

Putting Concepts to Work: Some Thoughts for the Twentyfirst Century

JESSE PRINZ AND ANDY CLARK

Abstract: Fodor's theory makes thinking prior to doing. It allows for an inactive agent or pure reflector, and for agents whose actions in various ways seem to float free of their own conceptual repertoires. We show that naturally evolved creatures are not like that. In the real world, thinking is always and everywhere about doing. The point of having a brain is to guide the actions of embodied beings in a complex material world. Some of those actions are, to be sure, more recondite than others. But in every case the contents of thoughts still look to depend, in some non-unique but vitally important way, on the kinds of doings they support.

1. Rationalism and Pragmatism

1.1 Introduction

Concepts, according to Jerry Fodor, are for thinking, not for doing. To command a concept of X is (Fodor claims) to be able to think about X's qua X's and not to accept (in your thinking though not necessarily in your practice) any substitutes. Fodor thus believes that, to a first approximation, possessing a concept is not about possessing any complex of abilities, or even any specific body of displayed knowledge. Instead, it is simply about being able to think about the things or states of affairs in question. And however this is to be cashed out, it is not to be cashed out in terms of what anyone can or can't do, or any body of knowledge they might be called upon to display.

Taken lightly, this can sound quite reasonable. And indeed Fodor makes some play with the avoidance of some rather brutal forms of behaviorism in his paper. But Fodor's project is not simply (and properly) to remind us that there is more to thought than brute action patterns. It is not to remind us that two agents may differ in their thinkings even while agreeing in all their observed actions. It is not to remind us that imagination and planning are as real as running and jumping. No. Fodor's project is to sever any putatively constitutive links between thinking and any kind of doing. But seen in this light it is clear that this is a doctrine we should not seriously contemplate, on pain of losing our grip on what concepts are for, and why we bother to ascribe them.

Address for correspondence: Cognitive Science Program, Sycamore Hall 0014, 1033 E. Third Street, Indiana University, Bloomington, IN 4705-7007; Department of Philosophy, CB 3125, Caldwell Hall, UNC/Chapel Hill, NC 27599-3125, USA.

Email: jesse@subcortex.com; andy@indiana.edu

According to Fodor, rationalists say that the possession condition for a concept is the ability to think about what that concept represents. The possession condition for MULE is the ability to think about mules. Concept pragmatists tell a different story. For them, possession conditions are defined in terms of other kinds of abilities. Fodor singles out two examples: the ability to sort things, and the ability to draw inferences about things. Both accounts fail. Fodor regards this as an indictment of concept pragmatism more broadly, and—given the prevalence of pragmatism over the last hundred years—an indictment of twentieth century philosophy. Concept possession, Fodor suggests, occurs (and must presumably somehow be assessed) independently of any considerations about epistemic matters, about what the agent knows or (if you favor knowings how over knowings that) what the agent can or cannot do. But there are problems both with the detail and the spirit of Fodor's argument, as we shall try to show. We think pragmatism can be defended. In fact, we favor a supercharged pragmatism that links concept possession to action. Concepts are action-oriented.

1.2 Useful Thoughts

As Fodor sets things up, rationalism and pragmatism are radically different views about what concepts are. To become a rationalist is to abandon pragmatism. We think this is a false division. There are significant differences, perhaps radical differences, between our view of what concepts are and Fodor's, but the rationalism/pragmatism divide is more blurred than Fodor would have us believe.

For the rationalist, having a concept is being able to think. This is not the whole story, however. Rationalists of Fodor's persuasion also claim that concepts are mental particulars. They are internal representations that can play a causal role in the determination of behavior. Otherwise thinking would be epiphenomenal. So the rationalist definition of concepts had better be supplemented: have a concept is to have a mental vehicle that allows one to think, where 'vehicle' refers to a concrete particular. On this reading, rationalism is a role-filler definition. It says concepts are things that play a role (allowing us to think), and it is more or less neutral about what those things are. This neutrality narrows the gap between concept rationalism and concept pragmatism. For it is possible that the thinking-role is filled by things that also underwrite other abilities. Perhaps representations linked in an inferential network play the thinking role. Perhaps sorting mechanism comprising perceptual templates do that job. Perhaps the capacity to think about a property is implemented by vehicles that govern our ability to physically interact with instances of that property (think of the property of being a Frisbee). Perhaps there are a variety of different kinds of vehicles that allow us to think about the same thing in different circumstances.

