Fodor on Imagistic Mental Representations
Daniel C. Burnston

Forthcoming in *Rivista Internazionale di Filosofia e Psicología*, Please Cite Final Version

Abstract

Fodor's view of the mind is *thoroughly computational*. This means that the basic kind of mental entity is a sententially structured or "discursive" mental representation. It also means that operations over this kind of mental representation have *broad architectural scope*, extending out to the edges of perception and the motor system. Still, however, in multiple epochs of his work, Fodor attempted to define a functional role for non-discursive, imagistic representation. I describe and critique his two considered proposals, both of which were initially formulated in *The Language of Thought*, but only the second of which is pursued thoroughly in later work. The first view says that images play a particular kind of functional role in certain types of deliberative tasks. The second says that images are solely restricted to the borders of perception, and act as a sort of medium for the fixing of conceptual reference. I argue, against the first proposal, that a broad-scope computationalism such as Fodor's renders images *in principle* functionally *redundant*, since all of their cognitively relevant properties must be represented in discursive form by both the producers and consumers of the representation. I argue, against the second proposal, that empirical evidence suggests that non-discursive representations are learned through perceptual learning, and directly inform category judgments. In each case, I point out extant debates for which the arguments are relevant. The upshot is that there is motivation for *limited scope* computationalism, in which some, but not all, mental processes operate on discursive mental representations.

1. Introduction

Fodor's influence on philosophy of mind is so widespread and so fundamental that it is often not directly acknowledged. I want to focus on just one issue in this essay – the role, if any, of non-propositional, imagistic representations within a generally (classical) computational approach to the mind. In Fodor's foundational work, the *Language of Thought* (Fodor, 1975), he considered, without clearly distinguishing them, two possible roles for such representations. First, they might be useful in certain kinds of problem solving. Second, they might provide a non-conceptual input for fixing the referents of concepts.

I think that there are major problems with both posited roles, but as often is the case, critiquing Fodor's stated views is formative. His admirable forthrightness in articulating his assumptions and commitments, and clarifying the scope of his positions, provides especially clear targets and contrasts for alternative views. I will argue, against the first proposal, that a computationalism as thorough-going as Fodor's leaves no distinctive role available for perceptual representation in reasoning. This explains why he largely abandoned the proposal in later work. The second proposal, however, runs into empirical problems, because there is evidence that non-discursive representations directly inform judgments and categorizations, which Fodor takes to be thoroughly conceptual processes.

The result of the analysis is that there is strong reason to accept *hybrid* accounts of cognitive architecture. Some mental processes operate on discursively structured propositional representations, some on perceptual and motor representation, and some on their interaction, but genuinely mental function is not restricted to any one form of representation.

I proceed as follows. In section 2, I articulate the central commitments of Fodor's computationalism. In section 3, I exegetically construct and criticize the first view, as Fodor lays it out in *The Language of Thought* (*LOT*). In section 4, I construct and critique the proposal, first raised in *LOT* and then furthered in *Langauge of Thought 2 (LOT2*, Fodor, 2008), that imagistic representations are only inputs to conceptual processing. In each of these sections I briefly point to current discussions for which the arguments are relevant. In section 5, I conclude by discussing the possibility of "hybrid" accounts of cognitive architecture, which genuinely embrace rich functional roles for non-propositional representation.

2. Core Commitments

In this section, I will articulate what I take to be the central commitments of Fodor's approach. For exegetical reasons, I will stick primarily to the formulations in *LOT*, with some moderate extensions. This is because I take these commitments to be ones that Fodor never relinquished. I will note parenthetically places elsewhere in his corpus where Fodor takes up similar themes.

2.1. Computationalism

Fodor's *computationalism* is based on the notion of *syntactically structured* mental representation. Computationalism is a view about both mental *entities* and *processes*. Mental entities are mental representations, and these are individuated according to their formal and semantic properties – a mental representation is a syntactically structured object that has a content. In this sense, mental representations are akin to linguistic statements; this is, of course, Fodor's guiding metaphor. To be syntactically structured is to be predicative, to attribute some property to an object. The content of a representation is its extension. Mental processes are inferences, which manipulate mental representations in virtue of their syntactic properties. A mental process transitions from given symbols to others according to a rule or function.

Comptuationalism underlies many other aspects of Fodor's overall picture, including his *naturalism*, his *anti-reductionism*, and his *propositional attitude realism*. I can only touch on these briefly here, but they are worth mentioning. According to Fodor, computation is the only kind of process we know of that is both naturalistic and genuinely inferential. Since inferences operate in virtue of the formal properties of representations, and since these properties ultimately are grounded on the physical states and transitions of the system, inferences are a proper part of physical causation. However, inferential processes are massively multiply realizable, guaranteeing the autonomy of psychological theory. Moreover, since propositional attitudes are themselves syntactically structured representations, and since folk psychology takes behavior to be the result of inferences over attitudes, a naturalistically kosher story can be told of how propositional attitudes cause behavior.

These joint commitments justify a kind of explanatory approach for Fodor, which takes propositional reasoning with folk psychological attitudes as the paradigm example of a mental process, and psycholinguistic data as the primary source of data about the language of thought.
Fodor does not think that every mental process operates on traditional folk psychological states like beliefs and desires. And he stresses continuously that the language of thought is not *just* a natural language, because these are neither coextensive nor ontogenetically similar. (That is, the language of thought is expressively richer than any natural language, and is innate whereas natural languages are learned.) But we are justified, according to Fodor, in starting from propositional reasoning and linguistic data and then expanding out from there.

These commitments are summed up nicely in the final chapter of *LOT*:

- "Mental states are relations between organisms and internal representations, and causally interrelated mental states succeed one another according to computational principles which apply formally to the representations. This is the sense in which internal representations provide the domains for such data processes as inform the mental life. It is, in short, of the essence of cognitive theories that they seek to interpret physical (causal) transformations as transformations of information, with the effect of exhibiting the rationality of mental processes." (*LOT*, p. 198)

2.2. The Broad Scope Commitment

I will call the second commitment the *broad scope* commitment. For Fodor, computationalism's scope is broad both phylogenetically, applying to any kind of minded organism, and architecturally, applying to all processes worth describing as mental. Indeed, they approach being *definitive* of the mental. It is having mental representations and performing rule-based inference which distinguishes the behavior of a minded thing from the behavior of, say, a planet on its orbit (Fodor, 1987). The architectural scope is primarily what we're concerned with here. To the extent that all mental processes are inferential, and to the extent that inference is best modeled as rule-based manipulation of symbols, "cognition is saturated with rationality through and through" (*LOT*, p. 172). Hence, computationalism applies both to more paradigmatically rational processes, such as decision-making, and to less paradigmatic ones like perception and motor control.

