

Form, Interpretation, and the Uniqueness of Content: Response to Morris

ROBERT CUMMINS

Department of Philosophy, University of Arizona, Tucson, AZ 85721, U.S.A.

Abstract. In response to Michael Morris, I attempt to refute the crucial second premise of the argument, which states that the formality condition cannot be satisfied “non-stipulatively” in computational systems. I defend the view of representation urged in *Meaning and Mental Representation* against the charge that it makes content stipulative and therefore irrelevant to the explanation of cognition. Some other reservations are expressed.

Key words. Mental representation, computation, formal condition, symbols, intentionality, computationalism, cognition.

Introduction

This paper divides into three parts. Part I is concerned with the first stage of Michael Morris’ attack on mental representation [Morris (1991), this issue]. I attempt to refute the crucial second premise of the argument – the premise stating that what Fodor calls the formality condition cannot be satisfied non-stipulatively in computational systems. In Part II, I defend the view of representation urged in *Meaning and Mental Representation* [Cummins (1989)] against the charge that it makes content stipulative and hence irrelevant to the explanation of cognition. Part III is a collection of quibbles about Morris’ discussion of the argument of *Meaning and Mental Representation*.

Part I: The Formality Condition

Morris rightly supposes that CTC (computational theories of cognition), and, probably, all other representation-invoking theories, require that representational types be non-semantically individuated. This constraint is what Jerry Fodor (1980) calls the Formality Condition (FC), and I will follow him in this. What Morris is arguing is that FC cannot be satisfied in a way that doesn’t trivialize explanatory appeals to representation. He concludes from this that there are no mental representations, for he stipulates that nothing is to count as a mental representation that doesn’t satisfy FC in a non-trivial way. Fair enough: nothing that doesn’t satisfy FC could be what is called “representation” in the CTC. If FC cannot be non-trivially satisfied, then the CTC and its associated notion of representation is hopeless.

Morris thinks FC cannot be non-trivially satisfied because he thinks it could only be met by stipulation. We could, he says, *assign* representational content to

things in virtue of their shapes or sizes or whatever. But, he thinks, it would be miraculous if there turned out to be a coincidence between form and non-stipulative content.

Given the general point that there is no such thing as a neutral individuation of objects, it should follow that if the concerns and characteristics of two levels of description are sufficiently different, it will seem, on the face of it, to be a miraculous coincidence if the individuation of the objects described by the two vocabularies should match. The points just made about word-identity look enough to show that the semantic and non-semantic levels of description *are* sufficiently different. It seems, then, on the face of it, that it would be a miraculous coincidence if semantically individuated objects could also be individuated non-semantically. [Morris (1991), pp. 13–14]

Of course, the semantic level of description and the formal level are very different. In general, they don't type things alike. But we don't care what happens *in general*. We care whether the formal and semantic levels type things alike in some specified computational system. It is not true that all and only triangles refer to cows. But all and only triangles that appear in the data structures of some particular system Σ might refer to cows, for that may be Σ 's way – its only way – of representing cows. The idea is that, given Σ 's computational structure, it will be a law – a law *in situ*, as Millikan (1984) has taught us to say – that, in Σ 's data structures, something represents cows just in case it is triangular. Moreover, in such a system, there is a clear sense in which something that represents cows does so because it is triangular, for it is a triangle that must get tokened if Σ is to get into the cattle business.

Why should an *in situ* correspondence such as this be miraculous? We can, Morris concedes, design systems like Σ that satisfy FC. But if *we* can do it, why can't evolution and learning? The idea behind AI, after all, is that we might be able, using computational techniques, to duplicate nature's design.

