

# Cummins—or something isomorphic to him

Brian Cantwell Smith<sup>1</sup>

*NB: This is very much a first draft—loose & disorganised. But maybe it will give a feel...*

## 1. Introduction

Everyone is right. Or anyway that's what I tell my students. "Look," I say, "these things you are reading were written by dedicated, intelligent people, who have devoted their lives to studying these issues. and are trying to tell us about valuable, hard-won insights. Think of those insights as paths in a forest—and of the text as the author's attempt to clear it off, and show it to us. Problem is, people write in *words*; and words are blunt instruments—intellectual bulldozers, big bruisers that cut wide swaths. Rare persons—poets, mostly—can wield words with enough finesse to clear a delicate path without doing too much collateral damage. But when most of us write, even if we think we're just cutting a trail, in fact we're mowing down trees, ripping up the earth, and sewing all kinds of destruction.

"So here's my advice," I go on. "Don't assume this text is written using *your* words, and then try to figure out whether it is true or false. It will almost certainly be false. Instead, assume that it's true—that it represents a genuine and important insight—and tell me what language it is written in. Ignore all the ancillary damage; that will grow back. Tell me what the author was on to—what they were so excited about. Tell me, if we were to follow their path further, where it would lead."

I rehearse these platitudes here only because they've proved bracingly difficult to honour in the case of Cummins *Representations, Targets, and Attitudes*. Like many of us, Cummins has worried about representations for decades. He is full of insight, detail, and systematic analysis. Yet I confess in this latest book he also says things that, at least if interpreted in my own dialectic, seem downright loopy. Such as the word 'two' cannot represent *two-ness*. In fact words can't represent at all. Representations, according to Cummins, have to be isomorphic to their targets. So the word 'two', being unary, can't represent duality. And that's not all. Cummins doesn't think that a picture can represent a left-hand, either, as opposed to a right hand. Nor, if Igor and Isaac are twins, can a picture represent Igor and not Isaac (or the other way around).

That representation requires isomorphism is a strong claim—so strong, in fact, that once you reach it, in chapter 7, it's tempting to go back and read the book a sec-

---

<sup>1</sup>Copyright © 1999 Brian Cantwell Smith  
Computer Science, Cognitive Science, & Philosophy  
Indiana University, Bloomington IN 47405 USA

Draft of August 10, 1999  
Comments welcome  
smithbc@indiana.edu

ond time, substituting the term ‘isomorph’ for ‘representation,’ to see if it makes more sense. But that strategy doesn’t work, either. It doesn’t illuminate an interesting path, that is, either to assume that Cummins by ‘representation’ what most of us mean, or to take his remarks merely as comments on the cognitive utility of isomorphic structures. Problem is, I have only a guess as to what his basic insight is. So this commentary is going to end up with a query. I’m happy that Cummins will be able to respond, to set us straight.

My strategy, in other words, is not going to be to ask whether Cummins is right that representation is isomorphism, but to figure out what Cummins could possibly think representation is, such that he think it is isomorphism. Why does he think it’s useful to peel off the isomorphic part of the intentional onion, that is, and give it special treatment?

## 2. Points of agreement

But before getting to questions, I want to start with some points of appreciation.

First, Cummins is interested not in an abstract, idealized conception of representation, but in finite, concrete, resource-limited concrete mechanisms. He is concerned with engineering, with the consequences for design of real-world tradeoffs between accuracy and efficacy (or efficiency). Better to suffer a few false positives in your predator detection routines, he says, and get the hell out of there, than to wait until you’re absolutely sure that that’s a piranha. He’s right: that’s what animals do, that’s how programs work, that’s the right design stance. And he’s right, too, that this kind of pragmatism has consequences for a theory of truth.