To take a couple of examples. For Fodor, decision is a (well,) decision-theoretic process. It involves explicit representation of one's beliefs about the world, including hypotheticals about what outcomes will result from what behaviors, as well as explicit representation of one's desires, including a preference ordering over outcomes. When one calculates what to do, one applies a decision-rule defined over these representations – i.e., one's decision is a "function" (*LOT*, p. 29) defined over these representations – and what one intends to do is to enact the behavioral option that is the output of the decision rule. In another important example, Fodor construes *concept learning* as "essentially a process of hypothesis formation and confirmation" (*LOT*, p. 35). When one learns a concept, one learns a rule that applies to the instances with which one is acquainted, and also applies to any other member of the concept's extension. One thus learns the concept by forming hypotheses based on instances and then confirming or disconfirming those hypotheses against new examples. (Fodor's worries about this notion of concept learning eventually led to his nativism; see below.)

It is vital for the discussion to follow that, according to Fodor's official view, both perception and motor control are just as computational as decision making and concept learning. Fodor's view of perception, which he expounds further in the *Modularity of Mind* (Fodor, 1983), is that perception is a dedicated device for taking transduced signal and converting that signal into the kinds of contents that are more stereotypical of beliefs. As Fodor says,

- "It is, I take it, an empirical question whether psychological processes are computational processes. But if they are, then what must go on in perception is that a description of the environment that is not couched in a vocabulary whose terms designate values of physical variables is somehow computed on the basis of a description that is couched in such a vocabulary." (*LOT*, p. 47)

So, perception's role is to take representations of the world that express only stimulus values – i.e., the physical quantities that transducers or purely "sensory" mechanisms are sensitive to – and convert them into the kinds of descriptions that can then be assented to in forming a belief. This conversion is computational because it implements rules for going from one kind of description to the other. As Fodor (1983) expands the view later on, perceptual inferences of this sort are distinguished from inferences more broadly by operating on a restricted set of encapsulated information and assumptions, which it uses to abduct from physical descriptions to conceptual ones. But these processes are still computational nature.[1]

A similar story goes for motor control. Motor control plays a very minor role in *LOT*, but Fodor addresses it explicitly in his paper, "The Appeal to Tacit Knowledge in Psychological Explanation," (Fodor, 1968) where he constructs, with characteristic humor, the following architectural view:

- "There is a little man who lives in one's head. The little man keeps a library. When one acts upon the intention to tie one's shoes, the little man fetches down a volume entitled Tying One's Shoes. The volume says such things as: 'Take the left free end of the shoelace in the left hand. Cross the left free end of the shoelace over the right free end of the shoelace . . . , etc.' When the little man reads the instruction 'take the left free end of the shoelace in the left hand', he pushes a button on a control panel. The button is marked 'take the left free end of a shoelace in the left hand'. When depressed, it activates a series of wheels, cogs, levers, and hydraulic mechanisms. As a causal consequence of the functioning of these mechanisms, one's left hand comes to seize the appropriate end of the shoelace. Similarly, mutatis mutandis, for the rest of the instructions." (Fodor, 1968, p. 627)

Fodor of course admits that this view is caricatured, but he thinks its general character can very well be correct for all that. The question is only how to flesh out the empirical details. The important notion is that there is a hierarchy of computational processes, which eventually bottom out in a non-mental process of firing the muscles (i.e., a purely mechanistic process that is akin to the operation of wheels, levers, and hydraulics). Importantly, this kind of architecture generalizes for Fodor. In general, one can expect that complex processes will involve a higher-level "executive":

- "One imagines a hierarchy of 'executive' programs which function to analyze macrotasks into microtasks. Such programs may 'call' both one another and lower-level problem-solving routines, though the extent of such cross-referencing is limited by the ingenuity of the program and, of course, the overall computational capacity of

[1] Following Fodor and much of the philosophical literature, I will primarily focus on *visual* representation in this paper. This is indeed an unfortunately oversimplification, given the importance of multimodality in perception (O'Callaghan, 2012), and the interesting question of whether and to what extent my points here apply to, say, the olfactory system (Barwich, 2014). For reasons of space, I will leave these further issues aside.

> the machine. When things go well the results of lower-level processes can be integrated to yield a solution of whatever macroproblem the system was originally posed. Whether, in a given case, things do go well is partly determined by whether the executive programs manage to select the right subroutines and to apply them in the right order." (*LOT*, p 164).

So, to sum up: what it means for cognition to be "rational through and through" is that computationalism holds across the entire architecture of the mind. In both perceptual and motor representation, the mental process goes all the way out, to the point where cognitive explanation stops and the mind interacts in a brute, non-mental causal fashion with the world – through perceptual transducers in the one case and through interaction with the body in the second case. Importantly, as will be explored at length below, this means that the language of thought – again, qua syntactically structured mental state – is computationally *sufficient* to account, in principle, for any process in the mind. That is, it has an expressive power which can describe any way in which perception and motor systems interact with the world.

2.3. The Difference in Kind Commitment

The first two commitments go very naturally together. The last commitment, I hope, will strike you as not-so-obviously-in-keeping with the others. (I hope that because that's what the rest of the paper is about.) The final commitment is a firm distinction in kind between "discursive" representations, which are paradigmatic ones in the language of thought, and "non-discursive" ones, which more closely approximate pictures, and are more closely associated with perceptual experience and imagery. Fodor calls these "iconic" representations, which comprise visual images and are "very much like pictures" (*LOT*, p. 184). While Fodor admits that, due to the broad-scope commitment, the representational architecture he is working with is "conceivably monolithic" (*LOT*, p. 172), he suggests that this is not in fact the case. Rather, empirical and introspective evidence suggests that "imaging plays some interesting role in thought" (*LOT*, p. 173).

Fodor struggles to articulate that role, however, because images don't have the right syntactic or semantic properties to contribute to a computationalist conception of thought – hence, "thinking and imaging can't be the same thing" (*LOT*, p. 173). While images can refer by resemblance, they do not have the predicative structure that allows for semantic evaluation – they therefore can't be truth-apt. Since, according to Fodor, thoughts just are the kind of things that can be true or false, images cannot constitute thought. Fodor thinks that this kind of failure is sufficient to doom any kind of empiricism, a theme that he will reiterate frequently throughout his corpus.

On the other side, Fodor thinks that it is precisely the *lack* of resemblance between discursive representations and their referents which enable them to state determinate, truth-apt contents. If thinking about John involves imaging John, how does one have the thought that "John is tall"? The *image* of John may show a tall man, but it will also resemble all kinds of *other* properties of John, for instance whether or not he is also heavy-set, whether he is standing or lying down, etc. Hence, non-discursive, imagistic representation fails to instantiate the connected properties of syntactic structure, truth-aptness, and abstractness that define discursive representations. Fodor later notes in *LOT2* that this same kind of problem disallows images from negating or quantifying.