I think Morris thinks of assigning interpretations as like spreading butter on bread: the interpretation is, as it were, just laid on afterwards, more or less as one pleases (though there are "some constraints"). He imagines being given a computational system, and then imagines simply stipulating a semantics for its data structures taking care that the result satisfies FC. But, of course, any proper interpretation of a computational system is bound to satisfy FC, because, as John Haugeland has emphasized (1985), a computational system just *is* an automatic *formal* system. You've ensured satisfaction of FC when you have written the program. Adding the interpretation doesn't add the representational contents, it just recognizes them. Programmers don't get FC satisfied by cleverly avoiding interpretations of their data-structures that don't satisfy FC. Interpretations that don't satisfy FC will simply be irrelevant to the system's behavior. It follows that, if nature can design any computational system at all, it can design one whose data structures satisfy FC, because every computational system satisfies FC.

Thus, premise (2) of Morris' argument – the claim that FC cannot be non-stipulatively satisfied – seems to be quite obviously false. Its falsity seems so obvious to me that I wonder if I have misunderstood something. With this in mind, let's look at what Morris says about FC and natural language word identity.

Morris thinks that natural language words, construed as semantic types, cannot satisfy FC. His reason for this is puzzling. It is that for any word w there are bound to be, in principle, if not in fact, things formally equivalent to w that are not tokens of the type because they deserve no interpretation at all (a noise made by a coffee machine, or a shape in the sand produced by an ant), or because they deserve different interpretations (tokens in another language). But why should we worry about this? Formal equivalents will, in fact, be recognized by native speakers as tokens of w , regardless of how produced. They will be interpreted accordingly as well: the interpretation process is so automatic and fast, you just can't help it. This seems to show that formal identity is what matters to the recognition and understanding systems. What we don't do, of course, is believe that ants produce scrawls in the sand with the intention of communicating to us in English. (Coffee pots are tougher: cars talk to us now, so I think I would at least hesitate over a coffee pot I didn't know well.) A language formally identical to English but with a different semantics might fool me for quite a while, depending on the topic and circumstances of utterance.

What all this shows is that deviations from FC cause problems for natural language processing. One supposes that one's language of thought will be less subject to these problems than natural language, but it would be silly to suppose it is noise free. The idea behind FC is not that semantic types be formally definable (except as an idealization), but that, in a computational engine, as in natural language, things will go smoothly to the extent that FC is satisfied, for the system is designed to react to symbols in virtue of their form. Noise is a practical problem for the system, not a philosophical problem for the CTC. The CTC's commitment to FC is not a commitment to the idea that semantic types can be formally defined, but to the idea that (other things equal) the system will function properly to the extent that FC is satisfied by the language of thought. And that much is pretty obviously possible because it is pretty obviously true of natural language as well: things work best when FC is satisfied. As users of natural language, we have both simple and sophisticated ways of detecting spurious forms; so, presumably, does our mental program. One wonders how someone so skeptical of FC as Morris is thinks symbol processing is done, whether the symbols be natural language, language of thought, or the data structures of a computer program. Does he think semantic properties are directly detected – transduced, as it were? For if they are inferred, as (nearly) everyone believes, they must be inferred from something else, and the “something else” is just what form is supposed to be. The idea that representations can be non-semantically identified with fair reliability seems the only alternative to the mystery theory of representation processing.

I think the real reason Morris thinks even *in situ* satisfaction of FC would be miraculous goes something like this: *real* semantic content – “original intentionality” – attaches to mental representations in virtue of the way the system harboring them is related to its environment and/or selective history. The idea is that representations have their content because of the way they are “grounded”,

not because of factors internal to the system such as form. Morris is right to think that this “externalist” picture of content won’t mix well with FC, but, as we will see shortly, the CTC has no use for externalist accounts of content in any case. It will take a bit of preparation to sort this out, however.

Part II: Stipulation and Objective Determinate Content

Morris believes in something he calls *objectively determinate content* (ODC hereafter). Content

... is objective, on these views, because it is not a matter of stipulation or decision what content a representation has. And it is determinate in the sense that there is always just one correct answer to the question, ‘What content does this representation have?’ These theories deny that there can be incompatible but fully adequate interpretation schemes for the representations in a system. [Morris (1991), p. 6, this issue]

Morris concedes that it is an objective matter whether or not a given computational system admits of a specified interpretation. His objection is that such interpretations are not unique, and hence that the sort of interpretational semantics that I argue is adequate to the explanatory needs of the CTC (Cummins, 1989) does not provide ODC.