This does not show that rationalism is equivalent to pragmatism. The neutrality of rationalism may entail multiple-realizability. That means that different kinds of vehicles can fill the thinking-role. Some of those vehicles may look like the kinds of things pragmatists talk about, but others may not. Perhaps God can think about

things without any sorting capacities or inferential roles. Perhaps we can think about our own experiential states in a similarly direct way. If so, pragmatists would be wrong to say that it is constitutive of being a concept that we can sort or infer. But what a shallow victory for rationalism! Suppose the vast majority of human thinking abilities were implemented by sorting mechanisms. There would be a token identity between most instances of most concepts and the mechanisms that pragmatists get excited about. If most concepts were implemented by sorting mechanism that were built up from perceptual templates, then it would turn out that empiricists were right about the vehicles of thought. Concept rationalism would be implemented by vehicular empiricism! This shows that Fodor's dichotomy between rationalism and pragmatism is misleading. The label 'rationalism' is inappropriate, and the theory itself may be too neutral to be of much theoretical use. The major competitor for concept rationalism, defined as the ability to think about things, would be a form of anti-representationalism, according to which we have no thoughts.

The interesting debate is not between rationalist and pragmatists, we contend. Rather, it is a debate about vehicles. It is a debate about *how* we think. It is a debate about the kinds of mechanisms that allow us to mentally moor onto the world. Fodor's view is that concepts are semantically unstructured symbols in an amodal language of thought. This atomistic theory of vehicles is not entailed by 'rationalism', but it is the major issue underlying his critique of the twentieth century. Fodor's three arguments against pragmatism are really arguments in favor of atomism. We think atomism is a bad theory of how most concepts work. Maybe God has atomistic concepts, but that's not the way that terrestrial creatures think. We favor a very different story according to which concepts are best understood as rough-and-ready complexes of abilities, whose vehicles are context-sensitive articulated inner structures of the broadly connectionist/prototype-style that Fodor so roundly rejects. This is a familiar debate, and our own positions are already well-expressed in print (Clark (1993), Prinz (2002)), so we shall keep our comments here (section 2.1) brief. Since the real point of Fodor's paper is to offer a single new argument, we shall concentrate on that (section 2.2). We will end with some more general observations, contending that pragmatism is not only viable, but essential (section 3).

2. Fodor's Objections to Pragmatism

2.1 Two Dogmas of Fodorism: Analyticity and Compositionality

Fodor's main arguments against pragmatism do not focus on the supercharged pragmatism that we will be talking about. His targets are more conservative: any theory according to which lexical concepts decompose into meaningful parts. Fodor has been arguing against such theories for the better part of the last 30 years (the twentieth century dies hard!). His arguments have distilled down to a

pair of major worries, sometimes presented as a kind of Scylla and Charybdis for fans of structured concepts.

If concepts have parts, those parts can be defining or non-defining. If they are defining, then one had better be able to distinguish conceptual relations that are definitional from those that are not. That requires an analytic/synthetic distinction. We know, from Quine, that such a distinction is not in the offing. We need to say that all true features are defining or none are. If we say none, we are back to Fodor's atomism. If we say all, concepts become unshareable, because they change identity with every new belief. So concepts cannot be built up from defining features. Suppose we go to non-defining features, for example the prototypical features used for sorting. The problem with these is they don't compose. We know that concepts must compose to explain productivity and systematicity, therefore concepts cannot be built up from non-defining features. The inferential role theories and sorting theories that get profiled in his critique of the twentieth century have traditionally fallen prey to the analyticity and compositionality problems respectively. We have discussed these worries elsewhere, so we will keep our remarks very brief (see Clark (1993), Prinz (2002)).

