Philosophy and Phenomenological Research
Vol. LX, No. 1, January 2000

Reply to Millikan

ROBERT CUMMINS
University of California, Davis

Introduction

Here's the problem: On the view I defend in RTA (Cummins 1996), representation is isomorphism. Since isomorphism is a relation between *structures,* it follows that only structures can represent or be represented. But, of course, structures are *abstracta*, whereas it seems that mental representations are states or processes in the brain, and the things they represent are, at least sometimes, objects.[1] Since one and the same brain state/process has lots of different structures, and one and the same object has lots of different structures, it looks as if my theory says next to nothing about whether or how a given brain state/process represents a given object. Let's see.

Begin with structures: A structure is a pair of sets, <R,D>, where R is a set of relations, and D (the domain) is a set of "objects", i.e., of arguments to the elements of R. <R,D> is isomorphic to <R*,D*> just in case there exists a one to one function mapping elements of R onto elements of R* and elements of D onto elements of D* such that an element of R holds of a sequence s of elements of D iff the image of R in R* holds of the images of s in D*. A function like this is an isomorphism. Generally there are many isomorphisms between two structures if there are any.[2]

Now let's ask what the relation is between an object (or state or process) and a structure. I am a realist about structures, meaning I don't think they are in the eye of the beholder. Space-time, for example, has an affine structure, and its having that structure is an objective fact about space-time. Similarly, an utterance token in English has a phrase structure, and its having that phrase-structure is an objective fact about that utterance. I don't propose to argue for this sort of realism beyond noting that, if you are not a realist about structure in this sense, then you ought to think most science worthy of the name is trivial, a major strategy since Galileo having been to get a grip on

1 It has become commonplace among causal theorists (e.g., Fodor,) to assume that it is *properties* that are represented, or perhaps *that property p is instantiated there*. What makes properties attractive to causal theorists, of course, is that they figure in causal relations: x's coming to have F causes y's coming to have G.

2 Isomorphisms are not what Millikan calls mapping or projection rules. See below.

the universe by finding mathematical structures that are isomorphic to the structures of nature.

Millikan, of course, has no quarrel with realism about structures. Although she sometimes writes as if she thinks anything goes, her worry is not that you can find any old structure you want anywhere you want to find it, but that there are lots of different structures in the same thing.[3] And I agree that there *are*; it just doesn't worry me, for reasons that will emerge as we go along.

It is one thing to be a realist about structures, and another to say what it is for an object or process or whatever to *have* one. But I don't think we need linger over this. For one thing, as just mentioned, it is standard scientific practice to mathematically specify structure in nature. Phrase structures, for example, are structures in good standing (a phrase structure is a set of constituents and a set of dominance relations), and no one worries overly much about what it is for a sentence token to have one. And the same sentence token also has other structures as well, e.g., a phonological structure and a logical form. Road maps specify the highway and intersection structure of a city, while a topological map of the same city specifies its elevation gradients. A phrase structure is not only specified but displayed by a phrase structure tree, a representation that itself actually *has* the structure in question "explicitly". Similarly, road maps not only specify highway and intersection structure, they *have* an isomorphic structure.

A (token) road map, of course, is a physical object--paper and ink. So it has lots of structures besides the one isomorphic to the highway and intersection structure of interest. Ditto for the city. So: what determines whether the map represents the city? It looks like there are only two choices:

> *Choice One*: Promiscuity: M represents C if there is some structure M has that is isomorphic to some structure that C has. The map gets credit, as it were, for every structure it has that is isomorphic to some structure the city has. Moreover, it gets credit for every isomorphism: if there is more than one way to map the graphical structure of the map onto the street and intersection structure of the city, one of which takes a certain line onto Pine Street, and another of which takes that same line onto Elm Street, then that line represents both Streets.
>
> *Choice Two*: Relativization to an isomorphism or "mapping rule" f: on this option, it is only that structure of M that is the domain of f that is significant, and only that structure that is the range of f that is represented. Rather than M representing C, we get something like

[3] Putnam (1988, Appendix) seems to be an anti-realist about structure, as does Searle (1992).

this: M's having the f-domain structure represents C's having the f-range structure. We then suppose that fixing the cognitive system somehow fixes f, with the consequence that only one structure of the representation matters semantically, and its target is just the range of f.