Second, Cummins is right to fight for a serious, substantive notion of content—one that doesn’t (vacuously) devolve into “**x** means whatever happens when it is used,” or “**x** refers to whatever it is pointed at.” And he realizes that honoring that may require rejecting the common methodological assumption that theories of cognition will necessarily be *causal*. I commend him for this paired object- and meta-level commitment to making room for content. It is quite likely, in my view, that intentional systems will ultimately prove to be theoretically distinctive, among natural entities, exactly in virtue of the fact that constitutive accounts of them will not (at least in any *local* sense) be causal accounts. It’s not hard to see why this might be true: intentionality and content are precisely achievements that allow systems to be existentially oriented towards that with which they are not causally engaged. So why should we expect constitutive accounts of them to be narrow?

Third, at least in some places Cummins has a salutary recognition of the normative character of intentionality. This is betrayed (among other places) in his insistence on focusing on representational *error*—on the gap that, in no small number of cases, comes between how things are, and how creatures take them to be. It would appear that Cummins is a straightforward (even promiscuous) realist, so for him acknowledging this gap, while crucial to his project, isn’t horribly expensive. Since exploring the metaphysical consequences of representation is one of our explicit work-

shop’s aims, however, and since I am not as sanguine a realist as Cummins, I want to go on record as saying that I take a recognition of the profound importance of the distinction between how we take the world to be, and what it is actually like, to be critical (if anything even more critical) for non-realists, too. So I welcome Cummins’ placing error—and gaps—on center stage.

Fourth, Cummins is interested in representational *use*—with *how representational systems actually work*. This is a critical point. Though he categorically rejects what he calls “use” theories of content, he nevertheless *is* interested in use itself—how concrete representations lead real-world mechanisms to behave. In fact I will tentatively suggest that his curious isomorphism suggestion arises from a combination of this mechanical intuition and a theoretical mistake. But even if that suggestion is wrong (very possible), his focus on concrete use is surely correct.

### 3. Situated Cognition

Focusing on use is not unique to Cummins, of course; it is a theme of many other writers, including those interested in **situated cognition**. But in situated company Cummins’ proposal is non-standard. It will be helpful in triangulating on his account, therefore, to list five characteristics widely agreed by situated theorists to hold of real-world representation (that’s their ‘representation’ bulldozer, not his).

The first, stemming from the commitment to use, is that representation is an **activity**—a property of processes, not of static structures. Many (but not Cummins) think that assigning content to an enduring structure is either a derivative shortcut, or else a recognition of a commonality that endures across, but is ultimately grounded in, activity. Second, representation is not only in general highly **context-dependent**, but in particular is often, and perhaps even always, **indexical**. As Barwise & Perry emphasised almost 20 years ago, words like *you*, *here*, *now*, etc., are the norm, not complicated deviants. (A note for the rest of the workshop: this inexorable indexicality, I believe, ultimately stems from the ontological fact that leads physical laws to be expressed in differential form—something I hope we can talk about.)

These first two “situated” properties—activity and indexicality—go together. For whereas general rules or patterns are associated with (relatively static) word *types*—rules that Barwise and Perry label “meaning”—the specific interpretation of a word use is a function not of the type, *nor even of tokens* (as shown by Perry’s example of two deaf mutes, so poor that they share a card on which is scrawled “I’m a poor deaf mute; can you give me some change,” which they alternately hand to passers-by), but of *actual situations of use*—of utterances, in the case of language; more generally, of *events*.

Third, going along with these issues of context-dependence, indexicality, and a focus on activity or events, situated representation is seen to be largely *implicit*. What the stomach sends the brain is at best something like “hungry!”—meaning *I am hungry*, of course, but the fact that it is the stomach’s owner who is hungry need not be represented explicitly, but rather is established by the inalienable context. In

fact the stomach needn't even say *hungry*; all it needs to do, as every programmer knows, is simply to raise a flag, or send a one-bit signal along a dedicated wire. In that context, the majority of the action's content—the part that Perry calls the “issue”—is established by the architecture. All that needs to be flagged “explicitly,” as it were, is the residual polarity: of whether (and possibly how much) the agent is, at that moment, hungry.