Suppose concepts are collections of action-oriented abilities. Suppose they are structured representations of features that allow us to sort, draw inferences, and physically coordinate behavior with objects in the world. How do concepts of this kind hold up against Fodor's dilemma? They certainly do not qualify as definitions, so they are not likely to combine compositionally. Take the concept 50 POUND EGG. If asked to consider such a monstrous entity, several features might come to mind. We might imagine that it is quite durable, in contrast to more familiar eggs, and we might imagine that it was layed by a huge creature, perhaps a dinosaur. Neither of these features, durability or layed by a dinosaur, are plausible components of the concept EGG or the concept of WEIGHING 50 POUNDS. They emerge from background knowledge. Fodor thinks this is a problem. He thinks the only thing we really know about a 50 pound egg is that it is an egg and that it is 50 pounds. Adding extra information prevents concepts from combining in a way that will support systematicity and productivity.

We are unmoved by this objection. First, as a matter of psychological fact, we do come up with emergent features. Giving advice to authors, Bill Gass put it like this:

Throw down any pair of terms like dice; speak of arrogant bananas; command someone, as Gertrude Stein once did, to 'argue an earnest cake,' and the mind will do more than mix them in its ear. It will endeavor a context in which the command is normal, even trite. Our grammar gives rules for doing that, but sometimes these are no more than suggestions (Gass, 1976 p. 35).

The phenomenon that Gass is talking about is pervasive. Even the most mundane word pairs will summon ideas from beyond their boundaries. This needs to be explained. Second, it would be very imprudent not to. If concepts are action-oriented, they need to prepare us for the interaction with objects in the world.

If we learn that there going to be a 50 pound egg around the next turn in the road, we better be prepared to cope with it. That may involve moving it, and, worse still, it may involve a confrontation with the creature that is looking after it. If we do not think of these things in advance, we could be in trouble. It would be brazenly irrational not to use any information at our disposal to anticipate what the egg is going to be like. This point applies quite generally. If you are going on a blind date with a Polish philosopher, you might plan on talking about logic. If you are shopping for an affordable car, you might look at compacts. If you have to give a 20 minute talk, you might want to speak quickly. Emergent features are a guide to life. A mind whose ideas are mixed like bricks rather than chemicals would be an inhuman mind, and it would be a mind unfit for this world.

This, of course, is not an answer to the compositionality objection, but rather a plea for the claim that the objection must be answerable. We clearly cope with emergent features, and cannot cope without them. The answer given in Prinz (2002) can be succinctly stated. Compositionality, like productivity and systematicity, is an ability. It is something we can do if we are in a pinch. If we are not sure what a Balinese chiropractor is like, we can just combine our default ways of thinking about Balinese people and chiropractors (a dancing charlatan?). In the limiting case, we can just combine labels. A Balinese chiropractor is a person who has the profession known as 'chiropractics' from the place called 'Bali'. These words serve as anchors that allow us to speak and defer to people in our linguistic community. As long as we have this much, we can generate a representation for an unfamiliar category using purely compositional means. That's all the compositionality we need. In ordinary cases, we have much more to go on. A prodigious concept combiner might imagine that a Balinese chiropractor cures patients by dancing on their backs.

Supercharged pragmatism may be able to accommodate the compositionality requirement, but, in so doing, it may run head-on into a cousin of the analyticity objection. If concepts are guides to interaction with category instances, then they must be adaptive. Instances vary. Eggs come in different shapes and sizes. A Balinese chiropractor is unlike a Samoan chiropractor, who you certainly don't want on your back. Theories that restrict concepts to defining features aim to keep conceptual constituents constant. Definitions tend to remain fixed. Once we give up on the definition project, and adopt a view of concepts that allows for considerable variation, we end up with concepts that are unsharable. As we move from person to person or context to context, the representation of any given category is likely to shift. That wreaks havoc on communication and psychological generalizations.