Relativization to an isomorphism is beguiling because it seems to solve all the problems at one stroke: the domain of the isomorphism fixes the constituents of the map that are to count, and the relations on them that are to count, and the range of the isomorphism does the same for the city. So it is fixed which structure of the map is to count, and which structure of the city is to count. And, finally, there is only one mapping between these structures that counts. Everything is neat and clean. When a representation R shows up in a system that uses the f-mapping, we know which properties of R are doing the representing work, and we know which properties of the target are represented.

Still, I rejected relativization in RTA. I rejected it because I think the only way to make sense of the idea that there is a "correct mapping rule" is the way Millikan makes sense of it: The map producer[4] is functioning Properly when it constructs maps that are isomorphic to its target structures according to that rule, and/or the map interpretation mechanisms of the map consumers are functioning Properly when they interpret the map according to that rule. This, I claim, makes representational content a function of use, and that in turn compromises the notion of representational error, and hence the normativity of representational accuracy. This story is part of the critical content of RTA that Millikan doesn't discuss. But it is essential to the point of dispute between us, because it is what is behind my rejection of relativization. The central claim of RTA is that you ultimately cannot make sense of representational error unless you have a conception of representational content that makes the content of a representation completely independent of how that representation functions in the systems that use it. The idea that the content of a mental representation could be completely independent of the way it functions in the mind strikes nearly everyone as mad. Mark Perlman (forthcoming) concludes, via the obvious modus tolens, that there cannot be any representational error. But most people simply deny the conditional: there just *must* be a way of making representational error compatible with a story about representation that makes the content of a representation in Σ depend on what Σ does with that representation.

---

4 The producer doesn't have to be whoever or whatever makes the map. If I ask you for a road map of Chicago, and you hand me a map, you are the "producer" in this transaction, and I am the consumer. The target is the street and intersection structure of Chicago. You get the representational task right if you give me something that is a road map of Chicago.

Before making my case that realtivizing representation to a particular isomorphism or mapping rule is bound to compromise the notion of representational error, I want to rehearse briefly why I think relativization is unnecessary in the first place. Perhaps Choice One isn't as bad as it seems.

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.

The first thing to notice about this story is that, since targets are phrase structures of current linguistic inputs, not just "things," there is no problem about which structure of the "thing" is the target. The structure the phrase-structure intender must represent to do its job is the phrase structure of the current linguistic input. Assuming that the current linguistic input actually has some definite phrase-structure, that is the structure against which R must be matched to judge its accuracy. The fact that R may match, or fail to match, a variety of other structures of the current linguistic input is simply irrelevant.

Meanwhile, R is, of course, itself some physical object or process or whatever, and, as such, will have a variety of structures. Which of these is to count as the *relevant* structure, the structure that does the representing? According to the doctrine of RTA, *all of them count*, as far as representation goes. If R has a structure isomorphic to the phrase-structure of the current input, then R accurately represents the current phrase structure. The information is there, and it is there in the way peculiar to what I call representation: it isn't, for example, arbitrarily encoded in a way that requires what amounts to a translation manual to access. But the consumers of R may not be able to exploit that information if it isn't in a structure of R that those consumers happen to be sensitive to. Thus, there is more to successful cognition than successful representation. The system must also be able to exploit whatever successful representations are produced. If an intender feeds a punch card that has a beautiful color picture of the target printed on it to a punch card reader, its efforts will be wasted because the reader is insensitive to everything but punch pattern.[5]

Why insist on the idea that structures of R to which consumers are insensitive nevertheless represent? Because we need to be able to account for the fact that a consumer may acquire (perhaps by learning) the capacity to exploit

[5] Even if R accurately represents its target in a way consumers can exploit, cognition may still be unsuccessful. Accuracy always has a price. Great accuracy can be too expensive. Sometimes better cognition results with less accurate but cheaper representation.

structures of R that it previously could not exploit. You cannot learn to exploit information that isn't there, so you had better not have a story about content that recognizes only content bearing structure to which consumers are already sensitive. You cannot have a representation driven story about learning if your story about representational content implies that the content that is supposed to do the driving is only there when the learning is complete. I'll return to this theme below in my discussion of attempts to relativize the representation relation to a particular mapping rule.