In general, working out the division of labour across these four contributors to total intentional content—(i) what part belongs to or is determined by a use or utterance, (ii) what part belongs to or is determined by the particular structural token or instance of a signal or structure that is created or employed in that use, (iii) what part belongs to or is determined by the type of which that token or instance is an instance, and (iv) what part belongs to or is determined by the general architectural configuration in which the whole scenario is embedded—these are extremely interesting, intricate questions, that have yet to receive much attention, let alone adequate treatment. There has been some initial work on the linguistic case, and (as Cummins well knows) computer programs embody tremendously sophisticated but untheorised approaches to more general data structures. But a general story awaits philosophical reconstruction.

Fourth, the engaged, dynamic, architecturally-fitted kinds of mechanisms that come out of this sort of picture are often so purpose-specific that it is not clear that any *static* semantic evaluation is applicable (such as traditional notions of content, reference, or truth). Rather, the appropriate semantic evaluation may be governed by norms applying to the activity in which the agent is engaged. This focus on activity—as for example in Millikan's “pullmi-pushyu” representations—has led some writers to believe that “representation” is not even the right notion under which to understand the applicable form of intentional force. This issue—by no means minor—ties into another of what our workshop topics. To give it a name, what is at stake is whether engaged activities are subject **statical norms**: i.e., norms on states, such as *true* or *false*, or whether they are subject to **dynamical norms**: i.e., norms on activity or process—such as *evolutionarily adaptive*, or *good*.

Fifth, this focus on dynamic, architecturally-embedded, implicit representation has fueled an increasing interest in what is called *non-conceptual* representation: representations that, perhaps in part because of being purpose-specific and having contextually-determined content, fail to meet standard notions of productivity, systematicity, etc. (for example, fail to meet Evan's Generality Constraint). No one denies (or at least not many people do—though Cummins, interestingly, is someone who would) that *some* of our human representational capacities are “symbolic”, “linguistic”, “conceptual”, recombinant, etc. (choose your weapon). But once one has embraced the idea of non-conceptual content, serious questions arise as to what *conceptual* content is for. Cummins answer, I take it, is that it is *only* for communication. (This is also, I take it, is one of Deacon's concerns, about which we will also be talking later in the week.)

I say these things about a situated approach to representation because Cum-

mins’ theory (representation as isomorphism) is so markedly different. Representations, according to Cummins:<sup>2</sup>

1. Are structures, and have content qua structures, not as *processes* or “in use”;
2. Have what structural content they do *qua* types (not qua tokens or in use)—and for that matter also *represent* types, not particulars (i.e., their content is no more than type-specific);
3. Are neither implicit nor context-dependent nor indexical;
4. Are neither implicit nor architecturally-dependent; and
5. Are not governed by dynamic norms.

Representations are also, as it happens, on Cummins view, not conceptual; at least most of them won’t be (since isomorphic structures are not in general combinable into more isomorphisms); but nor are his forms of representation exactly what people interested in non-conceptual content are primarily focusing on.

Does this mean that Cummins is uninterested in the concerns of situated theorists—even if he doesn’t identify his project with theirs? Not at all, I believe. Rather, he moves all these issues—indexicality, the determination of particular content by contingent local circumstances, architectural dependence, etc.—onto targets and applications. In that setting, he must surely agree that they are important concerns. For example, I think I might be able to convince him that that there are issues to work out in his account of targets, between what is true of particular tokens, and what is true of particular *uses* of particular tokens, what is true of creations of tokens that differs from what is true of subsequent uses of previously-created tokens, and so forth. For some kinds of content go with one of these things, some with another. (For example, imagine keeping a “Don’t erase this!” sign in your desk, which you put up, from time to time, on the whiteboard, when you don’t want the cleaning folks to clear it. In this case, what matters to the interpretation of the word ‘this’ is the *use* of the sign; virtually nothing holds of the token qua token. Whereas in other cases—for example a photograph—the referent is determined by the token, not by its use—presumably since referents are causally implicated in the process of creating photographs. And so on.)