In response, we note, first, that as long as we agree that concepts have broad content, we have as much concept sharing as Fodor has. Second, we have quantifiable similarity of concepts, because we can talk about two people representing a category using some of the same features. Third, why should we insist that concepts are shareable in a strong sense? We need to explain coordinated behavior. For example, we need to explain why people point to the same object,

when they use the word ‘mule’. We need to explain why people go to chiropractors when their backs hurt. We need to explain why people follow the same steps to boil an egg. None of these similarities require perfect identity of concepts. And differences are easy to multiply as well. If you think chiropractors are charlatans, you will go elsewhere with your backache. Once we start to regard concepts as action-oriented, variation becomes as essential as commonality.

We think that the compositionality and analyticity objections only have force for those who are in the grips of two dogmas. Concepts do not compose most of the time, and concepts are not perfectly shared. Perhaps Frege is to blame for these dogmas. Twentieth century pragmatism departed from Frege in various respects but, we think, it did not go far enough. The notion of immutable senses in the third realm must be replaced by senses that are sufficiently pliable to guide us through the terrestrial realm.

2.2 Jumping through Hoops: The Circularity Objection

Fodor’s analyticity and compositionality arguments are old news. The major innovation in his paper is the circularity argument. It has two forms. One is directed at pragmatists who identify concepts with sorting abilities, and the other is directed against pragmatists who emphasize inferential roles. We have something to say about both.

As it is most naturally interpreted, the argument against sorting abilities goes as follows. If concepts are sorting abilities, then concepts must be individuated by ways of sorting. But it cannot be the case that ways of sorting are extensionally defined. After all, necessarily co-referring concepts, such as TRIANGLE and (CLOSED) TRILATERAL, correspond to the same extensional sorts. So ways of sorting must be individuated conceptually. What kind of sort manifests a concept C? It cannot be one that simply sorts Cs, because a necessarily co-extensive concept would do the same. It must be a concept that sorts Cs in virtue of their C-ness. If TRIANGLE is a sorting ability it must be one that sorts triangles according to their triangularity. But to do that, we must possess a concept of triangularity. So sorting in virtue of triangularity requires possession of a TRIANGLE concept. Sorting in virtue of triangularity cannot, therefore, be the possession condition of a TRIANGLE concept. That would be circular.

We have several replies. First, circularity isn’t always a bad thing for concept possession. In particular, it is not a bad thing for simple recognitional concepts. Red things might just be those things that we sort in virtue of their redness. Second, even in this case, there may be a way to cash out of the circle. Sorting by redness is sorting by a color experience, and color experiences may have functional or physiological bases that can be specified without mentioning the colors that they represent. Third, in the case of triangles and trilaterals, there is a straightforward account of what distinguishes the sorting abilities. In one case, we sort by first identifying angles, and in the other we sort by identifying lines. In other words, sorting abilities may be structured, and two co-referential sorting abilities may be

distinguished by the sub-abilities that they comprise. Fourth, Fodor's locking theory is notoriously vulnerable to cases of necessary co-reference. The nomic correlation between a Mentalese symbol and triangularity is symmetrically dependent on the nomic correlation between that symbol and trilaterality. If concepts are atomic, there is no way, on Fodor's view, to distinguish TRIANGLE and TRILATERAL (or WATER and H₂O). To accommodate these cases, he is forced to abandon atomism and argue that at least one concept in each of these pair is structured. Which one? Who knows. In sum, the cases that Fodor presents seem to have a simple explanation on sorting theories, but no explanation on atomistic theories.