Of course proper interpretations are not unique, but neither are the representational contents required by the CTC. Morris assumes, however, that the CTC requires ODC, and he therefore imagines “completing” my account of representation in the CTC by adding the idea that content is rendered “determinate” by stipulation: the content of *r* is the content assigned to it by a proper interpretation singled out by stipulation as *the* correct interpretation.

I think it is clear enough how Cummins would fill this gap. A device is an adding machine if and only if two conditions are met: first, that it is all right to treat it as an adding machine (that interpretation is *proper*); and, secondly, that it is *in fact* treated as an adding machine.

The result of this is the following simple suggestion: what it is for a representation to have a given content is a matter of stipulation, though, of course, there are restrictions – provided by the account of what it is for an interpretation to be proper – on what one can legitimately stipulate. [Morris (1991), p. 8, this issue]

This is not my view at all. In *Meaning and Mental Representation*, I was at pains to argue that the CTC has no use for what Morris calls determinate content.¹

II.1. WHY UNIQUENESS OF REPRESENTATIONAL CONTENT IS IRRELEVANT TO FC²

Certain ways of mapping numbers onto things demonstrate that the range of the mapping has a certain structure. The Kelvin scale is a way of mapping numbers onto temperatures. What the scale gives you is this: ratios of numbers assigned by the scale are ratios of the temperatures associated with those numbers. The units of the scale are arbitrary, of course, but the fact that temperature admits of a

ratio-scale is not. That's an important objective fact about temperature. One way to specify that fact is this: A ratio-scale is a proper interpretation of temperature. We could say: it is an important objective property of temperature that it is ratio-scalable.

Things have unique temperatures, but they do not have unique measures. Imagine someone who got confused about this. *"To deal with temperatures, I need to know their numerical 'contents.' The only way to determine the ratio between t_1 and t_2 is to determine the ratio between their contents. But you tell me content isn't unique. How can I do physics if I don't know what the real, unique numerical content of each temperature is? Does water boil at 270 Kelvin or at 77 Schmelvin? Which is it?"*

Answer: Both. The point is to get the ratios right. Any way of assigning numbers to temperatures that does that is all that is required. There is a "determinate content" *only relative to some scale*. But, of course, this does not mean that there is no determinate temperature to be measured.

Propositions, like numbers, are abstract objects with lots of interesting relations defined on them, such as entailment relations, confirmation relations, and so on. Certain ways of mapping propositions onto things demonstrate that the range of the mapping has a certain structure. For example, when propositions are mapped onto the data structures of certain computations in the right way, the computations emerge as proofs. According to the CTC (i) there are computations going on in the head; (ii) properly interpreted, these computations emerge as cognitions – i.e., as manifestations of the very cognitive capacities that Granny and other cognitive psychologists have discovered to be characteristic of the human mind. In short, the CTC is the hypothesis that the human mind is a computer whose computations are cognitively interpretable.³

For this to make sense, interpretations need not be unique any more than scales need to be unique for thermodynamics to make sense. The feature of a system's data structures that a propositional interpretation "measures" may be unique even though the interpretations are not. The point of interpretation is to allow one to see that the computation in question has the structure of a cognitive process. Cognition is epistemic constraint satisfaction, something defined on propositions, just as ratios are defined on numbers. To see ratios among temperatures, you have to properly "interpret" them – i.e., assign them numbers in a certain way via the process we call measurement. To see epistemic constraints being satisfied by computations, you need to properly "measure" them – i.e., assign contents to their data structures via the process we call interpretation.

Of course, the CTC may be wrong: it may be that computational capacities are not interpretable as interesting cognitive capacities,⁴ or that human brains don't support the necessary computations. But the possibility of multiple interpretations won't make the CTC wrong any more than the possibility of multiple ratio scales for temperature makes thermodynamics wrong. If you think interpretations must be unique, you mistake the explanatory role of content in the CTC.