In sum (this may be too strong, but in that case Rob can disabuse me of it) *most of the issues that other people interested in use have focused on under the heading of “representation”—especially those people in the loosely-affiliated situated cognition contingent—has to do with what Rob calls targets, applications, and attitudes, not with representations.*

So Cummins is using the term ‘representation’ for something else. Does he just take it to be a conceptual synonym for ‘isomorphism’? No; it is a substantive thesis, for him, that representation is isomorphism. But clearly the word ‘representation,’ in that claim, doesn’t mean, for him, what it means for the rest of us. So what

---

<sup>2</sup>At least this is how it seems; the book isn’t entirely clear, but what he says and the examples he uses suggest this.

does it mean? What is going on?

#### 4. Use, Error, and Independence

The basic structure of Cummins' theory is that the phenomena (states, events, propositions, whatever) in the world that systems are directed towards—the phenomena that many of us would say they *represent*—are what he calls **targets**. In response to a forthcoming review by Ruth Millikan, he says (at the end): “I really have no dispute with her theory construed as a theory *whose target is target fixation*.<sup>3</sup> My point is simply that the same machinery won't do as a theory whose target is representational content.” *Targets*, that is, not representations, are what systems are intentionally directed towards; they are the things that the contingent particularities of agent's situations play a role in fixing; they are what is concrete and particular; they are what are indexically specified by perhaps implicit aspects of system architecture. What representations are, in contrast, are *what the systems represent those targets as being like*. So representations, on Cummins' view, are not, as it were, “what we talk about,” but “what we say about what we talk about” (except of course it isn't *talking*, but general intentional directedness, that is at issue—more on that later).

The way Cummins gets at representational content is via a discussion of error. At the beginning of chapter 7, where he introduces his positive view,<sup>4</sup> he says “The central thesis of the foregoing is that error is a mismatch between the content of a representation and the target of a particular use or application of it.” So far that seems coherent—or at least familiar (it has something of a subject / predicate ring to it, to say nothing of the topic/frame distinction from linguistics). And error is ubiquitous—not universal, but always, according to Cummins (and I agree with him on this) potentially present. He strongly objects to the idea, implicit in various familiar causal theories, that *any* situation can be identified, even ideally (such as an “optimal perceptual setting” or “normal conditions”), in which error can be assumed to be absent, and representational content therefore *identified* with target. What the target *is*, roughly, is the situation in the world that it is the function of the system to represent when it “tokens”—which is to say, I take it, when it *produces* or perhaps *uses*—a representation. And as we have seen, we won't go for wrong (for now) if we assume that targets are fixed in something like a Millikanesque way (though he argues for a spate of alleged improvements to her account).

It is Cummins' next sentence that is the kicker: “A theory of representation content that does not explain how *r*'s content is fixed in a way that is *independent* of how *r*'s targets are fixed is bound to mishandle error” [emphasis added].

That representational content must be *independent* of target content is one of the central themes in Cummins' book. Except that we can state a stronger version. Cummins repeatedly says that representational content must be specified (deter-

---

<sup>3</sup>Emphasis added.

<sup>4</sup>“Representation and Isomorphism”

mined) *independent of use*. (Since targets *are* determined by use—targets are what it is that it is the function of a use to represent—the broader claim that representational content be specified independent of use in general subsumes the narrower claim that it be specified independent of target in particular.)

Now here we start to get places where I either fail to understand what is going on, or else strongly disagree. In particular: That representational content not be *identical* to target or use I understand. That representation content be *independent* of target and use I do not.

Start with the agreement. What Cummins claims, and what I believe, is that if the representational content of every micro situation or event is simply *identified* with either the target or the use itself (e.g., the uses' causal consequences), then sure enough there will be no gap—hence no error, hence no substantive notion of content, hence the project will fail. But—and this I suppose is my first real criticism—there is surely a vast distance between non-identity and independence.