This makes us think we must have Fodor's objection wrong. How could the solution be so obvious, and so damning for the theory that he favors? So we feel compelled to offer another interpretation of his argument. It is hinted at by his brief remarks on Wittgenstein. Wittgenstein wants to analyze the concept of PAIN in terms of behavioral criteria. As we read Fodor, this analysis faces a dilemma. If PAIN is analyzed in terms of behaviors described without reference to pain, such as wincing, the analysis will fail. Someone can understand what wincing is without knowing what pain is. If, on the other hand, pain is analyzed in terms of behaviors known to be indicative of pain, then the analysis will be sufficient for knowing what pain is, but it will be circular (amongst other problems). Fodor relates this argument to his circularity argument against sorting theories. Perhaps he thinks that argument can also be understood as a dilemma. If concept C is a sorting ability, either it is an ability to sort Cs as such, in which case it is circular, or it is an ability to sort Cs as Fs (where F is a feature other than C), in which case it is a bad analysis.

The second horn of this dilemma is hard to interpret in the case of TRIANGLE. For that concept can be decomposed into the complex feature of HAVING THREE INTERIOR ANGLES, and the features in that complex are not equivalent to triangularity, but they are jointly sufficient for triangularity. The analysis succeeds. Likewise, MULE might comprise THE OFFSPRING OF A DONKEY AND A HORSE. But what will the sorting theorist say about HORSE? It cannot be analyzed in terms of HAVING HORSEISH APPEARANCES, so they will say it must be analyzed in terms of having appearances thus and such, where thus and such is specified without the concept HORSE. One might sort horses by sorting NEIGH SAYING QUADRUPEDS. The problem with this is the analysis fails. One can have a NEIGH SAYING QUADRUPED concept without a horse concept. Put differently, HORSE is either analyzed into defining features or non-defining features; the only defining feature for HORSE is HORSE; and non-defining features won't get the concept right. It is this last bit that we want to challenge.

On Fodor's own theory, HORSE is an atomic concept that represents horses in virtue of being locked onto horsehood. Locking is achieved by some naturalistic relation, such as carrying information in a particular way. Call this relation R. Now, forget about atomism for a minute, and suppose that HORSE is made up of NEIGH SAYING QUADRUPED along with some other features. It is perfectly possible for this bundle of features to bear relation R to horsehood. On Fodor's own account, the bundle would now qualify as a HORSE concept. It would represent horsehood despite the fact that it is built up from features that are not defining of horsehood.

Would it be *the* HORSE concept? Fodor likes to use the definite article when talking about concepts, but we think that such talk must be taken with a grain of salt. There are many ways of having a horse concept, because there are many ways of having thoughts about horses. Non-defining features can certainly do the trick if unstructured atoms can. Neither has the resources to uniquely single out horses by description, but both can get locked onto horsehood.

In sum, we think Fodor's circularity argument against sorting theories may boil down to the contention that bundles of non-defining features cannot be equivalent to the concepts that they are supposed to explain. That contention rests on Fodor's failure to appreciate that non-defining features may play the same informational role as atoms, and, thus, they too may allow us to think about properties that they do not define.

What about the circularity argument against inferential role theories? Fodor focuses his critical energies on Peacocke's inferentialist analysis of the concept AND. This is a good strategy since AND is one of the few concepts that inferential role theories seem to handle exceptionally well. But it is an odd strategy for Fodor, because he wrestles with AND on his own account. For him, most concepts are individuated by the properties to which they are nomologically locked. But it seems very odd to think that there is a property of conjunction out there that lawfully triggers AND tokens in us. So, Fodor has generally conceded AND to the inferentialists. He usually says that they get AND right. His new argument is designed to show that inferential role theories get AND wrong, and that leaves him without an account of AND too. If his argument goes through, it means that everyone is in trouble. Fortunately, the argument is not decisive.

For Peacocke, a person possesses the concept AND if she finds certain forms of inference primitively compelling in virtue of their form. For example, an AND possessor would be compelled by transitions from $p \text{ C } q$ to p , where p and q are sentences, and C is a word in their language, such as 'and'. Fodor wants to know how a person comes to find such transitions compelling. The only plausible answer, he thinks, is that they are compelling to those who interpret C as AND. The inferences cannot be the possession condition for the AND concept, because finding them compelling presupposes possession of that concept.

