

ON TELEOSEMANTICS AND NATURAL MAPS

Comments on Cummins et al., “Representation and Unexploited Content”

Joe Cruz
Williams College
Philosophy & Cognitive Science

1.

Let me begin by signaling my enthusiasm both for the specific case offered by Cummins et al. against teleosemantics and for the overall framework from which this work derives. If the first approximation of the idea is that there will be material implicit in a representation that can be exploited by a cognitive agent that later acquires the right abilities to extract this material, and if this material looks a great deal like content, then the teleosemanticist will find accommodating it challenging. Moreover, the distinction between representation and indication is intriguing and important, and the discussion of structural transformation and isomorphism is illuminating. While Cummins has been urging these themes for some time now, it seems to me that they have not been sufficiently appreciated in the literature.

Here I will concentrate on some provocative puzzles raised by the argument against teleosemantics. My aim is to press these while maintaining a broad sympathy for the view. What we are left with, I think, is a rather more radical conclusion that Cummins et al. appear to be ready to accept.

2.

The Cummins et al. argument against teleosemantic accounts of content is straightforward. They propose that representations sometimes encode content that cannot be exploited by the cognitive system within which the representation is embedded. That the content is present is shown by the fact that the cognitive system can *later* acquire new abilities in exploiting content. These innovations may well enhance the overall system's adaptive fitness, but what made the content *content* in the representation before it was properly exploited was evidently *not* its role in selection.¹ That is why they say, "In general, content only becomes adaptive, hence a candidate for the kind of content-fixing selection contemplated by teleosemantics, when and if the ability to exploit it is acquired. ...There can be no selection for an ability to exploit content that isn't there" (p. 3). If this is right, teleosemantics faces a serious problem. Here is a kind of content — the unexploited kind — which cannot get its status as content from anything having to do with its adaptiveness because it has not ever actually been used for anything. Again, once it is exploited, it might be useful, but, as Cummins et al. elaborate, "That requires content to predate selection, and teleosemantics requires selection to predate content" (p. 4).

If we step back a bit more from even these modest details, we can readily see what Cummins et al. are up to. They evidently view the capacity to represent as implicating three relevant parts. Firstly, there is the environment that is being represented. Secondly, there is the representation of that environment. And thirdly, there

¹ The hope here and throughout — as it is in the Cummins et al. piece — is that these simplistic ways of talking about evolution and teleosemantics does not compromise the broad points made.

is the cognitive agent that exploits the representation in order to remain more-or-less well coupled and reactive to the environment that is represented. This tripartite division is clear in the map case. There is the terrain depicted by the map, there is the map itself, and, finally, there is the map-reader. The possibility of unexploited content emerges in light of the potential limitations in the map-reader's ability to extract information from the map. As they point out, someone might simply not know that their map encodes distance information.² The crucial point is that the information is already in the map, and whether the map reader knows it is utterly *irrelevant*.

What is attractive about the tripartite picture is that it answers to ready intuitions about what the problem of representation *is* in the philosophy of mind, and it also fairly well reflects the ways the problem has been theorized over the last thirty years, what with, for better or for worse, so many of the moves and methods coming from reflection on the philosophy of language. Notice that we need not be overzealous in our understanding of the third (representation exploiting) element. While an ordinary map reader is a sophisticated agent with dramatic representational powers quite independent of the map, what Cummins et al. have in mind is potentially far more sub-personal. They describe the system that exploits — or fails to exploit — a representation as a “harboring system.” That harboring system need not be aware that it is “using” representations, nor need it be powerfully cognitively autonomous, nor need it be otherwise fancy in any way. Presumably, a harboring system is any system that can reliably and usefully extract information from representations.

² After all, not all maps do, e.g., the maps on the placemats at my local diner.

I said above that the picture that Cummins et al. are working with fairly well captures the way that the problem of representation and content has been theorized recently, but there is an important caveat to that claim. Cummins et al. have a proprietary and revealing way of articulating the error made by the teleosemanticist, but a different way of putting it is that teleosemantics either mixes together two things that are alleged to be separate (the harboring system and the representation), or they mistakenly apply the apparatus of explanation by selection for an ability had by a harboring system to a thing without abilities, namely the representation. Millikan, at least, is clear on what she is up to. In *Language, Thought, and Other Biological Categories*, she claims that,

...A sentence, and every other typically intentional device, is intention in part by virtue of certain proper and Normal [sic] relations that it bears to a ‘producer’ or ‘producers’ and to an ‘interpreter’ or ‘interpreters.’ ‘Producers’ and ‘interpreters’ are cooperating devices that produce and use the intentional device and *that sometimes are and sometimes not* contained within the same individual organism (p. 90, italics in original).

Later she writes, “There are no signs without potential interpreters” (p 118). I understand the sense of ‘potential’ here as claiming that there must be some actual interpreters to have been relevant in driving the selectional explanation, even if there are no interpreters right now present to interpret a sign. And this plays into the hands of Cummins et al. who offer interpreters that acquire the ability to exploit content long after that content is properly representational.

