1984). This fails because, to get selection, what you need is more effective representation than your competition, and this might require less accuracy, not more. Turning this point around, we can see that the account will work only if it is sometimes ok to collapse accuracy and effectiveness. So that is how use theories undermine the accuracy/effectiveness distinction.

Do all use theories do this? I think so, But I haven't the space to argue it here (see Cummins, 1996). What about Haugeland's theory of representational content? Does it do this?

Haugeland doesn't explicitly endorse any theory of representational content. His concern is with intentionality and objectivity. He certainly seems to endorse holism about representational content, but this is difficult to pin down

---

[9]This is why all use theories give you some version of the analytic-synthetic distinction.

because the distinction I make between representational and intentional content is not in the forefront of his writing. Still, I will risk being irrelevant and ask whether holism about representational content is likely to grease (or clear) the path to Haugeland's position on intentionality.

Holism about representational content is essentially the idea that a representation gets its content from its role in a theory-like cognitive structure. It is, in short, conceptual role semantics, the idea that a representation's content is determined by the set of epistemic liaisons it enables.[10] The theory in which a representation r figures determines what r will be applied to. False theories, of course, will not determine accurate applications of r. But we cannot, without circularity, require that r applies to — is true of — those things a true host theory determines it will be applied to. So, what conditions should we place on T?

Notice that this is just the crux of every use-theory, namely, specifying non-question begging conditions under which content-fixing applications can be assumed accurate. What are the options here? If we say that every application licensed by the theory is content-fixing, we leave no room for error at all. (Perlman, 2000) Or rather, since the theory may license application of r to t yet be mismanaged or improperly deployed in some way, the only error allowed is whatever lives in the gap between well-managed and ill-managed theory.

> But misrepresentation should not be confused with improper deployment on the part of the using system, nor bad luck in the results. These can diverge in virtue of the fundamental holism underlying what can count as a representation at all: the scheme must be such that, properly produced and used, its representations will, under normal conditions, guide the system successfully, on the whole. In case conditions are, in one way or another, not normal, however, then a representing system can misrepresent without in any way malfunctioning. (Haugeland, 1991, p. 173)

---

[10]The term 'epistemic liaisons' is Fodor's (1990). It is important to realize that representations do not, by themselves, have epistemic liaisons. An attitude — an application of a representation in some particular cognitive role such as belief, desire or intention — can be justified or rational in the light of others, but a representation cannot. So representations link to epistemic liaisons only via the attitudes they make possible. (See RTA, pp. 29ff. for a full discussion.) I am going to slur over this in what follows because I can't see that it makes any difference.

Well, what is proper deployment? Not deployment that licenses accurate applications, for this returns us in a circle again. It is rather something like this: all the inferences and calculations are done right. 'Done right' threatens circularity again. The only way out is to say something like this: done right means (a) done as they were designed to be done, which leads back to the teleological theory or Descartes' God, both of which lead to (b) done effectively, i.e., done in a way that leads to cognitive, behavioral, or reproductive success, as argued above.

I think it is safe to conclude that there is no way to carry out step two of the use theory schema above without assimilating accuracy to effectiveness. I don't claim that Haugeland has done this; only that whatever pressures there are for use theories generally, and for holistic conceptual role theories in particular, are pressures to assimilate accuracy to effectiveness, hence to under-estimate representational functions and over-price target fixation and intentional content. I don't think there is much pressure for holism about representational content once it is distinguished from intentional content. But there are pressures for use theories of content if you think that representations are structurally arbitrary in the way that words are, i.e., that any structure of formatives could, in a suitable context, represent just about anything.[11] If representations are like words in this respect, what could fix their content other than use? My own view is that this is a good reason to reject the idea that representations are structurally arbitrary, but that is another story.

REFERENCES

Alston, W. (1964) Philosophy of Language. Englewood Cliffs, NJ: Prentice-Hall.

Cummins, D.D. (1999) "The impact of the social environment on the evolution of mind." Synthese.

Cummins, R. (in press) "Reply to Millikan." Philosophy and Phenomenological Research.

[11]There are some hints of this in Haugland's Representational Genera. He clearly thinks that representational content is scheme-relative, and this suggests that a token with a given set of intrinsic structural properties could have arbitrarily different contents in different schemes.

Cummins, R. (1996) Representations, Targets and Attitudes. Cambridge, MA: MIT Press. A Bradford Book.

Davidson, D. (1967) "Truth and Meaning." Synthese 17:304-323.

Fodor, J. (1990) A Theory of Content and Other Essays. Cambridge, Mass.: MIT Press. A Bradford Book.

Haugeland, J. (1991) "Representational Genera." Reprinted in Haugeland, 1998a, 171-206. First published in W. Ramsey, S. Stich and D. Rumelhard, eds., Philosophy and Connectionist Theory. Hillsdale: Laurance Erlbaum, 61-89.

Haugeland, J. (1992) "Understanding Dennett and Searle." Reprinted in Haugeland, 1998a, pp. 291-304. First Published in A. Revonsuo and M. Kamppinen, eds., Consciousness in Philosophy and Cognitive Neuroscience. Hillsdale, NJ: Lawrence Erlbaum, 115-128.

Haugeland, J. (1996) "Objective Perception." Reprinted in Haugeland, 1998a, 224-265. First published in K. Akins, ed., Perception: Vancouver Studies in Cognitive Science, Vol. V. New York: Oxford University Press, 268-289.

Haugeland, J. (1998a) Having Thought. Cambridge, MA: Harvard University Press.

Haugeland, J. (1998b) "Truth and Rule Following." Reprinted in Haugeland, 1998a, 305-361.

Millikan, R. (1984) Language, Thought, and Other Biological Categories. Cambridge, Mass.: MIT Press. A Bradford Book.

Perlman, M. (2000) "Conceptual Flux: Mental Representation, Misrepresentation, and Concept Change. Springer, New York.

Sejnowski, T.J., & Rosenberg, C.R. 1987. "Parallel Networks that Learn to Pronounce English Text." Complex Systems, 1, 145-168.