The Crude Use Theory (CUT)

I propose to start by making my case against a very crude use theory. We'll see later whether Millikan's sophisticated theory avoids the problems I raise for its crude relative.

So: the Crude Use Theory (CUT) says (i) that the content of a representation is a function of (among other things) a "mapping rule" that specifies how a representation is supposed to be mapped onto its target, and (ii) that the mapping rule that determines the content of r is the one consumers of r actually use when they put r to work.[6] If, for example, the representation is a punch card, and the consumer is sensitive only to the relative locations of the holes in the card, then it is the holes and their locations that count, and nothing else. The fact (assuming it is one) that the card has a map of Chicago printed on it is irrelevant to assessing its representational content. What could possibly be wrong with *that*?

There are two things wrong. The first is that CUT fails to make conceptual room for the possibility that a consumer might learn to exploit a given representation type in a novel way. To get a sense for this possibility, let's begin with an analogy. Suppose you learn to use road maps simply to generate directions. By looking at the map, you can see that if you make a left at the third intersection, than immediately make a right, you will be on the street you want. It never occurs to you to extract distance information until, one day, someone suggests that the map shows a shorter route than the one you generated. Prior to this insight, you exploit the map according to a mapping rule that would be insensitive to various distortions of the map, such as shrinking the vertical axis relative to the horizontal one. If we assess representational content through the lens of the mapping rule its consumer assumes, we will have to say that there is no distance information there to be exploited. And this will make it impossible to give a sensible explanation of how, as the result of learning, you come to use tokens of the same formal

[6] Notice that what Millikan calls a mapping rule isn't an isomorphism, but a kind of wannabe isomorphism: it is a recipe for attempting to map a representation onto a target. "Interpretation rule" might be a better name for the item in question. An example of a rule of the kind in question is: intersecting lines on the right of the map are intersections to the east of those represented by intersecting lines in the middle.

type to effectively compare routes for length. To deal effectively with a case like this, we need to assume that, prior to learning, relative distances *are* represented, even though you, the consumer, never exploit them. We will have to allow for this even if you (and your ancestors) never use these maps to compare distances, on pain of losing track of the truth that, but for various contingencies, you could have learned to do it. Your capacity to learn this is as much a fact about the map as about you, and the relevant fact about the map is that it represents relative distance whether or not this fact is ever exploited.

The case just rehearsed is just an analogy. Is there any reason to think that natural cognitive systems learn to exploit previously unexploited structural features of their representations? I think there are.

The idea that a brain can learn to exploit previously unexploited structure in its representations is absolutely fundamental to connectionist conceptions of brain functioning. Connectionist learning consists in adjusting connection weights so as to respond properly to input activation patterns. Those input patterns are assumed to represent stimuli prior to learning. To use a famous example, a NetTalk (Sejnowski and Rosenberg, 1987) input consists of a pattern of activation over 189 input units as indicated in figure one. Each pattern represents a "window" of seven letters, NetTalk's task being to generate a phonetic representation of the middle letter.

Initially, weights are set randomly, and the network is trained up using the back propagation algorithm and a training set consisting of a representative text for which the phonetic values have been computed for each letter. Output activation patterns are compared with correct patterns (i.e., patterns encoding correct phonetic values), and weights are adjusted according to the generalized delta rule until accurate pronunciation of the text is achieved.