In keeping with the broad scope claim above, there is no kind of thought which occurs solely in an imagistic medium. Fodor is "in fact, strongly inclined to doubt the very *intelligibility* of the suggestion that there is a stage at which cognitive processes are carried out in a medium that is fundamentally nondiscursive" (*LOT*, p. 176). (Fodor is talking about developmental stages here, but his claim would apply equally to parts of cognitive architecture.) And indeed, Fodor says that any non-discursive representation needs to be used "under a description" (*LOT*, p. 188). But, again, that doesn't mean that there is no imagistic representation. Fodor cites psychological studies which he takes to suggest the presence of imagistic representation. For instance, when subjects are asked to trace an imagined figure and answer questions about it, there is an interference effect when they must give their responses by pointing to visual targets, rather than verbally. And these effects are modality specific – they reverse when subjects are answering questions about auditory presentations of words (see descriptions and citations in Fodor, 1975). These results suggest the *existence* of imagery and its functional usefulness. The question is how to characterize that functionality.

In the next two sections, I explore and critique the two proposals that Fodor developed to characterize the functional role of imagistic representation.

3. The Specific-Contribution Proposal

3.1. Exegesis.

The first functional construal suggests that non-discursive representations play a particularly important functional role for *certain kinds of tasks* precisely *in virtue of* their representational form. Fodor considers a variety of examples and possible task-types that might benefit. For instance, he suggests that constructing an image might allow for *perceptual matching.* So, when an auditory description of a letter (e.g., "capital P"), must be matched to a subsequent stimulus, "it is extremely natural to assume that what happens during the … interstimulus interval is that the subject constructs a letter image to fit the auditory description, and that it is that image which gets matched to the visual display"

(*LOT*, p. 185). This kind of use might be extended to *comparatives* generally. One might use imagistic representations the same way one might point to a picture of a man and say, "I'm looking for someone *dressed like this*" or "*taller than this*" (*LOT*, p. 183). Imagistic representations thus might be more efficient for certain kinds of processes, for instance those requiring parallel rather than serial search (*LOT*, p. 190).

Fodor does not offer a definitive characterization of what exactly imagistic representations offer, to what kinds of tasks, or why. What is very clear in Fodor, however, and as should be clear from the commitments above, is that non-discursive representations *being employed* to perform certain tasks is *due* to a computational process. Recall the hierarchy of executives from the previous section. Fodor suggests that it is "intelligent management" (*LOT*, p. 163) that organizes sub-processes in a goal directed way. As he states:

- "There would seem to be a variety of representations that a given input may receive, and which representation it does receive depends, inter alia, on the demands of the subject's task. Second, the subject's achievement in matching the exploitation of his representational capacities to the exigencies of the [task] situation is itself a form of intelligent behavior." (*LOT*, p. 171).

Of course, via the broad scope commitment, these executive processes are themselves discursive and computational. Put all this together, and you get a view on which executive processes, which are computational, *construct* visual images in order to perform certain tasks. And indeed this is what Fodor proposes:

- "The explanation thus implies (what common sense also suggests) that we have psychological faculties which can construct images which display the information that corresponding descriptions convey discursively; i.e., faculties which permit us to construct images from descriptions. The experiment demonstrates that having the information displayed as an image facilitates performance in certain kinds of tasks" (*LOT*, p. 191)

Moreover, the images *themselves*, per the previous section, are not capable of much. It is only the images along with the symbols that interpret them. This falls out of the computational commitment – images, recall, do not have a determinate content, which thought requires. So images can only contribute to thought if they can be interpreted according to discursive symbols, i.e., "under a description."

The important upshot is as follows: for whatever role the images are supposed to play, discursive representations *must be expressively adequate* to describe the relevant contents in those representations. This is true both of the generation and use of the image. Since computational processes are *generating* the images to meet certain task goals, they must be capable of describing exactly the needed properties of the image. And, since only discursive representations can specify the contents of the images in a way that they can be used for other processes (which, via the broad scope commitment, are computational), any content that can be of use in the image must be expressible by the discursive representations that describe them. These combined commitments, as I argue in the next section, create a deep problem in finding a functionally distinctive contribution for imagistic representations.

3.2. Assessing the Specific-Contribution Proposal

The problem I have been exegetically working towards is this: if the entire functionally relevant content of the image can be specified by the discursive system *constructing* the image, and if that same content is what is read off by the discursive system that uses the image, then the content of the image is *functionally redundant*. The image plays no role that couldn't be played by directly feeding the descriptive content into the interpreter. So, on the kind of hierarchically structured executive system Fodor posits, along with the broad scope commitment, there can't be a functionally distinct role for imagistic representation. They are, and must be, computationally inert.

To flesh out this problem, we can allude to a famous discussion from cognitive science: Larkin and Simon's (1987) wonderful paper, "Why a diagram is (sometimes) worth ten thousand words." Larkin and Simon primarily focus on reasoning with external representations – diagrams on physical paper that are viewed with the visual system – but they clearly think that something similar goes on in the mind when people construct and use visual images. Their discussion is particularly relevant for two reasons: they assume that images and sentential representations can be *informationally equivalent*, which is a corollary of Fodor's "expressive adequacy" commitment; and they assume that the functional advantage for reasoning with an image are primarily in terms of *efficiency*, which, as we saw above, is one way that Fodor describes the possible advantage.

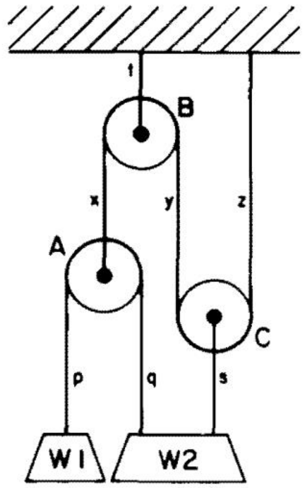

Figure 1. A pulley problem. From Larkin and Simon (1987).

Larkin and Simon consider the pulley problem given above, where the goal is to figure out the *ratio* of the two weights W1 and W2. They assume that this requires a way of representing the data, a program to reason about the data, and a database of information about the domain. They suggest that, even if performing the same set of inferences based on the same database, diagrammatic and sentential representations will differ in how they structure the data and in the kinds of steps their programs perform.

Larkin and Simon consider both types of architecture. On their framework, the two share the same database of knowledge, which explicitly states the relation between tensions of strings and weights, between single and multiple supports on a weight, and between pulleys and each of these properties. For instance, each starts from the assumption that the weight of W1 is one, and that weights and tensions are proportional. Therefore, the tension on string p is also one. From these assumptions, along with assumptions of how tensions distribute over and under pulleys, and how weight is determined with multiple strings of support (it is the sum the two tensions), one can work out that the ratio is 5:1. I am not going to go through the details for reasons of space; interested readers should consult the original paper. The important point is that the differences in data structure and program equate to a difference in *processing efficiency*. The data structure for the sentential program is simply a list of all of the relevant facts about the ropes, the pulleys and their arrangements. The data structure for the diagrammatic program, however, indexes the information to *locations*.