We offer the following way out for the inferentialist. Suppose AND is an innate operation over mental representations that serve as thoughts (thought representations). It is an operation defined by its role. When two thought-representations are linked by this operation, one is disposed to token either of those thought-representations independently. Learning the word 'and' is a matter of mapping that word onto this mental operation. To find sentences with 'and' compelling, one does need to have the concept. But having the concept is just a matter of having this operation over thought representations. Its possession conditions can be specified without mentioning conjunction (think of the engineer's description of a logic gate). Primitive compulsions in forming natural language inferences are evidence for possession of AND, not constitutive, though linguistic mastery may be constitutive of concept possession in other cases. Conjunction is an inferential skill.

Its mastery is constituted by the inferences we draw in thought. Fodor's rationalist proposal, according to which every concept is an ability to think about something, makes little sense here. Would Fodor let his theory of concepts commit him to the existence of a property of ANDNESS to which we become nomically locked? If the investigation of concepts were a guide to first philosophy, then Fodor would be right to call himself a rationalist. But surely his psychology cannot bear such metaphysical weight.

3. What Are Concepts For?

3.1 Cognition and Action

Now let's turn to the general spirit of Fodor's argument. The guiding idea, recall, is to sever the apparent links between concept-having and any kinds of abilities or dispositions to act and (presumably) speak. One immediate upshot of this general position is to open the conceptual doors to some truly bizarre possibilities.

Imagine that we encounter a creature (call it EGG) that, by whatever non-epistemic criteria Fodor favors, counts as having the concept 'egg'. And suppose that this creature one day asserts that it is wrong to boil eggs, and that all living beings should always do their utmost to stop such barbaric practices, whatever the cost. But despite this, and despite the lack of any perceptual or motor impediments, or countervailing beliefs (such as the belief that people will always do more and more of anything that EGG tries to stop, or whatever) or other mental impediments (such as chronic apathy) EGG never lifts a finger to stop people seen dropping hen-made spheres into boiling water. Nonetheless, EGG is in full command (as far as we can tell) of the notions of boiling and of wrongness, deploying the concepts much as we do, and having no trouble being accepted into the community of folk who talk about wrongness and boilings. Should we simply nod wisely, accepting that despite all this, the concept of 'egg' is indeed fully present in the mind of EGG? To do so would surely undermine the point of concept ascription, and talk of concepts, in the first place. A child displaying this behavior would be corrected, and efforts made to bring their usage into line. (This practice of bringing usage into line might even be called 'word learning', and Fodor is happy (see his footnote 7) to translate claims about concepts into claims about words.). Yet despite this whole constellation of belief, desire, opportunity and motive, EGG never acts, never even has so much as a disposition to act. This gives us no cause, on Fodor's account, to withdraw our ascription of the concept. Since speaking is, of course, itself a mode of acting, we may add that this same conceiver will never (after that one original utterance) join any community of discourse on the topic of egg-boiling-avoidance, even assuming it has all the other concepts and knowledge needed to do so (and all the hardware too) and that it honestly and sincerely believes that this is the most important issue in all creation.

If you think this scenario is conceivable, you just might be Jerry Fodor. But we hope you don't find it conceivable. If it is inconceivable, then Fodor's story has to

be wrong. We will not try to prove that the scenario is inconceivable. It is sufficient if we can get you to believe that EGG doesn't have a mind like ours. We think that is obvious. EGG might satisfy the minimum condition for having a concept (for all we can prove), but he does not have a concept of the kind we possess. Our concepts are action-oriented.