This whole process only makes sense on the assumption that input patterns represent letter strings prior to learning. Indeed, the representational content of input patterns remains the same throughout learning. But, prior to learning, the network cannot, of course, exploit those representations properly. That is precisely what the process of weight adjustment is designed to achieve. Since, however, the way the network "uses" input patterns varies continuously throughout learning, CUT will force us to say that the contents of the input patterns changes continuously as well. We shall have to say (as Andy Clark said to me in correspondence) that the network creates the content of its input patterns as learning progresses. But if we say this, we have no reason to say that early responses are errors. And if early responses are not errors, why change the weights in any particular direction? Indeed, why change them at all? Evidently, CUT makes it impossible to even articulate a

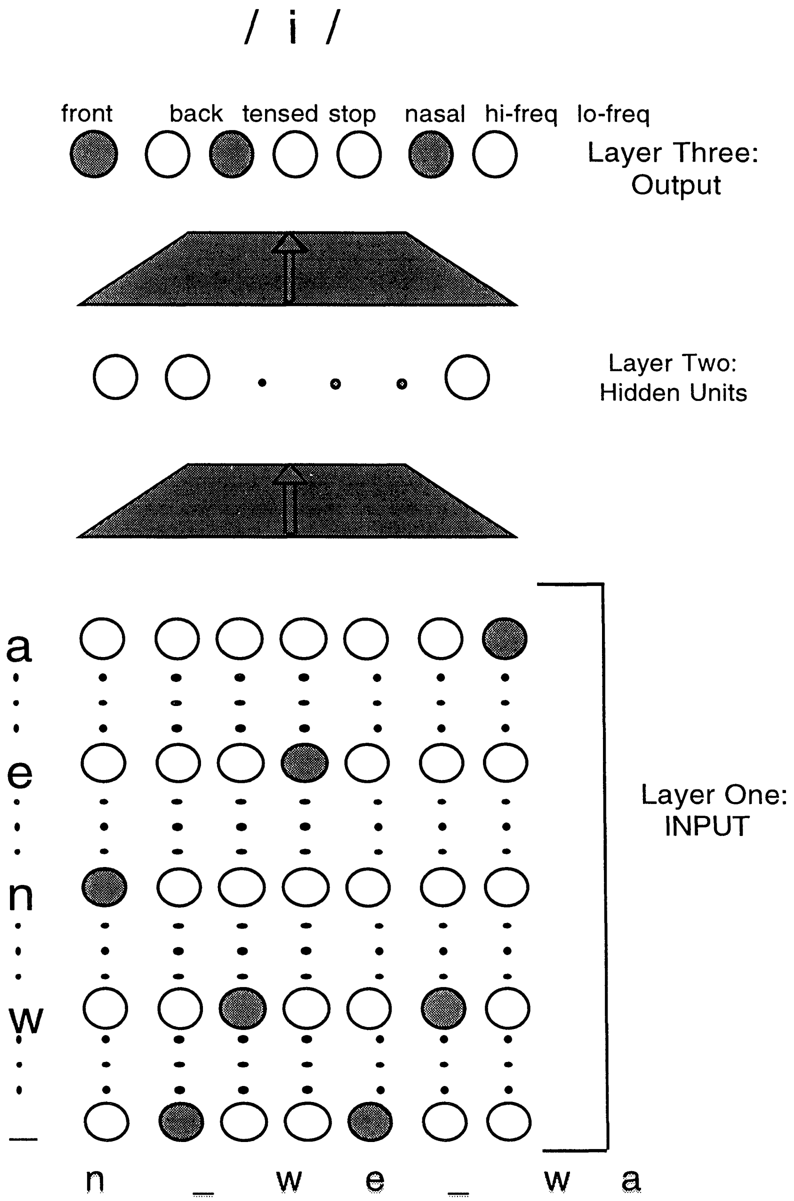

Figure One: Schematic Diagram of NetTalk:

connectionist learning problem.[7] The whole point of a representationalist framework is to be able to articulate psychological processes in a way that allows us to track the operation of semantic constraints.