Both programs start at W1 with its weight of one, and the associated tension of one on rope p. And each proceeds step-by-step, determining the tensions and weights as they go according to the rules. The difference is that the diagram program has a basic and, Larkin and Simon assume, computationally cheap *shift of attention* operation. They also assume that all of the information at a location is immediately available to the diagram program. This leads to a massive change in the search costs for the two programs. Since the sentential program has no privileged location for search, it must search through every bit of information for what is relevant to each processing step, and that includes the new information it infers at each step. This is shown in the table on the left of Figure 2, where each 'o' represents non-needed information, and each 'x' needed information, at each processing step P1-P6.

Original Element	Steps / Productions	P1	1 P2	2 P3	3 P2	4 P2	5 P3	6 P4
2.	(Weight W1)	x	o	o	o			o
	(Rope Rp)	x	x	x	o			o
	(Rope Rq)	o	x	x	o			x
	(Pulley Pa)	o	x	x	o			o
	(hangs W1 from Rp)	x	o	o	o			o
	(pulley-system Rp Pa Rq)	o	x	x	o		o	o
	(Weight W2)	o		o	o	o	o	x
	(hangs W2 from Rq)	o		o	o	o	o	x
2.	(Rope Rx)	o		x	x	o	o	o
	(Pulley Pb)	o		o	x	o	o	o
	(Rope Ry)	o		o	x	x	x	o
	(Pulley Pc)	o		o	o	x	x	o
	(Rope Rz)	o		o	o	x	x	o
	(Rope Rt)	o		o	o	o	o	o
	(Rope Rs)	o		o	o	o	x	x
	(Ceiling c)	o		o	o	o	o	o
	(hangs Pa from Rx)	o		x	o	o	o	o
	(pulley-system Rx Pb Ry)	o			x	o	o	o
	(pulley-system Ry Pc Rz)	o				x	x	o
	(hangs Pb from Rt)	o					o	o
	(hangs Rt from c)	o					o	o
3.	(hangs Rz from c)	o					o	o
	(hangs Rs from Pc)	o					x	o
	(hangs W2 from Rs)	o					o	x
	(value W1 1)	x					o	o
New Elements								
1.	(value Rp 1)		x	x				x
2.	(value Rq 1)			x				o
3.	(value Ry 2)				x			o
4.	(value Rx 2)					x	x	o
5.	(value Rz 2)						x	o
6.	(value Rs 4)							x
7.	(value W2 5)							o
Total Elements Searched: 138		25	7	20	19	14	22	31

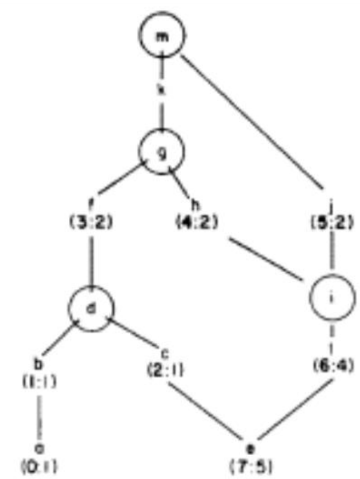

Figure 2. Search processes in sentential versus diagrammatic programs. From Larkin and Simon (1987).

The diagram program, however, only has to move its attention as it computes each step, and all of the needed information is immediately available. This search pattern is shown on the right of Figure 2 (the number on the left at each node is the step at which that node's value is calculated; the number on the right is the value obtained).

So far, there is nothing out-of-keeping with broad-scope computationalism. The differences in efficiency are determined only by differences in the number of sentential representations that must be held in mind and searched through. The diagram provides a distinctive *input* to the program, and the program can employ basic operations (e.g., shift-in-attention) to realize the gain in efficiency. And the gain is significant – from 138 search steps down to seven.

The problem arises when we think about the nature of executive control for Fodor, and the fact that imagery cannot operate strictly on perceptual input. Problem solving, recall, is hierarchical for Fodor, with commands passing between levels of executive control. Higher-level control processes organize types of representations in the best way to solve given tasks, and these representations include images. Imagery is the result of a process that generates the image from a description – in imagery, as opposed to in perception, the content that is imaged must be generated internally by the system. In this scenario, I submit, there is no way for the image to play a non-redundant role in the problem-solving process.

In going through this, it will be helpful to have an idealized architecture in mind. Call the higher-level executive just "*EXECUTIVE*" for short. Call the image-generator the "*DRAWER*" program. And call the program that reasons about the problem the "*SOLVER*." By the broad-scope commitment, and its corollary that the discursive system is expressively adequate for describing anything contained in the image, we can have two versions of *SOLVER*, which correspond to the two programs given in Larkin and Simon. Given the informational equivalence assumed by both the broad-scope commitment, and by Larkin and Simon, we must assume that the overall system has the resources to solve the problem either by generating an image or via the purely sentential method. Let's call the purely sentential program *SOLVER(SEN)* and the diagrammatic one *SOLVER(IMG)*. In this setup, we can see that there are a variety of reasons why the image can play no functional role.

Consider what *EXECUTIVE* has to do. *EXECUTIVE* has to look at the problem and *determine* the best way to solve it – i.e., it has to determine which representations and programs are the best ones for the job. The primary computational advantage attributed to the image is one of efficiency. *Executive* doing its job would consist in its recognizing that, for this kind of problem, the kind of program run with a diagram by *SOLVER(IMG)* is more efficient than the kind of program run by *SOLVER(SEN)*. Given this determination, *EXECUTIVE* instructs *DRAWER* to generate the image and runs *SOLVER(IMG)* on that input.

This sounds nice enough, until we realize that this purported gain in efficiency does not require any image at all. In order to make the determination, *EXECUTIVE* has to know exactly how many steps are involved in the search process, and have a rubric for preferring one over the other. But, as Fodor notes in *LOT*, the only way to measure complexity in a computational system is through *syntax* – the complexity of the representations and operations that the program will have to use to solve the problem. In order for *EXECUTIVE* to be able to make this kind of accounting, it must know what the data structures are and how many steps will be needed by each program. But, according to broad-scope, the executive system *is a sentential system*. *EXECUTIVE* itself *cannot recognize* non-discursive representation. So in order for it to judge efficiency, it would have to already know all of the steps that would be performed over what data structures, and would have to know this in discursive format. Now we can phrase the functional redundancy worry. What is the point of actually running *DRAWER*, given that it already has exactly the information that *SOLVER(IMG)* needs to represent discursively, in exactly that format? Recall that, in order to be used, the image must be converted into discursive representation, but given that the information is already available in that form to *EXECUTIVE*, there is no point to running *DRAWER* rather than directly feeding the information, in discursive form, to *SOLVER(IMG)*.