II.2. SYMBOL GROUNDING

I know from experience that this won't convince The Opposition, because The Opposition isn't worried about how the CTC or any other empirical theory explains things. The Opposition wants to know what intentionality is. They know, "from their own case," as one used to say, that thoughts have ODC, and they know, intuitively, I guess, that this cannot be a matter of physical processes in their brains having the kind of abstract structure that is made manifest by the viability of a certain interpretation.⁵ The Opposition has a certain explanation in mind of the alleged fact that thoughts have ODC. I call it the Representational Theory of Intentionality (RTI) – the theory that a thought that p is a relation to a mental representation with the content that p . The RTI won't buy you thoughts with ODC unless you have mental representations with ODC. *That's* why The Opposition wants mental representations to have ODC, and they won't be happy with any account of mental representation that doesn't supply it.

We can get another handle on why ODC is immaterial to the CTC, however, by looking at what might actually supply ODC and asking what explanatory role, if any, that could play in a computational theory. So: Let's suppose that representations have ODC. Let's suppose they get this by being "grounded" somehow, as the jargon goes, either in adaptational history, or causally, or whatever. What I'll try to argue is that grounding is irrelevant to the explanatory needs of cognitive science, that something like what I called s-representation in *Meaning and Mental Representation* is all that matters to the explanation of cognitive behavior.

Here is a twin earth fable. There are two identical cities, Opolis and Twinopolis. I make a map of Opolis by taking an aerial photograph. I use this map to get around in Opolis. WW-III occurs, wiping out airplanes and the like. The map is carefully copied by successive generations of my family, giving them a huge advantage over others living in Opolis, etc. (The point is to hook me and my use of the map to Opolis in whatever way might be thought essential to grounding the map in Opolis.)

Now: Does the map represent Twinopolis?

Of course.

- (a) No map could be a better map of Twinopolis than this map, so grounding isn't *necessary* for representing.
- (b) Grounding also isn't sufficient for representing. Suppose Opolis changes over the years so that the map isn't quite accurate for Opolis, the city in which it is "grounded". The map remains a perfect map of Twinopolis, however. Surely the map represents Twinopolis, not Opolis, even though users of the map take it to represent Opolis, not Twinopolis.

Given (a) and (b), it is evident that accuracy is independent of "grounding",

i.e., independent of any facts about the user of the representation. But it is also evident that, so far as cognitive explanation goes, the crucial feature of a representation is what I've been calling accuracy, and that is independent of grounding.

So grounding doesn't matter to cognitive explanation. What does matter is just that representations have the right structure. The right structure is any structure that is interpretable as the relevant geography. No structure that isn't so interpretable will do. Proper interpretability as the relevant geography is therefore necessary and sufficient. The fact that a representation that is properly interpretable as the relevant geography will also be properly interpretable as a lot of other things is obviously irrelevant. Uniqueness of interpretation – what Morris calls “determinate content” – has no role to play in the CTC; its explanatory needs are utterly indifferent to the possibility of alternative interpretations, just as thermodynamics is indifferent to the possibility of alternative “scales”.

I think philosophers often lose their way here because they confuse what a representation actually does represent with what a system uses it to represent. For example: A chess system, at a certain point in the computation, needs a representation of the current state of the board. It constructs (or “activates”) r . It then uses r to represent the current state of the board. But, of course, r , the thing the system constructed, may not represent the current state of the board. That's how representational error arises: Representational error lives in the gap between what r actually represents, and what the system uses it to represent. But this rather obvious observation has the consequence that representation is independent of use, and hence of “grounding”.

