Review of Cummins Representations, Targets and Attitudes

Ruth Garrett Millikan
University of Connecticut
Millikan@UConnVM.UConn.edu

Robert Cummins, Representations, Targets and Attitudes, Cambridge MA: MIT Press 1996 pp. Vii, 153, US$25.00 (cloth).

Cummins offers us a theory of mental representations, intentionality and propositional attitudes designed to answer the question what cognitive error is and why reference to mental representations can help to explain behavior. A large portion of Targets also deals with familiar positions on the naturalization of representation, such as conceptual role and other "use" theories of mental content, and causal\indicator theories, in an effort to show the need for a different sort of theory and what such a theory should accomplish. There is a great deal that is new and engaging in these critical sections, including an especially interesting interpretation of Fodor's asymmetrical dependency theory of mental content, and fresh ideas about holism. I will not discuss these critical parts of Cummins' book, though they compose much of its bulk.

Cummins' theory has as basic ingredients (1) the claim that a representation is always isomorphic to what it represents and (2) the use of teleology to make room between what a representation should represent and what it actually does represent. These are the same ingredients that I used in Language, Thought and Other Biological Categories[1] (hereafter LTOBC) to similar ends, but Cummins claims that he employs them differently, and such as to mend fatal flaws in my own approach. Graham Macdonald, in his review of Targets,[2] has explained how Cummins' criticisms rest on misstatements of my position, relieving me from having to do so myself. However, I am not convinced that Cummins' analysis differs from mine in the fundamentals, nor where it does differ, that the changes are an improvement. What I will do, then is to tell something of Cummins' positive position, clarifying its relation, as I see it, to my own, and suggesting where certain clarifications are needed and where certain inadequacies may lie. The book would have benefitted, I believe, had a much larger portion been devoted to clarifying its positive theses. A great deal remains mysterious.

A representation, on Cummins' view, is an isomorph. Full stop. A representation represents everything to which it is isomorphic, and represents each of these things equally. Cummins says that this is mathematician's usage of "representation" and he argues that it "presents no problem, provided it does not make representational content indiscriminate" (p.102). (More on that later). Nothing is gained by arguing over words. But everything in Targets that Cummins claims about "representations" can be understood better, I think, if one substitutes the word "isomorph," and our job should be to understand what Cummins is thinking. So I will just say "isomorph" in most cases where Cummins would say "representation" in what follows.

Cummins claims that reference to mental representations is explanatory of

---

[1] MIT Press, Cambridge MA, 1984.

[2] Forthcoming in The British Journal of Philosophy of Science.

behavior exactly because representations are isomorphic to what they represent. They "stand proxy" for what they represent, guiding the cognitive systems that use them in conformity with the things they represent. On the other hand, an isomorph of something, taken simply as such, is not the sort of thing that can be said either to be correct or to be in error. Only when there is something definite which an item is needed to be isomorphic with, only when there is an intended "application" for it, can it be evaluated as correct or incorrect. Correlatively, it is only when an isomorph is "applied" that there arises any "intentionality" or any "aboutness." An isomorph may be "true of", in the sense of isomorphic with, many things, none of which it is about. What an isomorph is about is whatever it is "intended" to be isomorphic with. What has intentional content is not an isomorph (Cummins says "not a representation") but rather an application or applying of an isomorph (p. 16). Cummins calls that which a given isomorph is "intended" to be isomorphic with its "target," and says that if we do not distinguish between what a given isomorph is in fact isomorphic with ("what it represents") and what it is intended to be isomorphic with ("its target") we cannot understand what "representational error" is.