Having come this far, we can see, even if we are not devout connectionists, that the phenomenon in question must be ubiquitous in the brain. Proximal perceptual stimuli must be well-represented before the appropriate cortical structures learn to exploit those representations in target location and recognition. For example, the retinotopic maps created in the visual cortex of an infant surely contain much information that the developing brain is able to extract only as the result of considerable visual learning (Spelke, 1990, §4.3). The visual system needs to learn to exploit the geometrical information that is there in the cortical projections, and this, surely, is largely a matter of *acquiring* a "mapping rule" that will do the job.

CUT, it appears, has things exactly backwards. It isn't the mapping rule the system uses that determines the contents of its representations; rather, systems revise their mapping rules in order to better exploit the contents of their representations. The game, in a nutshell, is to find mapping rules that allow one to exploit whatever isomorphisms exist between one's representations and the environment. It is obvious that CUT undermines all of this. CUT assesses representational content through the filter of the mapping rule that representation consumers actually exploit, and hence must treat what are evidently cases of learning to exploit previously unexploited structural features of their representations as cases of a change in content. The result is incoherent: you cannot learn to exploit information that isn't there to begin with. Millikan's theory is not CUT, of course. But I'll argue shortly that it fares no better with the sort of learning just rehearsed.[8]

The second problem with CUT is that it fails to distinguish improperly exploiting an accurate representation and properly exploiting an inaccurate one. Once again, it is useful to begin with the analogy discussed above. Suppose, as before, that you use a road map of your city to find your way around, but never make distance judgments based on the map. You use a mapping rule that is insensitive to distances, as revealed by the fact that various distortions of the map make no difference to you. CUT implies, as we have just seen, that the map doesn't represent relative distances, and hence implies as well that these not are misrepresented when the map is distorted. What *does* count as misrepresentation according to CUT? A simple case would be one in which the map leaves out a street. As a result, you will do such things as take the third left when you should have taken the fourth. The

7 This point generalizes to any sort of connectionist learning. In all such cases, we must assume that the contents of input patterns remain unchanged as the weights are adjusted on pain of having no grounds for saying that early responses are errors.

8 The forgoing argument doesn't appear in RTA, but it was conveyed to Millikan in the correspondence we had when she was writing her review.

number of intersections between you and the street you want is misrepresented, and you will be misled because the mapping rule you use involves counting intersections. So far, so good.

But now consider a case in which you use the map upside down because you are confused about which way is north. (Everyone has done this: you turn the map upside down because you think you are headed south.) In this case, the mapping rule you use will lead you astray. But the mistake is yours, not the map's. If we say that the content of the map is determined by the mapping rule that consumers use, we shall have to say that the map misrepresents the city when you are mistaken about which way you are headed. But if we say this, how shall we assimilate the following fact: *Had you rotated the map 180°, you would not have gotten lost*? CUT has the consequence that the map mysteriously changes content (continuously?) when it is rotated. But, surely, the counterfactual above is true partly because the map is accurate (no missing streets, etc.). Because it is accurate, it is is capable of conveying the correct information if it is used correctly. CUT makes no room for the obvious but crucial fact that there is a mapping rule that will allow you to consume the accurate information in the map. CUT makes no room for this fact because its central feature is that *the map has no determinate content independently of the use consumers make of it*. Yet there is surely a structural property the map has, independent of use, that accounts for the fact that it can be used to accurately plot routes and relative distances, and that property, according to the theory of RTA, is (part of) its representational content.

The rotated map is an illustration of a more general phenomenon. Cognitive processing is often facilitated by the occurrence of content preserving transformations on representations, transformations that, while preserving content, alter form in a way that makes that content accessible to consuming sub-systems. Logic and mathematics are full of such transformations, and it is difficult to think of any theory of cognitive processing that doesn't capitalize on this idea in one way or another. But the idea makes sense only if we are able to distinguish between the case of an accurate representation that is improperly used and an inaccurate one that is properly used. By undermining that distinction, CUT undermines one of the most fundamental explanatory tactics that make up the representationalist strategy.