Where does this leave us? If I'm right about the explanatory role representation plays in the CTC, then there is no reason to think that the representational contents it requires are stipulative, for there is no need, as Morris thinks there is, for stipulation to select a unique interpretation from among the many possible ones. The only issue is whether genuine cognitive functions are among the possible interpretations. This is no trivial matter: Finding *some* proper interpretation of a computational capacity is, I suppose, trivial, but it is not trivial to show that the brain (or some other physical system) has a computational capacity that can be properly interpreted as, say, natural language parsing.

The sense that interpretation is trivial comes from neglecting the fact that the point is to explain some antecedently specified cognitive capacity such as natural language parsing. The availability of trivial interpretations for any given computational capacity is just irrelevant to the question whether some (or any) computational capacity can be properly interpreted as the capacity to parse English. Philosophers tend to lose sight of this point because they really care about something else entirely, viz., explaining how consciousness can be full of what seem to introspection to be determinate thoughts. The CTC doesn't care about that at all, so far as I can tell. A corollary is that the CTC isn't going to be much help to philosophers interested in intentionality.

II.3. FC AND ODC

Certain accounts of representation *would* render satisfaction of FC pretty miraculous. It is no accident that these are all attempts to “ground” representations in a way that gives them ODC. Morris is right to think there is a problem about reconciling such accounts with FC, though, as we’ve seen, he is wrong to suppose that only grounded contents are “real”, and hence that the CTC can’t explain genuine cognition.

To see why externalist accounts of representation have difficulty with FC, we need only recall the affair of the grounded and ungrounded maps. Given that an ungrounded map is exactly as useful as a grounded one of the same “form”, it is difficult to see how evolution or learning could select only for the grounded representations. Given the fact that the form is what counts, it *would* be amazing if it turned out that there was, in some particular system, a correspondence between form and ODC.

Fodor recognizes a version of this difficulty when he rightly complains [Fodor (1987), p. 76ff.] that so-called two-factor functionalist theories cannot explain why causally grounded denotations should correspond to the denotations implicit in the satisfaction conditions that come along with the propositional contents assigned to representations via functional role. What he doesn’t acknowledge is that this makes a problem for *any* theory that wants to “ground” representations and also to acknowledge FC. Forms (and the functional roles that depend on them) do their stuff regardless of how grounded, as we’ve seen, and, according to FC, form is what counts. What mechanism guarantees that items with the right form (= computational role) have the right causal grounding? Of course, there is no such mechanism because it would be utterly pointless, as the parable of the maps shows. A system incorporating such a mechanism would have no conceivable advantage over one that did not.

You can see now where I agree with Morris. If you think representations must have ODC to do their stuff, then you’d better abandon FC and the empirical theories that, like the CTC, require it. If, like me you think ODC has no explanatory value in theories like the CTC, then you will cheerfully abandon it, and cleave tightly to FC. Thoughts, widely construed, may have ODC, but it is just bad philosophy of science to confuse these Putnam–Burge thoughts with the representations that loom large in the CTC.

II.4. INTENTIONAL SPECIFICATION AND THE EXPLANANDA OF THE CTC

But aren’t the explananda of cognitive science intentionally characterized, specified via attributions of ODC? If this is supposed to mean that cognitive psychologists self-consciously indulge in the language of wide content in characterizing the empirical effects they discover, then, no, I don’t think the explananda of cognitive science are intentionally characterized. (Try a twin-earth story on a

psychologist and you'll see how self-conscious they are about it.) On the other hand, it may be that linguistically expressible contents are essentially wide, in which case they can't be avoided in the specification of cognitive capacities.

Let this be so. Does it follow that a theory that hopes to explain cognitive capacities must suppose that mental representations have ODC? Of course not. What's required, to repeat the point again, is just that there be computational capacities of the brain that are properly interpretable as cognitive capacities. It matters not a jot to this hypothesis whether cognitive capacities are widely specified. It's whether the computations can be properly mapped onto the right propositional structures that matters. If they can, then we know that the computational discipline respects the epistemic constraints that define the target explananda. If they can't, no amount of "grounding" will help. If C has a proper interpretation in E (Earth), it is bound to have an analogous one in Twin-E. The availability of *both* interpretations is precisely what the CTC needs to explain why Σ would behave with equal success/failure in both environments.