In fact it is madness to think that, in order for one thing  $\mathbf{x}$  (representational content, in this case) not to be *identical* to something else  $\mathbf{y}$  (target or causal consequence),  $\mathbf{x}$  must be *independent* of  $\mathbf{y}$ . This is so even if we recognise that 'independent' is a bulldozer of its own—a veritable Caterpillar D9. There are presumably as many flavours of dependence and independence as there of possibility and necessity: at least physical, nomological, metaphysical, and logical—if you believe that modality comes in stripes. But my point holds, I believe, on whatever variety you choose. Consider my arm. My arm is manifestly not *identical* to my body. But neither is it *independent* of my body. In fact notice that thinking of my arm's being independent of my body is rather gruesome—which shows that you have to break the part-whole relation in order to get independence to apply. In general, that is, my arm is *partially constitutive* of my body. And modulo *recherché* machinations about mereology, part-whole relations are completely unproblematic.

Here's a more serious example. Consider fitting a curve to a set of points. And ask about the relation between the curve and the points (at, say, each point's x-value). If you draw a random wavy line subject to no constraints except that it pass exactly through each and every point, then, sure enough, there would be no distance—no gap—between the curve and the point. But if, as usual, one fits a *constrained* curve—a straight line, say, or a Gaussian, using a least-squares method—then there *will* be gaps. In fact the gaps are critical: they figure in the measure of the total error, the very error which is minimised and used to determine the line. In general, the line won't go through *any point*, exactly; there will *always* be gaps. But from the existence of those gaps can we conclude that the line is determined *independently* of the points? Absolutely not. That would be crazy. In fact the situation is as close as possible to the opposite of that: *in spite of the gaps, the line is determined by the points*.

Here's the point: while Cummins and I agree that content (more generally: normative virtue) cannot be identified with "whatever happens," I see no reason to suppose that one cannot put *constraints* on content so as to introduce the requisite ten-

sion between content in a particular case and the target or use in that case. For example, consider something that John and I, at least, are interested in: objectivity. My sense is that normative virtue stems ultimately from a (non-standardly-rich) sense of objectivity and ‘stabilising’ the world: that what makes representations true, overall, is whether they add up to a picture of “the whole world”<sup>5</sup>—or rather, more radically, whether they add up to an oriented form of living in the world, since I think the most profound norms are dynamical, not statical. But perhaps someone else would prefer some other overarching constraint than objectivity. The point is merely this: if there is nothing to constrain the content of a given use of a representational token, then sure enough you don’t have enough machinery to have a substantial gap. But tying the content of that particular use to the content of other uses, tying it to overarching constraints such as stabilisation or objectivity, fitting it into a coherent scheme—any of these things could be exploited in order to make gaps (and hence error) substantial. *Independence*, in particular, is not required. As far as I can see, that is, Cummins’ sequitur from “needing a substantive notion of error” to “requiring independence of content and use” is simply a bad argument—the conclusion doesn’t follow from the premise.

Let me make a methodological note, in passing. This inference of Cummins seems to me an instance of something rather common: a regrettable but widespread tendency in analytic philosophy to assume that things that are different must be independent. It reflects a form of **conceptual absolutism**—a kind of “forced clarity” (vaguely Cartesian, perhaps) that, far from illuminating genuine issues, instead (I believe) takes leave of phenomena and encroaches on a worrying kind of scholasticism. In “Analog and Analog” (a favourite paper of mine), Haugeland talks about “second-order digitality”: the idea is one of concepts or types being perfectly clear, wholly distinct, and completely independent—in a way that charge and momentum arguably are, for example, and that chutzpah and braggadocio are palpably not. An absolutist commitment to conceptual independence, I believe, is a case of mistaken (or at least over-eager) second-order digitality. In my own work, I have wrestled for a long time with a similarly absolutist independence assumption: the thesis underwriting the formal symbol manipulation construal of computation that syntax or operations proceed independently of semantics. For reasons with which Cummins would sympathise, syntax is not the *same* as semantics (holding that they are the same requires evacuating semantics of substance). But to believe that the non-identity of syntax and semantics warrants an *independence* claim is (in my view) far too strong—ultimately theoretically misleading, and again leading to incipient scholasticism. When I first noticed this similarity, it struck me as an interestingly parallel; on reflection, I am not convinced that it is a *different* case, at all. I am not entirely convinced, that is, that the two cases are not, at root, the same: that “use” is