So used, "intended" is clearly a metaphor. Cummins proposes to cash this metaphor using naturalized teleology, for example, using some theory such as my own in (LTOBC, Millikan 1993) that derives what a mechanism "intends" from what it (or its producer, etc.) was selected for doing during its evolutionary history. "The semantic content of an application [i.e., of the application of an isomorph] is that the representation [isomorph] hits the [is isomorphic with its] target" (p. 16). Using Cummins' flagship example, suppose that during a certain calculation made by a chess-playing system in the midst of a game, an isomorph of a certain board position is (internally) produced by the system. And suppose that the purpose of producing this isomorph is to make available for the system's use an isomorph of the board position now current in the game. This application or applying of that isomorph in that context has as its semantic content that the current position is as shown by the isomorph. Similarly, Cummins says, the visual system has as its function to produce an isomorph of the current layout of the surrounding environment. The isomorph that it produces has no intentionality taken in isolation, but its application carries the semantic content that the layout of the current environment is so and so.

Cummins takes it that whether an application of an isomorph is then employed in forming a desire rather than a belief or, say, a hypothesis or some entirely different sort of intentional attitude, depends on its role in thought. His formula is that a propositional attitude is determined by an isomorph applied to a target, this application yielding a semantic content or satisfaction condition, and <u>subsequently</u> given a role appropriate to its designated attitude.

Now, sometimes it might be the job of a representation to be an isomorph not of a property or set of relations but of a "proposition." For example, "an intender in the language-understanding system whose function it is to represent the meaning of the current sentence" (p. 19) may have a proposition as its target. The target might be "that Scrooge is rich" while what I actually produce is an isomorph of <u>my</u> being rich (Cummins says I might produce instead "an │I am rich│", p. 19). Then the content of the application would be <u>that</u> the meaning of the current sentence was that I am rich.

One thing this shows, Cummins says, is that even when the "representational content" [the isomorph's content] is a proposition, "representational correctness," that is, the correctness of the application, is not the same as the representation's being true. (Indeed, aiming an isomorph at a targeted proposition would seem to be, just, embedding it in intentional context, so it is not to wonder if the content and truth value of the isomorph would fail to equal the content and correctness versus error of the application.) The main point Cummins wishes to stress, however, is that a propositional attitude does not consist in an attitude toward a representation and does not derive its truth value from the representation [the isomorph] that it contains, not even when that representation happens to represent a proposition.

Having got this far, let me try, for purposes of comparison, to translate some of this thesis into the terminology of LTOBC. The idea seems to be that the isomorph-producing mechanisms have, on each occasion of use, a definite job to do, and that job is to produce an isomorph of something in particular. This is also the claim in LTOBC. LTOBC claims further, and Cummins seems to agree, that what it is the function of the isomorph-producer to produce is what the system needs to use in order to proceed with its activities, for example, "what [the system] proceeds to use the data structure [isomorph] to represent" (p. 7). That would seem to be, in the vocabulary of LTOBC, what the "consumer" aspect of the system needs to have to function properly when guided by the isomorph. In LTOBC this whole thought is expressed by saying that the function of the isomorph-producer is to produce an "icon" [isomorph] that the consumer can use to perform its proper functions, where helping with, being part cause of, the fulfillment of these same functions is, of course, a more ulterior function of the producer as well. (Things can have lots of proper functions, each helping to perform the next.) What this isomorph should be an isomorph of is thus determined by its user and its uses—not, of course, by its actual uses but by what its proper use would be. The isomorph that is actually produced may thus be either a correct one or an incorrect one.

Now Cummins calls the isomorph a "representation" whereas LTOBC calls it an "icon". And Cummins calls the isomorph-qua-produced-with-the-proper-function-of-being-isomorphic-to-a-certain-kind-of-thing—i.e., qua having the function of "hitting a certain target"—the "application of a representation", whereas LTOBC calls it an "intentional icon" or, in certain sophisticated cases, a "representation". What Cummins calls the "semantic content of an application," LTOBC calls the "semantic content of an intentional icon." So far, only a translation seems to be needed to make Cummins isomorphic to LTOBC, the most important thing being that in LTOBC intentional icons are typed using "shape"- (as taken from a certain reproductively established family) plus-target, not by "shape" (form) alone. That is, so far the difference lies only in the use of the term "representation," not in the analysis given.