One natural response to this worry, which is unfortunately not available to Fodor, is to deny that *EXECUTIVE* has to have access to everything that *SOLVER(SEN)* and *SOLVER(IMG)* do in order to make its determination. Rather, it might operate on a kind of heuristic – e.g., "when solving pulley problems use *SOLVER(IMG)*". On this scheme, all *EXECUTIVE* would require, on the front-end, is an ability to recognize task-types. It wouldn't need to count representations and operations, and hence wouldn't need to have access to them in discursive format. It could simply call the right program for the task. Now, the heuristic solution is not open to Fodor because he thinks that central cognition is isotropic (Fodor, 1983). What it is for a process to be isotropic, vis-à-vis some domain, is to be open to information from anywhere in the system in solving problems in that domain. That is, there is no demarcation of the domain "pulley problems" that determines exactly the right set of information or program that should be used to solve that problem. Fodor argues very strenuously in *The Mind Doesn't Work That Way* (Fodor, 2001) that the isotropic nature of central cognition (for our purposes, *EXECUTIVE*) prevents the use of heuristics to solve the frame problem (i.e., determining which process, over which information, should be used to solve a given task).

I don't want to focus on this aspect of Fodor's work here, since I think that, even given the possibility of a heuristic program-selector, no non-redundant role for imagery is guaranteed. To see this, now consider what *DRAWER* has to do. *DRAWER*'s function, recall, is to generate images from descriptions. On the non-heuristic account, presumably the description comes from the executive. On the heuristic account, however, the executive does not have access to the information that *DRAWER* uses to generate the image. Let's assume that *DRAWER*, itself, can generate the description it needs in order to produce the image. So, when being called for *pulley-problem*, it produces an image of a pulley system under the description given by the empirical data.

Here's the issue: *DRAWER* can't produce just *any* diagram in relation to the description. The diagram has to be the kind of one that *SOLVER(IMG)* can use to solve the problem. How are we to guarantee this? The only way to do so is if *DRAWER* either represents, or is wired such as to only produce, exactly the kinds of images that *SOLVER(IMG)* can read. But, again, the image is generated *from a description,* and therefore these properties *must* be represented discursively in *DRAWER*, or *DRAWER* must be somehow hard-wired to only produce just the discursive representations (since these generate the images) that *SOLVER(IMG)* reads. For instance, in the diagrammatic solution from Larkin and Simon, there is a basic operation of attention moving between *locations*. And given that *DRAWER* has to produce images in the relevant diagrammatic space, it must be able to represent *locations* in sentential format. Since *location* is just an index for information which allows for ordered search, and since that information must be represented discursively on both the front and back end, *DRAWER* already can express all of the information that *SOLVER(IMG) uses discursively* in discursive format. Similarly for the other properties that *DRAWER* supposedly depicts in the diagrammatic space.

But if *DRAWER* can already represent all of the properties that *SOLVER(IMG)* uses to solve the problem in discursive format, then there is no point to generating the image. After all *SOLVER(IMG)* is itself a computational process, one that has to operate over sentential representation. The only difference is that the diagram makes certain inputs "immediately available" to it. What informational equivalence, plus the image-from-description nature of *DRAWER* entails, however, is that all of the relevant information needed by the consumer process (*SOLVER(IMG)*) is represented in discursive format by the front-end one (*DRAWER*). If this is the case, then no computational advantage is gained by having the image. *DRAWER* might as well skip drawing and feed the sentential representations directly as input to *SOLVER(IMG)*.

We can briefly consider some potential responses. One is to question informational equivalence. On this response, there is *more* to the image than what is contained in the description that is used to generate the image. Perhaps there is information that is *implicit* in the description that the full image conveys *explicitly*, and perhaps this is what a program like *SOLVER(IMG)* uses to run its program. The problem with this proposal is that it is left mysterious why the particular information that is implicit in the description, but explicit in the diagram, should be of any functional use. As Fodor stresses over and over about icons, their content is *indeterminate* if not used *under a description*. It is only discursive description of their properties that makes icons usable for reasoning. What *DRAWER* has to do, in order to make something that *SOLVER(IMG)* can use, is to make an image that will precisely match the description *under which SOLVER(IMG) uses the image*. Otherwise, there is no guarantee that the implicit information will in fact be useful. Fodor in fact admits this, saying that, "The image that gets produced may be quite schematic since how the image is taken—what role it plays in cognitive processing—is determined not only by its figural properties but also by the character of the description it is paired with" (*LOT*, p. 192).

Two quick other responses are worth considering. First it might be argued that the diagrammatic representation serves a memory function; it might be easier to remember information in a diagrammatic format. Second, it could be argued that the image plays a *communicative* role that is not playable by discursive representations. Perhaps, for instance, *SOLVER(IMG)* can only take images as inputs, and cannot receive direct sentential input from either *EXECUTIVE* or *DRAWER*. As such, an imagistic intermediary between *EXECUTIVE* (or *DRAWER*, on the heuristic solution) is required to generate the benefits of efficiency.

The thing to note about these proposals is that each relies, not on a general view about the architecture of cognition, but on quasi-empirical claims about what the memory limits of architectures for particular data structures might be, or on what kinds of programs can share what information. It is hard to know how to even assess this kind of claim without discussion of particular systems, and it is thus out-of-keeping with Fodor's general kind of argument – namely, that the project of naturalizing the mind, along with the nature of computation, *require* that cognition operate in a way that meets the broad-scope commitment. The broad-scope and distinct kind commitments are, thus, incompatible with images serving any non-redundant functional role in thought.

As noted, Fodor almost completely abandons the specific-contribution view after *LOT*. I can only find one place in *LOT2* where he references this possibility (p. 22). I suggest that this was for good reason – it is not really compatible with broad-scope computationalism. In the next section, we will consider the possibility that he pursues further. First, however, I want to point out that the specific-contribution view is still relevant to current discussions.

3.3. Relevance to current discussions.

To be perfectly clear: I am *not* attempting anything close to exegetical thoroughness here (or in section 4.3 below). The point in this section is just that the considerations raised by close analysis of Fodor's proposals are not restricted to Fodor's view.

Consider, first, Carruthers' (2015) recent account of cognitive architecture in his book, *The Centered Mind*. Carruthers' considerations, at first, seem very far removed from Fodor's. He is interested primarily in what the science of working memory can tell us about the stream of consciousness, and Fodor is persistently quiet on the nature of consciousness. However, the architecture Carruthers proposes is Fodorian in many respects. While Carruthers thinks that all conscious experience is sensory in character, he posits that the activation and control of these sensory images is due to amodal representations, which are "active in the background of the stream of consciousness, causing and controlling the latter's contents" (p. 18). Moreover, they "determine" those contents (p. 7). Once sensory representations are thus activated, they are selected for global broadcast, hence becoming conscious. However, the users of these representations elsewhere in the mind are also amodal, and amodal contents can be "bound into" sensory representations and broadcast along with them. Lastly, Carruthers' pictures of decision and categorization are similar to Fodor's, in that each are computational processes involving amodal concepts. So, actions are the direct result of an entirely amodal decision process, and seeing an image of a Dalmatian in a page of blotches involves applying the concept 'Dalmatian' to one's percept.