---

<sup>5</sup>There are very substantial issues of how to reconcile this claim with the requisite respect for inexorable and ultimately ineffable pluralism. Solving that is one of the main things I try to do in *On the Origin of Objects*.

not effectively syntax, and that “content” is not semantics. But I digress; maybe this is a question we can discuss over drinks. For now—to come back to the topic at hand—let me simply say that, as far as I can tell, with respect to the intricate interplay of use, content, target, etc., the authentic phenomena in the world are nowhere near as “clear and distinct” as is dreamt of in Cummins’ philosophy.

## 5. Isomorphism

Where are we? Well, my story is about half done. On the one hand, I feel as if I have an inchoate sense of why Rob feels he needs something on which to ground a notion of content that is independent of use (even though I believe it is based on a flawed inference). But I haven’t yet said anything about why isomorphism should recommend itself as that independent answer.

My intuition is that in order to answer this, we need to look at Cummins’ commendable commitment to use (not to use-based content, remember, but to use itself). Based on my sense that use is what matters to him, I am going to go out on a limb and hazard a guess as to why he is attracted to isomorphism. This sort of rank speculation would never do in a scholarly publication—but since Rob is here to set me straight, I hope that even if my guess is wildly off-target, his correction will be illuminating. I might say that one reason I am making this guess is because it goes some ways towards explaining Rob’s stance (e.g., as illustrated in his final chapter) on why language is purely communicative, and not representational at all (on his notion of ‘representation’).

The idea is this. Symbols, language, and the like are extremely abstract—very “distanced” from the entities in the world they are about. If you are actually going to build something to do something, it has to *interact* with the world—physically engage with it. To do so (this intuition underlies some accounts of non-conceptual content) you need representations—activities, structures, whatever—that are massively more detailed than mere words. Even vastly interconnected word- or concept-based “knowledge networks” aren’t (in general) remotely detailed enough to do the task (this is something, incidentally, that has been borne out in AI practice). And if, as a kind of tour-de-force, you hand-build linguistic representations that *are* sufficiently detailed to serve as the final arbiters of action, they will in general be so verbose as to fail the effectivity norm: they will be too complex for simple mechanisms to handle. Imagine trying to put all the information on a map of London into words, for example—allowing yourself talk of splines, bezier curves, x and y coordinates, etc. All you would get would be a useless tome of impenetrable axioms, that no one could reasonably use.<sup>6</sup> Maps are just *so much easier*.

One reason maps are easier is that direct physical engagement often involves physical *coupling* with the world. And coupling, it turns out, as a consequence of the way physics works, *itself tends to involve isomorphism*. Your route, through London, is isomorphic to the layout of the streets you are travelling on—a platitudinous but

---

<sup>6</sup>Unless, of course, they were simply to use it to reconstruct the map.

nevertheless substantial fact about physical travel. The shape of your hand, similarly, at least at the points of contact between it and the thing you are picking up, is isomorphic to that part of the thing you are picking up that your hand is in contact with. And so on.

So my hunch is that Cummins realises—and is motivated by the fact that—real-world agents need structures that enable them to act in ways that are “close to the details” of the worlds they engage with. Moreover, such agents don’t give a hoot about whether the structures that engender that appropriate behaviour have a causal (or teleological or functional or whatever) history tying them to the part of the world they engage with. All they need is something that will “steer them correctly.” And the *simplest adequately detailed structure for steering* is something that is isomorphic to the world it steers you through.

So here is my guess about the basis of Cummins’ isomorphism intuition. As we saw above, he thinks that he needs something *independent* of use,<sup>7</sup> that will nevertheless *guide* use. Isomorphic physical structures—free of teleological or design or other baggage—seem like the simplest candidate. Hence, I believe, his proposal.