But there are disagreements just around the corner. LTOBC claims that the proper function of the producer of an indicative or fact-stating representation is not merely to produce an isomorph of a certain sort of thing, but to produce exactly that sort of isomorph of that thing that the consumer knows how to use, hence to produce something that is isomorphic in accordance with a definite projection rule. And LTOBC doesn't call what is produced by the producer an "intentional icon", nor claim that it has any "semantic content", unless it is the sort of thing that is in the domain of the

projection rule for icons the consumer has been designed (by natural selection or by learning) to read. Cummins, on the other hand, takes the semantic content to be determined merely by isomorph-plus-target, whether or not the isomorph user can use that particular kind of isomorph, or can use it given the user's actual relation to the isomorph, for example, whether or not the isomorph is upside down vis a vis the proposed user (p. 99 ff.). This last does seem to me to be peculiar. Semantic content has turned out to be a function of intended use, yet to be oblivious of the intended mechanics of that intended use. The application of an isomorph can be correct even though the system that tried to use it breaks down as a result. Also, it seems that semantic content will not be determinate unless the case happens to be such that there is only one possible rule of projection mapping the isomorph onto its target. More on that later.

Another important divide is that Cummins takes the sort of difference in employment between isomorphs that constitutes their being caught up in different intentional attitudes not to be a factor in determining targets for the embedded isomorphs. This move is quite different from LTOBC, according to which reference to the role determining the containing attitude is essential in determining the semantic content of an intentional icon, indeed, where semantic content is determined by a different principle for indicative and imperative intentional icons. That is, for Cummins, applications first have semantic contents and then are given attitudinal roles, while in LTOBC the semantic contents are essentially abstractions from the attitudinal roles. Cummins does not explain why the intended role constituting the attitude an isomorph participates in differs from (or how it is related to) the rest of the intended "function" of that isomorph, such as to be irrelevant to determining its "target".

Cummins says that sentences in human languages are not "representations." They "convey messages" without being isomorphs of the contents they convey. He calls his theory the "PTR" or picture theory of representation—after Wittgenstein's famous theory in the Trachtatus, one would naturally assume. Yet language was the first subject of Wittgenstein's picture theory. Why does Cummins think that language does not picture when Wittgenstein thought that it did. It is not that Cummins thinks propositions can't be pictured, as we have seen. The answer lies, I believe, at the very root of Cummins' position, in his interpretation of the notion of picturing or "isomorphism". The difference between Cummin's and Wittgenstein's interpretations of picturing is important, I believe, so let me try to spell it out.

Cummins tells us that he is using "representation" and "isomorphism" in the mathematician's sense. What mathematicians are typically concerned with, however, is not mapping concrete existents onto other concrete existents, but mapping mathematical structures onto mathematical structures. A mathematical structure is a set of given abstract objects, plus a set of given relations between or among them, plus (or alternatively) a set of given functions or "operations" that transform these objects one into another. (Thus +2 can be considered as an operation that can be performed on any real number yielding another real number, rotate 20° clockwise as an operation that can be performed on any positioned oriented plane figure, move it three inches left as another operation on positioned plane figures.) For one mathematical structure to be isomorphic with another, there must be at least one possible one-to-one mapping

from the one structure to the other (often there are numerous such mappings) such that each object in the one set corresponds to a designated or given object in the other and each designated relation and/or function in the one structure corresponds consistently to a designated or given relation or function in the other, correlated relations or functions relating correlated objects. Relative to such a particular given mapping between isomorphic structures, each object in the one structure has a unique "image" in the other. Notice how many times the word "given" has entered here.

Now it is possible to generalize this mathematical notion of isomorphism, applying it not to sets of abstract objects but to sets of real concrete existents, and hence even, very loosely, to the concrete containers of such sets. Cummins talks of a "map" (he seems to mean a concrete physical object, a certain piece of paper with ink marks) being isomorphic to Chicago, but he also notes that strictly speaking it will not be Chicago but, say, some properties of the streets of Chicago or the elevations of various areas in Chicago (things Chicago contains) with which the map (its blobs of ink) are isomorphic (he says, which they "represent"). But as we have noted, a set of objects can be said to be isomorphic to another set of objects only when considered as part of a designated or given "structure". This means that one must designate not only exactly which objects are in each of the sets to be correlated, but also which of the various relations exemplified by them are to be given as part of each structure. Moreover, since there is often more than one way of mapping two given structures one onto the other so as to preserve an isomorphism, before any object or relation in the one structure can be said to be "the image" of any object or relation in the other, often certain initial correlations between objects and/or relations must be "given" as well as the structures themselves.