By keeping a representation's content distinct from facts about its use, including which mapping rule describes its use, the theory of RTA allows for the psychological possibility that cognitive systems can learn to exploit representations in new ways, and for the closely related logical point that you cannot appeal to content to explain why a representation can be used successfully if you explain content in terms of successful use. We have just seen that CUT, the "crude use theory", gets it wrong on both counts. It remains to be seen whether the more sophisticated use theory advocated by Millikan fares any better.

Millikan's Account

The Crude Use Theory (CUT) says (i) that the content of a representation is a function of a "mapping rule" that specifies how a representation is supposed to be mapped onto its target, and (ii) that the mapping rule that determines the content of r is the one consumers of r actually use when they put r to work. Millikan's account agrees about (i), but holds that it isn't the mapping rule that consumers of r *actually* use currently that determines content, but rather the mapping rule they use when functioning Properly. Moreover, her theory requires that the producers were selected for producing r in accordance with the same rule. Given Millikan's account of Proper Functions, this amounts to the claim that the mapping rule that fixes content is the one whose past use conferred selective advantage on the ancestors of current consumers. So, rather than (ii), we get (ii*): mapping rule M determines the content of r iff the mechanism(s) in producers/consumers responsible for production/exploitation of r were selected for (replicated) because they used M. (i) and (ii) together give us what I will call Millikan's use theory, or MUT.[9]

MUT *is* a use theory. The mapping rule that producers/consumers use when functioning Properly in their production/exploitation of r that determines the content of r. (The requirement that r be produced in accordance with a specified rule is as much a requirement on use as is the requirement that consumers interpret in accordance with that rule.) It is worth pausing a moment to see why this is an improvement over CUT. The idea is that the content of a representation is just its target on occasions of correct use. If every occasion of applying r to a dog is a correct application of r, and every occasion of applying r to a non-dog is an incorrect application of r, then the content of r is *dog*. As it stands, this formula is unhelpful because it appeals to the semantic notions of correct and incorrect application. But suppose you could assume that users of r were using it correctly when functioning Properly. Then you could say: If x applies r to dogs and only dogs when functioning Properly, the the content of r (for x) is *dog*. Then you could give a selectionist (or some other suitably naturalistic) account of Proper Functioning, and you would have an account of content. All use theories work this way, i.e., by trying to find some naturalistically specifiable condition under which the application of r is bound to be correct, and then identifying the content of r with whatever it is applied to under that condition. MUT is better than CUT precisely because it takes a serious shot at specifying a condition under which applications can be presumed to be correct.

---

9 The full account is in Millikan, 1984. Only certain very general features of her rich and nuanced account will relevant to the discussion here.

So how does MUT fare with respect to the two objections just leveled against CUT? Start with the problem about learning to exploit previously unexploited structural features of a representation, and consider again the person who uses a map to determine routes but not to compare distances. I see no reason to think that this is not a case of Proper Functioning, i.e., a case in which the mechanism responsible for map use was replicated because it used street maps in accordance with M, a mapping rule sensitive to line intersection but not relative distance. The responsible mechanism is, we may suppose, a script for map use. That script, we may suppose further, was passed on, hence replicated, because it provided for the successful plotting of routes and in spite of the fact that it makes no provision for comparing distances. Given these assumptions, MUT will rule that the map does not represent relative distances. But this is incompatible with the fact that users could learn to use the map to estimate and compare distances. A defender of MUT will have to say that the map acquires this content during the learning process. But surely the learning process is possible precisely because the content is already there, waiting to be exploited. MUT, it seems, is no better off than CUT when it comes to giving a coherent description of the kind of learning that, if connectionists are even remotely right, must be ubiquitous in the brain.

We are not, of course, entitled to assume that it is a Proper function of the producer to produce maps representing distances. Since, by assumption, there is no history of exploiting such information, we can hardly suppose that producers were replicated because they produced maps that had it. But we are entitled to assume that it is a Proper function of the producer to produce maps representing the structure of the city. If it happens to do this in a way which encodes distance information, then a consumer could learn to exploit that information. MUT will simply rule that there is nothing there to exploit that counts as content about distance unless and until some consumer actually does exploit it. But if a structural property of a map that can be exploited to determine distances isn't a representation of distances, it will do until the real thing comes along.