Now I’ve already said that I don’t think the “independent of use” requirement is right (it is vastly too strong). Moreover, if I am right in my guess about steering being the issue, then surely what matters is ultimately not that a structure *itself* be isomorphic, as that it lead to *behaviour* that is isomorphic (I’m going to relax the constraint that the behaviour be isomorphic presently—but bear with me for a moment). To see this, let me suggest a modification of the autocart that Cummins describes in his chapter 7. In his original example, a cart is driven by a card that happens to contain on it a slit that is isomorphic to the track that the cart needs to follow, and the cart is engineered so as to exploit this isomorphism. The card, according to Cummins, is a representation of the track, because the slot on the card is isomorphic to the track.<sup>8</sup> But suppose we changes things so that the track that the cart needs to follow has a repeating structure—say, a sequence of path fragments that all happen to be geometrically identical. Imagine as well that there is a grooved wheel—a little like the wheel in a music box, except with continuous rather than discrete markings—such that one rotation of the wheel is able to exactly lead the cart through one of the periods in the geometrically-repeating pathway. Surely, I would have thought, in this particular case the wheel would be just as good a “representation” of the whole path as the card—since it would lead to exactly the same behaviour. And note that it would lead to the right behaviour over the *whole* path, not just one period of it, because of being circular. By construction, however, the wheel is not isomorphic to the path. So on Cummins’ account, it cannot be said to represent the path—or any multi-period segment of it. Nor can you say that one *rotation* of it

---

<sup>7</sup>Falsely, of course, as I argued—but I am trying to follow him here.

<sup>8</sup>It is also a representation of every similar track, of arbitrary different scale, orientation, and reflection—to say nothing of all other objects of that shape, and so on ad infinitum. But leave that aside for now.

represents one period—since rotations are kinds of uses. And so Cummins would be forced to deny that the wheel represents the path, even though, from a pragmatic point of view that I would have thought he would like, it is just as good as the card at driving the cart along the path.

If this example cuts any ice with Cummins, it seems to me that it should suggest that (again, just for a moment) isomorphic *behaviour* is what really grounds his intuition, not isomorphic *structure*. Behaviour is use, of course—and Cummins is committed to eschewing use, even to being independent of use, so this is not a formulation he will be likely to accept. But I've already argued that he needn't cut himself off from use so absolutely, and so I hope to have softened his allergies on that score a little bit. Isn't it just possible, I might suggest to him, that he should adjust his definition of representation so as to allow the wheel to be a representation of a repeating, periodic path?

But of course Cummins is a savvy guy, and will recognise that I am offering him the eighth day in the Scopes monkey trial. Because, as soon as he agrees to go down *this* path, it becomes easy to argue that it isn't *isomorphic* behaviour that matters, but *appropriate* behaviour. Except that "appropriate behaviour" is a big-ticket item—what does *appropriate* mean? What kind of bulldozer is *that*? But that takes us back into norms—and, for us here at the workshop, maybe into drinks.

## 6. Conclusion

Here, then, is my best shot at understanding the path Cummins' is trying to lead us down:

1. He wants us to appreciate the overwhelming importance of the fine-grained causal driving of things, and yet
2. He recognises that identifying the content of a representation with the fine-grained behaviour that it engenders (without any other constraint) vitiates any attempt to have a full-blooded sense of content.

On the face of it, those two requirements together pose a challenge. Isomorphism, I suggest, is his proposed solution. My response is (i) that both requirements are critical, (ii) that the second doesn't imply that content be independent of use, and (iii) that if we exploit over-arching dynamic norms, we can be open to vastly more intricate, use-based accounts of representational content.