We think two aspects of Fodor's own philosophical program actually pull in this direction. First, as discussed already, Fodor thinks concepts represent by being informationally tied to the world. That requires that concepts get their content through representations that tell us how to recognize things. Any knowledge that tells us how things appear can potentially be used to drive action. In creatures like us, perceptual information carries motor affordances. When we look at an egg, we see how it can be grasped. Thus, creatures like us cannot have egg concepts, even on Fodor's view, without having information that would facilitate egg-directed behavior.

Second, Fodor believes in mental representations because he believes that intentional explanations of behavior work. If Jones wants an egg, and believes there is one in the refrigerator door, Jones will open the refrigerator and gingerly lift out an egg. Notice, that this behavior cannot be explained by mere sentences in a language of thought. There must be a link between the concept of an egg, and the ability to find one and lift one. There must be a link between refrigerator thoughts and the ability to open the door. For Fodor, all of the stuff that links thoughts to the knowledge that drives action is independent of the thoughts that we ascribe when explaining action. But that undermines Fodor's explanatory project. If thoughts explain action, then having thoughts should provide us with resources on the basis of which we can act.

Before concluding, we want to explore the link between thinking and acting with one more case study.

3.2 Voluntary Body Motions

Consider the act of willing a voluntary body motion. For concreteness, take the event of your willing your left index finger to move. This event involved, *inter alia*, some neural activity in motor and pre-motor cortex. We take it as plausible to suppose that somewhere in that neural activity there lurks a (context-variable but identifiable) pattern that might reasonably be said to be the most local bearer of the content 'left index finger wiggle'.¹ How does that pattern get to bear that content?

¹ Recent experiments go some way towards establishing this. The experiments studied the way signals from cerebral cortex control the motions of a monkey's limbs. An owl monkey had 96 wires implanted into its frontal cortex, feeding signals into a computer. As the monkey's brain sent signals to move the monkey's limbs, this 'neural wire-tap' was used to gather data about the correlations between patterns of neural signal and specific motions. The correlations were not simple, and turned out to involve ensembles of neurons in multiple cortical areas. But the patterns, though buried, were there in the signals. Once these mappings were known, the computer could then predict the intended movements directly from the neural activity. See Wessberg *et al.* (2000).

In an impressive series of well-publicized experiments a paralyzed stroke victim was enabled to control a computer screen cursor using signals from two neural implants. The signals go to an amplifier which relays them to the computer. The patient is then, with effort and practice, able to learn to use their own neural activity to control simple cursor movements. Notice that in order to get the cursor to move, the patient first needs to *experiment with her own motor signals*. The implants, recall, are lodged in the motor cortex—the area controlling bodily movement. The nerves that feed into the electrodes thus carry signals which, normally, would participate in the control of movements such as the raising of an arm, or a leg, or the wiggling of a finger. To successfully move the cursor by thought, the paralyzed patient first tries to will the motion of various bodily parts. When such efforts yield a signal that the computer hears, a buzzer sounds, so the patient knows to concentrate on ‘that kind of thought’. After a while, having ‘that kind of thought’ becomes ‘second nature’²—when you want the cursor to move, you just will it to do so, exactly as you might will your own leg to move.

After sufficient practice, then, the kind of ‘thought control’ exerted by the paralyzed patient over the cursor, and by the normal subject over her biological arm look to be pretty much on a par. But suppose that someone (a Fodor of the motor cortex perhaps) insists there is a significant difference. Surely, they might say, the neural activity that normally causes you to raise a leg actually *means* (in virtue, perhaps, of that other, non-epistemic, content-fixing stuff) ‘raise that leg’, whereas the neural activity that drives the cursor really *means* let’s say, ‘wiggle your left finger’, even though, courtesy of the new circuitry, it now causes the cursor to move.

But why should we believe this? Imagine now that you are a young infant, above whose cot there dangles an attractive mobile. You desire to touch it. But you do not yet know how to issue the correct motor commands to do so. Your brain, however, generates many bursts of essentially random neural activity. Some of these bursts seem to move your hand closer to the target. After a while, you learn how to generate this kind of neural activity at will, and hence how to control your own limbs so as to carry out your project.