Now Cummins' "representations" seem to be concrete containers of concrete sets of objects considered as objects in structures. A map of Chicago or a data structure in a computer or a configuration of neural firings in someone's head is a concrete existent containing parts, properties and relations of various kinds, and if it is to be spoken of as bearing an isomorphism to something else, it must be considered as articulated into some definite set of contained objects with some definite set of relations on them. How is this to be done? Cummins asserts, for example, that if you take a certain map, it either is or is not perfectly isomorphic with Chicago! But there is an indefinite and perhaps infinite number of ways to cull from the things contained in that concrete map a set of objects to consider. Nor does the notion of isomorphism require that the members of the set be designated in any principled way. They could be any objects one felt like listing. And there is an indefinite number of relations one might consider among members of these sets of objects. (At the limit, every map contains the set that is itself, and given that monadic relations are relations too, it follows that any object in the world can be considered to be an image of any other, given that each has at least one property). I imagine that Cummins wishes us to consider only relations that are at least diadic. And perhaps there must be some "principle" that picks out the objects?

Let me digress for a moment to apply this to a point raised earlier. Suppose that we have a determinate target. And we have a concrete complex object (or event or whatever) that is to be used as a representation of it. But there is an indefinite number

of different projection rules by which the object can be interpreted as being isomorphic to the target, or not being isomorphic to the target, or indeed, as being isomorphic to anything at all you please. We are not to look at how the user part of the system is going to react to the object, nor to how it was designed to react to objects of that family of objects (for example, that reproductively established family of objects--see LTOBC). But still the application of object has a semantic content. How is this semantic content determined? The applying of the isomorph to the target is supposed to have the content that the representation hits the target, that is, that it is a correct representation of the target. What will determine whether this particular one is indeed a correct one? For example, what determines whether that structure is a correct representation of the current chess position, given that it could be a representation of any chess position whatever, on the right reading, and given that the way the representation consumer will or is designed to read it does not equal the right reading? I leave it for Cummins to explain in his response to this review (this journal, this issue).

Returning now to the difference between Cummins and Wittgenstein on picturing, Cummins often talks as if his representations [his isomorphs] were concrete isolated objects, such as a particular ink-and-paper map of Chicago, applied one by one to targets. On the other hand, Cummins sometimes talks not just about one concrete object or set "representing" another, but about a "representational scheme". He says, for example, that certain systems may be forced to make errors because these errors are "unavoidable given the expressive power of the representational scheme" (p. 23). This would seem to be a reference to some particular scheme of mapping the system is designed to use, some particular way of setting up correspondences between the representations and their images. Given this, one could talk more reasonably about whether this particular map is or in not isomorphic to Chicago.

A representational scheme makes determinate which abstract or concrete objects and which relations and/or functions are contained in each of the paired isomorphic structures, and it must determine enough correspondences between objects and between relations or functions to make determinate all the rest of the imaging. The most common kind of representational scheme correlates a set of abstract objects, members of a set of representational forms or types, with a another set of abstract objects, represented forms or represented propositions. Here the inner structure of each representation is not what carries the main burden of the mapping, but there are relations between or transformations of the (possible) representations that correspond to relations between or logical transformations of the things represented. For example, the transformation rotate the bee dance 20° clockwise corresponds to move the angle of the direction of nectar clockwise 20° closer to the sun, and so forth, for the set of all possible representations in the bee dance scheme. The complexity that makes for interesting isomorphism in the bees' representational scheme, unlike that of a good map of Chicago, shows up in the relations among possible well formed representations rather than within each representation.