The basic point to keep in focus here is that natural selection can only select what is already there. A producer mechanism cannot get replicated because it produces representations that encode a certain kind of information unless those representations already encode that information. Once the information is encoded and there to be exploited, the mechanism can be selected because it produces representations encoding that information, but not before. This puts selectionist stories in a bind. They must say that the mechanism is selected for producing representations that do *not* encode the crucial information, since that information is not *there* until selection has occurred.

How does MUT fare with respect to the second objection leveled against CUT, namely that it fails to distinguish improperly exploiting an accurate representation and properly exploiting an inaccurate one.

Unlike CUT, MUT does make room for the possibility that a representation is being applied in accordance with a mapping rule that does not fix the content of that representation, since it allows for the possibility that, in using r in accordance with M, a consumer of r is not functioning Properly. An all too familiar case of this is treating a grid line on the map as a street representation. Behavior in accordance with the rule that treats grid lines as street representations is, presumably, not the behavior that got the consumer mechanism replicated, so MUT will not treat this as a content-fixing case.

Nevertheless, MUT still gets the original upside down map case wrong. A consumer using the standard projection rule with an upside down map will correlate intersecting lines at the left side of the map with intersections in the west rather than with intersections in the east, and will get lost as a consequence. Since, by hypothesis, the projection rule being used is the the "right" one, i.e., the one MUT says is content fixing, it follows that the upside down map misrepresents the intersection structure of the city. And, just as before, it will follow that rotating the map changes its content. But, to repeat what was said above, surely the fact that everything would work fine if the map were rotated 180° is to be accounted for by the fact that the map is accurate as is, by the fact that the information is already there, in a structural property the map has independent of use, ready to be consumed if it is just exploited properly. Like CUT, MUT makes no room for this fact because its central feature is that *the content of the map is not fixed independently of the use consumers make of it.*

It is possible to object to this line of criticism on the grounds that the disoriented map user is not in fact using the "right" projection rule, since the "right" projection rule maps certain specific line intersections onto certain specific street intersections, and that mapping is not in force when the map is mistakenly rotated. But this idea really won't stand scrutiny: The disoriented user gets confused precisely because the projection rule he uses is one that maps line intersections *on his left* to street intersections in the west. If he in fact used the other rule, he wouldn't get lost. Nor is it plausible to insist that this is not the "right" rule according to MUT. There is no reason to suppose that our map user ever used any rule but the one in question. Indeed, we can simply stipulate this without running afoul of MUT, since MUT will accommodate any rule that allows for enough success to ensure replication, and the rule in question surely satisfies that constraint.

It is also possible to object to the idea that a representation has a content independently of any particular projection rule if one thinks that having a content means having one particular interpretation to the exclusion of all others. PTR certainly doesn't provide for contents in this sense. But I cannot see why it should. A representation is valuable on a particular occasion if it has *some* structure similar to that of its target on that occasion, and the subsystems that consume that representation are sensitive to the structure in question. The fact that the representation has other contents as well, in virtue of having other structural properties, has no tendency to undermine a content explanation that appeals to the presence of one of these structural properties. Moreover, we have just seen that there are compelling reasons to treat as yet unexploited structure in a representation as content, assessable for accuracy. We need a conception of content that makes room for an assessment of what's there prior to, and independently of, any attempts to exploit it.

A more serious worry is that the notion of structure itself is not determinate, a worry with which Millikan appears to have some sympathy. I have none. If the notion of structure is not determinate, then a sentence has no determinate phrase structure to represent, and so it will be of no concern that a state or process in the brain has no structure determinate enough to represent it. If, on the other hand, things in the world that need representing have determinate structures, as all natural science assumes, then states and processes in the brain, being comparable things in the world, assuredly have them too. I freely confess that I have no general account of structure in the physical world, but surely a theory of mental representation should not be expected to solve a metaphysical problem that affects all of science equally. Motions, molecules, DNA, space and natural languages all have structures. Why not representations?