It should be clear, then, that I do not, as Morris claims, hold that the CTC does not apply to intentionally specified processes of reasoning. What does worry me is that it may not be possible to specify cognitive functions in any vocabulary. By specifying a function, I mean fixing the value for each argument in a finite non-ostensive way. The restriction against ostension means you aren't allowed to do this: $f(x) = y$ if this human (or humans typically) return y on x . If that, or an infinite list, is the best you can do, then the process of analysis we call programming – breaking down specified functions into sub-functions with known computational instantiations – can't get started. No program, no computational theory. I hope it is obvious that this worry about the specifiability of cognitive functions is not designed to insulate the CTC from concern for intentionally specified capacities, as opposed to capacities specified some other way. The worry is that cognitive capacities might not be specifiable at all, not that they might only be intentionally specifiable. There is no reason to think intentional specification creates a special problem for the CTC.

Part III: Quibbles

The guiding ideas of *Meaning and Mental Representation* can be put like this.

(A) To understand what representation is, we'd best begin with what it does. As a philosopher, I have no theories about what representation does. There *are* theories about it, though, empirical theories that appeal to representation to explain various other things. So philosophers can do this: they can ask what representation must be if it is to do what one of these theories says it does.

(B) Moreover, there are philosophical theories about what representation is. So one can ask: if representation is what such-and-such philosophical theory says it is, could it do what such-and-such empirical theory says it does? One can ask, in effect, whether anything Philosophy has to offer can ground what Science has to

offer. I set out to discover whether current philosophical accounts of representation could provide foundations for the computational theories of cognition (CTC) that largely motivated current philosophical interest in mental representation in the first place.

Part (B) motivates the critical parts of the book, (A) the positive parts.

Morris appears to misunderstand the strategy when he claims that the argument for the positive theory – my account of what representation amounts to in the CTC – rests on the elimination of alternatives. That was not my intention: the account is supposed to stand on its merits, not solely on being the only straw floating. Nevertheless, he is right in thinking that my view is that there are no other straws floating. Someone who doesn't agree with me about that is, of course, likely to overlook the merits of interpretational semantics as an account of representation in the CTC while pursuing one of the other straws. Morris himself shows some sympathy with adaptational role theories, or, at any rate, a lack of sympathy for the arguments I urged for thinking content in the CTC must be ahistorical.

The basic point is simple: the CTC entails that computationally identical systems are cognitively identical. But computational states are not historically individuated. Hence the CTC entails that cognitive states are not historically individuated either.⁶ Morris thinks,

It is natural to continue from the points made so far that *content* is itself irrelevant to the kind of explanation which the CTC aims to provide. This seems in tune with Cummins' own account of the role of the semantic description of a process: namely, that it serves just to fix the *explanandum* for a computational explanation. [Morris (1991), p. 9, this issue]

It should be evident from the discussion above that I don't hold that content comes in only in the specification of the explanandum of a computational theory. Interpretation is a kind of redescription: Under interpretation, a computational capacity emerges as something else – the capacity to construct proofs, say, or to parse natural language. The right interpretation is like a pair of glasses that allows you to look at a computational process and see, say, language parsing. Thus, interpretation is the very core of the strategy, which is to show that computing – the right computing – is cognition, just as computing – the right computing – is calculation. Without the interpretation, you just have the computation; you have no way to tell if it is the *right* computation. Once you see what the role of interpretation is, you can see immediately both why content is essential to the CTC, and why historical (or any other) “grounding” is irrelevant.⁷

Appreciating the explanatory role of interpretation, and hence of content, in the CTC can, ironically, help us understand why Morris is right in supposing that content explanations aren't causal explanations. The CTC doesn't ring in content as a causal factor, it rings in content in order to redescribe certain causal factors – the computational ones – so as to be able to see them as cognitive factors. Like most really explanatory theories, the CTC isn't what I called a transition theory in *The Nature of Psychological Explanation* (Cummins (1983)).