Again, this is not even a half-hearted attempt at a full exegesis. But I want to show how someone could be motivated by the present concerns to pose questions to an account like Carruthers'. For one, it is not clear that the notion of an amodal concept being "bound into" a sensory one is compatible with the distinct-kinds commitment, because the nature of this binding is obscure. It makes sense for how an amodal concept can fit into a sentential structure – i.e., by playing a syntactic role within it. But what 'slot' in an iconic representation does an amodal concept fill, and how does it get there (cf. Green & Quilty-Dunn, 2017)? Secondly, the worry about functional redundancy is particularly pertinent here. Carruthers suggests that images play a particularly useful role in prospective decision-making, for envisioning the possible consequences of one's actions. But if these contents are determined by amodal representations, and must be conceptualized as such, then the functional redundancy worry looms.

Now, there are other aspects of Carruthers' account which he could draw on, which are not available to Fodor. For instance, he thinks that the modal-amodal interactions in the mind are a result of evolutionary conservation of elements of cognitive architecture. This is the kind of empirical claim that I suggested above might be needed to get out of the redundancy issue. However, it is famously not open to Fodor due to his skepticism regarding evolutionary theorizing about the mind (Fodor, 2001), and I too am skeptical that evolutionary considerations can confirm hypotheses about cognitive architecture (Burnston, 2016). But this is where the conversation would have to go.

As another example, consider current empirically informed philosophy of action. Recently, several theorists have begun to recognize the importance of motor imagery in action, and this has led to the question of how these kinds of representations interact with traditionally defined, discursive propositional attitudes. There is no consensus on how to solve this "interface problem" (Butterfill & Sinigaglia, 2014; Ferretti & Caiani, 2018; Mylopoulos & Pacherie, 2017, 2018; Shepherd, 2017). My own diagnosis of the state of this discussion (Burnston, 2017a) is that it is in somewhat of a quandary, with philosophers attributing all of the genuine *decision* processes to propositional attitudes, and then trying to figure out how those decisions get embedded into perceptual and motor representations. But I think this is susceptible to redundancy worries.

To bring this out, consider a kind of case inspired by the discussion in Butterfill and Sinigaglia (2014). Imagine a mountain climber who begins to fall, and can grab either at a sturdy ledge or a bit of loose gravel. Butterfill and Sinigaglia suggest *both* that motor imagery is able to select the ledge as the goal for a grasping action, *and* that this process has to be coordinated with a purely discursive process which is also capable of representing that goal. As they put it, the outcome is "multiply determined." It seems to me, at least, that I would rather not have to coordinate multiple cognitive faculties in a sophisticated way if (i) that coordination were redundant, and (ii) I'm at risk of falling off a mountain. Of course, this is an intuition pump and not an argument, but the functional redundancy worry informs it.[2] I'll say more about this in the final section. For now, just note that this is *not* a problem if you posit that no genuinely mental stuff is done by motor imagery – if, following Fodor, you posit that propositional representation goes all the way out the edges of the motor system (Brozzo, 2017).

4. Images as a Reference-Fixing Medium

4.1. Exegesis.

The second kind of function is explored peremptorily in *LOT*, but comes to be the major view Fodor puts forward in *LOT2*, and its influential companion paper "The Revenge of the Given" (Fodor, 2007). Put briefly, the view is that non-discursive representations are useful as a way of *fixing the reference of concepts*. The idea is that non-conceptual representations serve as a kind of medium in which the fixing of conceptual reference can occur. This process is not computational, but serves as a starting point for computational thought. This view is bound up in complicated ways with Fodor's nativism and conceptual atomism, and it is worth exploring these ways briefly.

In *LOT*, the idea that the language of thought is innate is primarily motivated by considerations about learning natural language, and learning concepts on the basis of perception. For Fodor, learning a concept (/word) involves learning a rule under which it applies. Learning a rule is a process of *hypothesis and confirmation*. But this is problematic: in order to *formulate* the hypothesis about the concept, one must be able to represent both the concept and the rule which is supposed to determine its extension. Hence, concepts have to be innate.

In *LOT*, Fodor was concerned with "ameliorating" the counter-intuitiveness of nativism, and attempted to do so in two ways. First, he wanted to show that you could make sense of how innate concepts are initially activated in response to the environment (so that they needn't be "literally present" at birth), and second he wanted to suggest that maybe not *all* concepts needed to be innate. He pursued the first point with the notion of an "exemplar," and the second with a distinction between simple and complex concepts. Exemplars are supposed to be perceptual and themselves non-conceptual, and thus don't have to be learned. The idea is that when one is in the presence of a certain kind of environmental input, one's innate concept is activated, and this process is itself a non-inferential one. Since

[2] To be sure, there is considerably more complexity to Butterfill and Sinigaglia's view. They take propositional states like intentions to *defer* to motor plans that are generated in motor imagery. But again, it is unclear why this is not redundant. If motor imagery can specify and select outcomes, why is the coordination needed? See (Burnston, 2017a) for further discussion.

it is not inferential, it is not learned (learning, recall, is one of the mental processes covered by the broad-scope commitment). Fodor reaches for the notion of "imprinting" as an analogous process, and he repeats this in *LOT2*. In *LOT*, the notion is that you have some innate, simple concepts that are triggered by exemplars, and then you can generate new concepts via conceptual combination. So, maybe you get "flying" and "machine" by imprinting, and then you can get "airplane" by defining it as "flying machine."

The notion of imprinting, while serving an ameliorative role in *LOT*, became intensely important for Fodor as his nativism became more pronounced. And his nativism became more pronounced as his atomism became more pronounced. Fodor subsequently developed doubt about there being any definitions for concepts, and hence about the notion that you can have a robust simple/complex distinction for concepts. The worry is that any definitional notion of concepts implies a disastrous semantic holism (Fodor, 1987; Fodor, 1998; Fodor & Lepore, 1992). I don't have space to go into these developments in detail, and doing so would take us relatively far afield. The point is that, once you take concepts to be atomistic and unlearned, the problem of how they connect to the environment is exacerbated. Fodor spends a fair bit of time dealing with this problem in *LOT2*, and his notion of non-conceptual representation was largely shaped by it.