MUT doesn't differ from CUT at all in the crucial respect that underlies the problem both face in accommodating the possibility of learning to exploit previously unexploited structure in a representation (the absolutely basic and ubiquitous case if connectionists are right about learning), or the problem both face in explaining how a good projection rule and a good representation can still go wrong. The difficulty with both CUT and MUT is the idea that the content of a representation is fixed in part by the projection rule that describes the way in which consumers (actually or Normally) exploit that representation. That idea makes information unexploited by the actual/Normal projection rule impossible, just as it makes it impossible that an accurate representation could be mis-exploited in accordance with the actual/Normal rule.

Success and Accuracy

Use theories, as pointed out above, 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 use.

2. Provide a naturalistic specification of a class of uses that can be presumed to be correct.

A very tempting way to go about (2) is to look for cases in which the application of a representation is effective, i.e., leads to success. This is what MUT does: it reasons that the projection rule in play in the Normal cases, i.e., the cases that account, via selection, for the replication of the mechanisms that generated those cases, must be a good projection rule, i.e., one that results in the correct application of the representation in question. The Normal applications of a representation, are, by definition, the cases that conferred selective advantage on the appliers, so they must have been correct applications. How could incorrect applications confer selective advantage?

One of the central points of RTA is to distinguish between representational error and cognitive error. Just as cognitive error need not be traceable to representational error, so representational error need not lead to cognitive error. Indeed, it often happens that a less accurate representation is more effective than a more accurate one because it is more tractable. A street map of Chicago that shows every street will generally be too big to unfold in the car. Reduction doesn't solve the problem because it results in lines so close together they cannot be distinguished. One common way around this dilemma is to show only the main streets. The result is a map that is less accurate but more tractable, hence often more effective, than a map showing every street. Natural selection must often prefer tractability to accuracy. When it does so, Millikan's theory of representational content will always get things wrong, because it identifies the content of r with the targets to which it was effectively applied. It will deem the main-routes map an accurate map of Chicago, and the all-streets map an inaccurate one. Once the content of a representation is identified with the targets to which it was effectively applied, we cannot even articulate the point about there being a trade-off between accuracy and tractability, a point that evidently provides Cognitive Science with some powerful explanatory leverage.

Representations and Intentional States (Attitudes)

Having said all this, I should point out that Millikan's theory is intended as a theory of the content of intentional states (intentional icons). I have no quarrel whatever with the idea that the content of a belief or desire or intention is in part determined by how a representation is being used. The doctrine of RTA is that you get the content of an attitude by looking at which target a representation is applied to on a particular occassion. That makes the content of an attitude a function of use (though not exclusively). I argue that theories like Millikan's give us the outlines of a theory of target fixation: they tell us what it is for t to be the target of the tokening of a particular representation. Error, Jerry Fodor once said to me, is when a representation is applied to something it doesn't apply to. My theory is a theory about the second concept of application. Millikan's theory, I suggesed in RTA, is best construed as a theory about the first. So I really have no dispute with her theory construed as a theory whose target is target fixation. My point is simply that the same machinery won't do as a theory whose target is representational content.

References

Cummins, R. 1996. *Representations, Targets and Attitudes.* Cambridge, Massachusetts: MIT Press.

Fodor, J. 1990b. *A Theory of Content and Other Essays.* Cambridge, Massachusetts: MIT Press.

Millikan, R. 1984. *Language, Thought and Other Biological Categories.* Cambridge, Massachusetts: MIT Press.

Perlman, M. (forthcoming) "Conceptual Flux: The Case against Misrepresentation." Klewer.

Putnam, H. 1988. *Representation and Reality.* Cambridge, Massachusetts: MIT Press.

Spelke, E. 1990. "Origins of Visual Knowledge," in D. Osherson, S. Kossly, and J. Hollerbach (eds.) *Visual Cognition and Action.* Cambridge, Massachusetts: MIT Press.

Searle, J. 1992. *The Rediscovery of Mind.* Cambridge, Massachusetts: MIT Press.

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