That is, the CTC isn't primarily an attempt to state the laws of cognitive state transitions. Those laws – specifications of cognitive functions – are its explananda. The CTC is first and foremost a theory about what cognition *is*. Explaining laws isn't a matter of identifying causes of events; it is a matter of identifying mechanisms whose design gives rise to behavior satisfying those laws. Once this point is appreciated, it becomes clear that it doesn't much matter to the CTC whether or not we say that content is a causal factor. What matters is that there be data structures in our heads, that these are causal factors in virtue of their forms, and that there is a way of interpreting them that reveals the computations that manipulate them as cognitive processes.

Notes

¹ I don't have any use for ODC either. Morris, in effect, chides me for not providing an account of ODC. Since I don't believe in ODC, I must be excused from supplying an account of it. In any case, the issue is whether the CTC has any use for ODC.

² This section owes a good deal to Paul Churchland's discussion in Churchland (1979).

³ Philosophers think cognitive interpretability is a property nearly anything has got. Bright people all over the country have told me that canopeners and typewriters have it. Really? Just which cognitive capacities are they interpretable as having?

Here's a way to get rich: find a way of properly interpreting your canopener as a natural language parser. Automate this interpretation (which must be possible if it is proper), and give it access to a simulation of your canopener. Voila! A natural language parser! And so cheap!

You might also be interested in some Florida real estate I know about. . . .

⁴ It is certainly an empirical question whether all (or any) cognitive functions are computable. Nature abounds in systems whose state transition function is not a computable function. The brain, qua cognitive engine, may be one of these.

⁵ I suspect The Opposition of conflating (a) the introspectable phenomenology of having a thought, (b) the content of the thought, (c) the communicative content of its standard linguistic expression, (d) the representational content of the representations (if any) that various theories suppose to underlie having a thought. All of these are quite different.

⁶ I also pointed out in *Meaning and Mental Representation* that anyone who takes the representational capacities of the mind seriously is bound to think that certain mental structures are adaptations because of their representational capacities. But if this idea is not to degenerate into a tautology, we need a notion of representation that is logically independent of adaptation. The idea is not, as Morris seems to suppose, that the adaptational process leading to a representational structure must itself be a cognitive process. The idea is just that we can't understand why certain mental structures are adaptive in the first place until we understand their representational roles, for it is their representational roles that make them adaptive. You can't hold *that* and also hold that it is adaptational history that makes a mental structure representational.

In passing, it is worth noting that Morris gets Millikan badly wrong by ignoring the fact that, in her account, it is how representations are consumed, not how they are produced, that matters.

⁷ Morris compounds the confusion about the role of interpretation by thinking of computation, not as mediating changes in representations, but in mediating changes in their contents. See his discussion of (B)(iv), where what he says entails that the causal processes that mediate changes in data structures are not computational. To repeat: the CTC holds that the causal processes in the brain responsible for cognition are computational processes. If this is right, then there are computational processes in the brain such that they can be interpreted as such things as language parsing. Interpreting a computation is just interpreting its data structures, the things the computations are defined over.

References

- Churchland, Paul (1979), *Scientific Realism and the Plasticity of Mind*, Cambridge, UK: Cambridge University Press.
- Cummins, R. (1983), *The Nature of Psychological Explanation*, Cambridge, MA: MIT/Bradford Books.
- Cummins, R. (1989), *Meaning and Mental Representation*, Cambridge, MA: MIT/Bradford Books.
- Fodor, J. (1980), 'Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology', *The Behavioral and Brain Sciences* 3, pp. 63–109.
- Fodor, J. (1989), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA, MIT/Bradford Books.
- Haugeland, J. (1985), *Artificial Intelligence: The Very Idea*, Cambridge, MA: MIT/Bradford Books.
- Millikan, R. (1984), *Language, Thought, and Other Biological Categories*, Cambridge, MA: MIT/Bradford Books.