In *LOT2*, Fodor, befitting his now rabid "mad dog" nativism, suggests that the very idea of learning a concept – any concept, now – is incoherent. But concepts still must be connected to the world (in other circles, this is known as the "symbol grounding" problem; see Barsalou, 1999). Here is where imprinting comes to the fore:

- "There would appear to be plenty of ethological precedents—from 'imprinting' to 'parameter setting' inclusive—where it's implausible that the acquisition of a concept is mediated by a rational process like inductive inference, but where concept acquisition is nevertheless highly sensitive to the character of the creature's experience. So, neither the *LOT 1* argument nor the present revision shows that concepts can't be acquired from experience. The most they show is that acquiring a concept from experience must be distinguished from learning it." (*LOT2*, p. 144)

The way this works, according to *LOT2*, is that we have a non-inferential ability to recognize stereotypical referents, and once we do this it triggers (via a non-mental, neural process) the activation of the concept. Here is where non-conceptual representation fits in. The stereotypes, since they are metaphysically distinct and ontogenetically prior to concept activation, cannot themselves be concepts. Hence, there must be a non-discursive medium in which they are represented. And, since iconic representations are already part of Fodor's mental ontology, they naturally slide in here. The considered view in *LOT2* is that we have an "attractor landscape" of recognitional abilities, defined non-conceptually, that each attractor corresponds to a stereotype, and that falling into the attractor triggers the concept.

This is, in effect, a *dispositional* rather than an inferential account of concept formation, where "*what's innate is the geometry of the attractor landscape*" (*LOT2*, p. 161).

Fodor is never very clear what this kind of geometry consists in, only suggesting that it allows us to pick up on "statistical" regularities in the environment. The idea that the geometry is predominantly innate thus minimizes the role of learning in imprinting on objects. Fodor will occasionally refer to the kind of fixing on exemplars he discusses as "statistical-inductive," where this is only characterized negatively – i.e., as *not* the kind of hypothesis formation and confirmation that he thinks must characterize *concept* learning. In both *LOT* and *LOT2* Fodor gestures towards a kind of learning that might take place within the perceptual system itself, but he clearly thinks this is limited, and bound to particulars (i.e., not concepts). He suggests in *LOT* that perceptual learning can allow us to *discriminate* more finely between objects, and that it can allows us to recognize certain sensory properties associated with objects – e.g., "what a steak tastes like, learning what middle C sounds like played on an oboe, and so forth" (*LOT*, p. 34). In *LOT2*, Fodor says that all his nativism requires is that the "initial layout" of the attractor landscape is set, where is possible that the geometry "can alter under the pressure of experience, learning, maturation, or any other mind-world or mind-body interactions" (*LOT2*, p. 163).

What Fodor is clear about, however, is that there are limits to what non-conceptual representations can represent. While they might *trigger* the application of a concept, non-conceptual representations themselves cannot *individuate* an object as falling under a category. That is to say, while imprinting is itself a process of going from non-conceptual contents to the tokening of concepts, non-conceptual representations don't *themselves* convey the content of the concept. Fodor is thus committed to a "shallow," "thin," or "conservative" view of what non-conceptual contents can represent. As he says:

- "In the case of vision, the icons register the sorts of properties that photographs do (two-dimensional shape, shading, color, and so forth) but not 'object' properties like being an animal or, a fortiori, being a cat belonging to Granny. You can, of course, see a cat as a cat that belongs to Granny; but that requires conceptualization. The present point is that a cat can't be, as it were, given as a cat that belongs to Granny" (*LOT2*, pp. 185-186).

So, for Fodor non-conceptual representations are bound to the particulars that the visual system has access to. They are the outcomes of the mechanisms of transduction. While these may be stored in a "buffer" so that imprinting can work over them, they themselves do not contain any content aside from the physical properties that can be transduced.

Fodor's positive view of iconic representation is the "picture principle," according to which every decomposition of an iconic representation represents a part of what the icon represents. So, if you carve a bit out of a picture of a dog, then that bit still represents (say) the leg of the dog, whereas if you pull out the 'o' from 'dog', it doesn't represent some part of the dog. This principle captures the general properties we've been discussing – the lack of determinate content, the lack of predicative structure, etc., because there is no "privileged" semantic decomposition of an icon, the way there is for discursive representations.

Two last exegetical points are worth noting before we move on to consider this function for non-discursive representation. The first is to note how fundamentally different this proposed function is, compared to the one in the last section, even though they are not clearly distinguished in *LOT*. The specific-contribution view suggests that non-discursive representations play a functionally distinct role for contributing to specific tasks involving deliberation and judgment. The concept-trigger view restricts non-conceptual representation to the very *borders* of the mind, in fact yet another significant step (the statistical-inductive one) from where concepts are applied. *Nothing* task-specific or deliberative is done in this medium. So, the second point to note is really just how much better this fits with the other commitments. I now move on to assess this position.

4.2. Assessing the Reference-Fixing View

The basic commitments of the reference-fixing view of non-conceptual representations are that (i) they are pre-computational, and (ii) they are "shallow" in what they can represent. They are, in effect, just the outputs of perceptual transducers, and a process of statistical-inductive learning allows for the "imprinting" of concepts onto particular patterns in these representations. All the work of forming perceptual judgments, updating beliefs, and *thinking* is still purely done in discursive format.

The basic problem with this picture is that Fodor underestimates the scope, the malleability, and the *impact on behavior* of processes that are non-conceptual and involve statistical-inductive learning. In this section I'll discuss empirical results from perceptual learning which suggest that, far from simply giving way to discursive representations, perceptual learning encodes non-conceptual structures that are actively employed in judgments. What marks these representations as non-discursive is that they are *holistic* and *dimensional*.

I'll discuss two examples, which bring out the holistic and dimensional nature of perceptual learning. The first is from a highly-cited paper by Schyns and Rodet (1997), who explicitly contrasted holistic models of perceptual recognition with models based on *logical construction*. The way they did this was to have subjects categorize items that, in terms of conjunctions of features, were equivalent, but vary the subjects' history of exposure to these combinations of features. Their basic stimuli and methods are shown below.