What does this have to do with language and communication? Only this: that words, I agree, or anyway I think (I don't know whether *Cummins* would agree) need fleshing out into vastly more intricate and fine-grained form before they can actually let us engage with the world's virtually ineffable detail. That's in a way a corollary of the first of the two points just listed. I don't think he's right that words don't represent. But I do think he's right that words are a high-level, abstract form of condensing and summarizing intentional directedness—a high-level, abstract form of orienting to and describing the world that our capacities for intentional directedness direct us at. They are useful; on that we're agreed. But in order to use them—in or-

der to ground them—they need to rest on a vastly more intricate, sensitive, detailed capacity for physical coupling. A creature that had only words, without that fine-grained detailed way of steering its way around the world, wouldn't be able to *do* anything delicate, wouldn't be able to find or follow trails.

Remember: words are bulldozers. They do serious damage.

---

What follows are fragments on other issues that should be integrated into the foregoing somewhere

### **X. Fragment 1**

At several points, Cummins says that the target is what it is the function of system to represent—and that this target fixation must be independent of representational content. He also talks about nested intenders, which give rise in turn to nested targets—his solution (along with indexicality) of how to provide productive intentional power. But I don't see how the nesting story goes—in particular, how it fits with his overarching independence mandate.

Suppose that, walking on the beach, I say to myself “There's John; I need to return his wife's library book.” And suppose I am mistaken: that what I have seen, and thought was John, is a piece of driftwood, not a person at all.<sup>9</sup> So far so good; the driftwood is the target; John is the representational content; and there is a gap (because I've made a mistake). But now what about “his wife.” What is the *target* of ‘his’ (or its isomorphic analog)? If I read Cummins right, the target of ‘his’ should be whatever it is the function of the tokening of ‘his’ to represent. But who (or what) is that? If it is *John*, as it seems to me there is good reason to suppose, then that target is a function of the *content* of the representation which causally preceded it—but that seems to violate Cummins' independence mandate. As much as I at least can glean from the book, Cummins would claim that the target of ‘his’ would be the *driftwood*. But that seems bizarre. And the oddness propagates, the further the inferential stream goes. What about “library book”? What is the target of that? Or suppose one entertains a whole sequence of thoughts, starting from a failed identification. It would seem, on Cummins' view, as if the entire sequence would be semantically malformed. But suppose one starts off on a wrong footing, but then comes to a very important realisation—as often happens, when meandering in thought (e.g., suppose, on reflecting on the library book, when was on non-conceptual content that you realise that it bears a striking resemblance to some Indian philosophy that you read back in college). Just because the thought *started* with a false identification doesn't mean that the subsequent realisation cannot be true, semantically sound, or normatively important. But I don't see how this would go, on a Cummins-like story.

---

<sup>9</sup>“John” is a symbol, of course, not a representation on Cummins' typology, but consider a case where diagrams or isomorphic pictures are used instead.

In sum, on Cummins' view, is the target of an agent's action totally determined by the nesting and interplay of the *targets* of the history and surrounding causal context? If that is so, then what is the point of the representational content at all? It starts to look as if it has no bearing—as if it is almost epiphenomenal. I would have thought that the target  $\mathbf{t}$  of an action  $\mathbf{a}$  (even if we believe in targets) would be a function perhaps not of the representational content of the representation  $\mathbf{r}$  applied to  $\mathbf{t}$ , but a function of the representational content of representations leading up to action  $\mathbf{a}$ . But if representational contents are in general to be determined independently of use altogether, then I take it that targets are to be independently of representational contents.

What is happening here?

### **Y. Fragment 22**

«Not yet written. The issue has to do with the general nature of (mathematical) isomorphisms, and how to establish the appropriate fine (intensional) grain that is necessary for semantic (representational) purposes: which properties and relations—and even objects—in the representation map to which properties, relations, and objects in the content. There are lots of issues: how the right mapping is determined, for example. This part of the critique could get technical; I decided to defer it for now, since the larger issue of whether isomorphism is the right subject seems prior, and of more general interest. But I confess that I am not quiet about whether, other questions aside, simply saying that “representation is isomorphism” is anything like *specific* enough. This is true notwithstanding Cummins' comments that it is not problematic for a given representation to represent lots of things, and that only one of them need be the one that the target matches.»

----- end of file -----