But apparently it is structure within representations, not between them, not structure in a representational scheme, that really interests Cummins. Wittgenstein would have said that "John loves Jane" was a "picture." It mirrors one nonsymmetrical

relation with another nonsymmetrical relation so as to contrast with "Jane loves John" (you can't use a symmetrical relation here). And it mirrors sameness by sameness and difference by difference as compared with "June loves John", difference of grammatical subject reflecting difference of agent, sameness of grammatical object reflecting sameness of patient, and sameness of relation to the same verb reflecting sameness of relation between agent and patient. But Cummins says that "expressions in a natural language...have semantic properties conventionally" and that "representational content [i.e., isomorphism] is irrelevant to their conventional meanings" (p.90-91). And over and over he emphasizes that "representations are isomorphic [individually] to what they represent" and he talks about "structures" as "the relations between elements in a representation" (p. 93) rather than between representations in a system.

But I do not see how to put content into Cummins claims unless one assumes that mental representations are defined as such relative to some particular scheme of representation that the biological system is designed to use. And if it is the consistent use of representational schemes that is important, then individual mental representations might be relatively unstructured affairs, and (certain portions of) natural language will surely count as exemplifying a representational scheme that helps to carry linguistic meaning. Sentences in natural language will function as isomorphs.

I think that one reason Cummins wants to separate the content of a representation from its target is that he initially thinks of a representation (rather as Fodor does) as some sort of structure type ("shape") that you should be able to pick up physically and move to another context without changing what it represents. For example, the same structure type should represent the same thing whether it is located in the belief box, the fear box or the wish box. But it doesn't look to him as though the full semantic content of a structure used for mental representation does remain the same through changes in context. For example, "X, Y, and Z detectors might all do exactly the same thing in the presence of their targets because detectors are, in the standard and simplest case, devices that just 'light up' in the presence of their targets" (p. 74). If there is any content that remains the same here through changes of context for detectors it must be "disappearingly small" (p. 63), something like a yes or a no, all the rest of the content of the detector's flash being contributed by the target, by what the flash is supposed to detect. Now add to this Cummins' vision that the capacity to represent must be a function of the inner structure of a representation. Then everything represented by the representation will have to correspond to some one of its variable parts. Consider again the representation by the chess playing system of the current chess position. Its variable parts correspond to aspects of a possible chess position, but there is nothing both variable and movable here that tells what actual chess position it is supposed to represent the location of.[3] That part seems so be contributed by the point of having the representation, and it seems that the representation itself could be moved around to another location where it might be supposed to represent, say, "the last game position but one".

---

[3] unless, as LTOBC explains (compare "intentional signals"), one considers the location of the individual representation as a representing variable).

Contrast this with the position of LTOBC. There the content of a representation is considered to be radically context dependent, so that it would never cross one's mind that the same "shape" has to represent the same when used by another part of the system for another purpose. And there, not just the variable parts and aspects but also the invariant parts and aspects that remain the same over the transformations that define a representational scheme are considered to have a representing function. There are no variants corresponding to the sun or the hive or the nectar that a bee dance represents. These parts or aspects of what is represented are invariant for all bee dances. That it is a certain sort of figure eight that talks about sun and hive and nectar in B-mese rather than a rectangle or a sound is quite arbitrary of course, apart from the peculiar evolutionary history and engineering problems of bee design. And it is arbitrary that it is, specifically, an angle that shows direction. What is not arbitrary is that the whole system of representations can be mapped one-to-one onto its chosen set of possible representeds, preserving certain relations among the objects in each domain. That is the objective isomorphism that holds between the two kinds of structures. That is the aspect of the representing that is not merely "conventional". And that is enough to allow us to use mental representations to do quite a lot of explaining of how action is conformed to the world, especially if we treat the dynamics among mental representations as displaying some of the isomorphism, that is, especially if we add some of Cummins' insights in his earlier book, Meaning and Mental Representation, to the mix.[4]

[4] My gratitude to Rob Cummins for his patient and good tempered help with this review. Thanks also to Virgil Whitmyer for helpful discussions about it.