3.

The discomfort we may feel at this point is due to the fact that we find something compelling about Millikan’s picture, too. In that spirit, I want to raise the question of

whether the tripartite picture Cummins et al. are working with is as innocent as it seems. What sense can be made of the middle element, the representation itself? Cummins et al. seem to require a stark separation between the representation and harboring system in the sense that, what enables the representation to have the content-bearing powers that it has, has nothing to do with the capacity for the harboring system to (presently) exploit that content. How, then, does the representation come to bear the content that it does? For an appraisal of this question, the map example surely leaves something to be desired. The problem with it is that that which remains potentially unexploited is content because someone else, even if only the mapmaker herself, was probably aiming to represent distance information. There are perfectly obvious reasons that a mapmaker might wish to represent both route-finding information and distance information in the same projection, but for representations in natural cognitive systems we are compelled to grapple with the general question of *why there would be* and *how there could be* exploitable but unexploited content at all. Two possibilities leap to mind, and both confirm what is compelling about teleosemantics.

One possibility is that the exploitable but not necessarily exploited content is an *accidental artifact* of the representational structure. It would be as if a simple-minded mapmaker set out to craft a representation that would allow for route finding, and accidentally drew a map that also encoded distance. Or, we can imagine a map designed to convey distance information on two dimensions that *accidentally* encoded information about elevation. These are not so likely in the case of maps, but perhaps we can stretch our credulity in envisioning these accidents. To see how remote the possibilities are, though, instead of drawing a map, think of writing out natural language directions in

terms of landmarks. The kind of media used to do this — words — are very unlikely accidentally to encode distance information. Words just do not do that without the person writing the directions intending them to do that. It is difficult to say precisely what is going on, but it looks very much as if words just do not do that sort of thing accidentally. Drawings have a better shot, but that is likely because we, as adults, might well encode routes and distances because we have access to both kinds of information to exploit in our drawing. It might well not be so in, for instance, the case of children's maps. The point is that an exploitable bit of content in a representation *may* have gotten there accidentally, but, if it is anywhere near as spectacularly useful as distance information on a map is, the serendipity will be shocking.

The other possibility is that the unexploited content appears in a representation because it is (something like) *entailed* by another, actually exploited, content. Take our route finding map. It looks nearly enough entailed by that map that the representation will encode a kind of temporal information concerning which landmarks will be encountered before other landmarks (when starting from a particular point). It may not occur to a map user to attempt to exploit this information, but it is there for the taking. Now, the status of this entailed content is troubling for the Cummins et al. case, because the teleosemanticist can seemingly rightly insist that the entailed unexploited content is content in virtue of being entailed by content that is actually exploited. This is to push the case back in Millikan's favor, where a sign is a sign only if there is a producer *and* an interpreter. The interpreter, on this view, need not be able to now exploit all the entailed features of the content, but a relation between the representation and the harboring system

secures the status of the entailed content as content again. There need not be any autonomous content had by the representation by itself.

We may draw these lessons in other media offered by Cummins et al.: Assuming the unexploited structure is not present accidentally in the neural network's input activation pattern, we should think that over learning trials the network extracts structure that was entailed by the input. How, though, did the input come to have a structure that would admit to entailments that would later be usefully exploitable by the harboring, more global network? Again, puzzlement over this question, I think, will tend to fuel a kind of teleosemantic intuition. It will seem as if the explanation is that the particular content entailments possessed by a structure are explained through an appeal to the adaptiveness of having just those entailments. Even if the *this particular* natural system needs some learning to come to exploit those entailments, the way they got to be there, the teleosemanticists might continue, is through having been actually and adaptively exploited.

4.

I think there is a non-teleosemantic reply to these worries, but they lead to a view that I am not sure Cummins et al. will be inclined to accept. The potentially radical conclusion that these elaborations of the Cummins et al. case are leading me to is this. We tend to view the structure of a representation as requiring some explanation with regard to its relation to the environment. Why not, though, instead view representations as more direct manifestations of the environment? On this view, they have whatever structure they have, and the contours of that structure are more or less continuous with

the contours of the structure of the environment itself. Thus, the explanation for the entailments possessed by the representation is that the environment itself contains those entailments. Imagine maps of the environment created not by something like imaging and environment, but by moving through that environment such that the movement literally shapes the parts of the cognitive system responsible for future route-finding. Since the movement was over distances, the representation contains entailments having to do with distance, even if the cognitive agent explicitly only exploits the route-finding properties of the imprinted map. This would avoid questions about entailment that invited a return to teleosemantics.

Alas, this may also be a road that leads to skepticism about representations themselves. I, at least, find that direction not unwelcome.

References

- Cummins, R., et al. (2005) "Representation and Unexploited Content." Presented at the 2005 Annual Meeting of the Society for Philosophy and Psychology. Wake Forest University.
- Millikan, R. (1984) *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.