If this is a good picture of how we first learn to control our own bodies then surely there is a sense in which the neural activity that yields some desired result (a certain kind of arm motion, say) *counts as ‘meaning’ ‘move that arm’ only because, as we attempt to act on our world, it is the kind of neural activity which, as a matter of fact, brings about just that kind of result*. It is not that the neural signal, in any independent or prior sense, means ‘move that arm’. Rather, it means what it means because of what we reliably use it to bring about. If, later in life, we learn to use a ‘finger-wiggling’ pattern of neural activity to control the cursor, we are simply doing more of the same. The pattern of activity that promotes a specific cursor movement

² For this quote and further details, see Graham-Rowe (1998).

should later count as having the content ‘move the cursor to the right’ (or whatever) in just as rich and real a sense as it once counted as meaning ‘wiggle the left finger’. (In fact, it could come to mean both things simultaneously, were both chains of effect functional and intentionally exploited.)

The moral of this second story is, once again, that content cannot and should not be fixed independently of its expression in possible actions. The idea that the intention to wiggle one’s finger has nothing at all to do with the ability, circumstances permitting, to actually wiggle a finger is every bit as bizarre (and for much the same reasons) as the notion of the anti-egg-boiling inactivist.

Of course, the example of finger wiggling biases the case in favor of action-oriented theories of mental representation. Not all concepts are built into the motor system, and not all concepts are primarily in the business of motor control. The point of this example is that representations *can* get meaning from the uses to which they are put. Consider another case: you lose a chess piece, and you start using a bottle cap to serve as a pawn. In the game, the bottle cap is a pawn. If you didn’t use it that way, it would have another meaning. If we start thinking seriously about examples like this then theories that separate concepts from action will begin to look very odd. We think research on concepts would benefit from looking to such examples as a starting place. We have representations in order to act, and the way we act, on the basis of our representations, may have some impact on what they mean. Much work in philosophy of mind and language has focused on concepts of natural kinds, where the main point has been that we are able to represent a category despite how little we know about it. If we place action-guiding representations at the center, we may end up with different theories. We may discover in hindsight that natural kind concepts, which once seemed so pristine are action-oriented as well.

4. Conclusions

Fodor’s theory makes thinking prior to doing. It allows for an inactive agent or pure reflector, and for agents whose actions in various ways seem to float free of their own conceptual repertoires. Naturally evolved creatures are not like that. In the real world, thinking is always and everywhere about doing. The point of having a brain is to guide the actions of embodied beings in a complex material world. Some of those actions are, to be sure, more recondite than others. But in every case the contents of thoughts still look to depend, in some non-unique but vitally important way, on the kinds of doings they support.

Fodor wants to reject the tenth century and return to the seventeenth century. We think the time is ripe to move into the twentyfirst century, the century of situated and embodied cognition. Thinking, from the situated and embodied perspective, emerges as a much more familiar beast, and one that can actually do some explanatory work. What active agents think is constituted by what they can

do. It is the role in (potential) doing that makes our cognitive contents what they are, not the other way around. Just ask Granny.

University of North Carolina
Chapel Hill

University of Indiana
Bloomington

References

- Clark, A. 1993: *Associative Engines: Connectionism, Concepts and Representational Change*. Cambridge, MA: MIT Press.
- Fodor, J.A. this issue: Having concepts: a brief refutation of the twentieth century.
- Graham-Rowe, D. 1998: Think and it's done. *New Scientist*, 17/10/1998.
- Gass, W.H. 1976: *On Being Blue: A Philosophical Inquiry*. Boston, MA: David R. Godine.
- Prinz, J.J. 2002: *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press.
- Wessberg, J., Stambaugh, C., Kralic, J., Beck, P., Laubach, M., Chapin, J., Kim, J., Biggs, S., Mandayam A., and Nicolelis, M. 2000: Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408, 305–306.