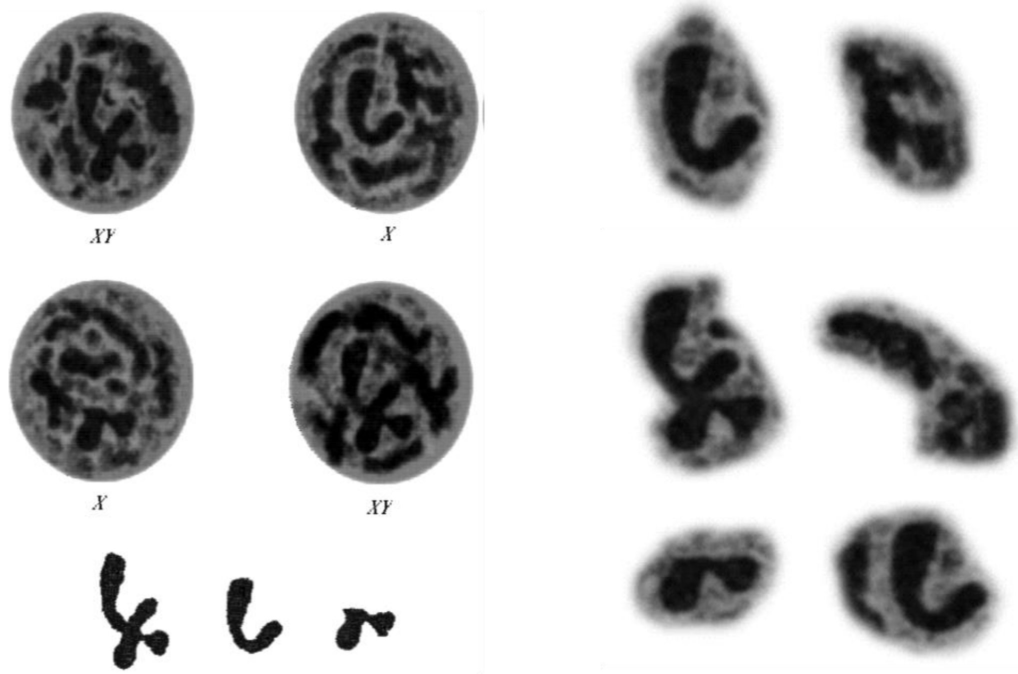

Figure 3. Stimuli and test conditions in Schyns and Rodet's (1997) study.

In the left panel are the training regimes and the basic features. The bottom shows the two basic features (X in the middle, Y to the right), and their conjunction (XY, to the left). The experimental groups were distinguished by their training regimes. One group followed the regime in the middle row, first learning to distinguish feature X, then learning to distinguish XY. The other group was first taught XY, then X (top row). Schyns and Rodet suggested that the group first trained on X would subsequently be able to recognize Y. That's because, when they have to learn XY, they have to *add* something (i.e., Y) specific to the feature X, which they've already learned. However, they suggested that the group who first learned XY would not subsequently recognize Y after learning X, because they encoded XY as a holistic unit.

They tested this via the experiment laid out in the right panel. The two parts of each row were presented to subjects successively, and subjects were told that they were two distinct parts of the same "cell," presented separately. The top row shows X paired with a distractor feature, the middle row shows XY paired with a distractor feature, and the bottom row shows the two parts of XY presented sequentially, which the experimenters referred to as X-Y. Subjects were asked what categories were presented in the combined snapshots. Both groups referred to the top row stimuli as X, and to the middle row stimuli as XY. However, they *differed* in the X-Y stimuli. While almost all of the group first trained with XY classified the bottom row as X, a significantly greater number of the group first trained on X classified the bottom row as XY.

What explains this difference? Schyns and Rodet suggest that the training history predisposed the XY→X group to encode XY *holistically* – that is, not as comprised of parts – whereas it predisposed the X→XY group to also distinguish Y. Thus, subjects in the X→XY group would be predisposed to see XY as a conjunction, and thus classify the X-Y stimuli as XY, whereas the XY→X subjects would perceive XY holistically, and thus classify X-Y only by the X feature. Amazingly, this effect persists even if XY→X subjects were *also* subsequently taught Y. It seems like, once XY was encoded holistically, it was just not subject to decomposition. To test whether this was just some weighting operation on distinguished feature, subjects were then shown example cells and simply asked to draw the features they'd seen before. When X→XY subjects saw a cell with a Y feature, such as in the second row of the right panel of figure XX, they drew a circle around the Y feature, but subjects from the XY→X group did not. It was as though they were "cognitively blind" to Y.

Why is all this important? What the experiment shows is that perceptual learning distinguishes categories that are *equivalent* in terms of their discursive descriptions. In both groups, XY can simply be defined as the conjunction X^Y. But if both groups of subjects in fact represented them that way, they would both classify the X-Y stimuli as XY. But they don't do that, and hence the discursive representation does not describe subjects' categories (at least for the XY→X group). This suggests that perceptual learning creates structured representations whose content is *distinct* from discursive descriptions – i.e., non-conceptual – and that those representations *shape perceptual judgments*. But, on the reference-fixing view, this is more function than non-conceptual representation should contribute. Equivalence of discursive representation should equate to equivalence of judgment.

I'll return to this shortly. First, I want to discuss one more example which shows that percepts are organized *dimensionally*. Consider the following figure from Gureckis and Goldstone (2008).

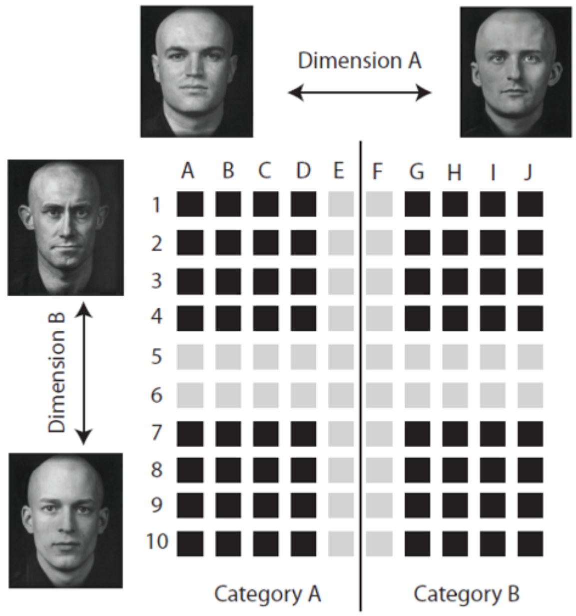

Figure 4. Exemplar space from Gureckis and Goldstone (2008).

The figure shows an *exemplar space* of faces. Each square represents an individual face. Each face is constructed by combining, in a continuous fashion, the features on either end of dimension A, with those on either end of dimension B. In a wide range of studies using faces, as well as objects like cars (Folstein, Gauthier, & Palmeri, 2012), it has been shown that experimenters can define an arbitrary category boundary in this space (such as the vertical line in the center of the space above), and subjects can learn to discriminate the arbitrary categories based on *generalization from exemplars*. Subjects are presented an example, guess which category it belongs to, and are given feedback. After training, subjects can categorize a *novel* example as belonging to the appropriate category. Moreover, if subjects are trained on the A versus B discrimination, they also, with no feedback at all, begin to distinguish along the orthogonal dimension as well (thus splitting the space into four quadrants as shown).

This kind of representational space is called a *morphspace* (Folstein et al., 2012) because not only do subjects discriminate along the relevant dimensions, but the dimensions are actually warped to accentuate differences between categories. This is shown by tracking subjects' *similarity judgments* before and after the training. A number of different types of similarity judgment have been used (Goldstone & Steyvers, 2001; Gureckis & Goldstone, 2008; Folstein et al., 2012), but their general result is that intra-category members are classified as *more alike* after training than before training.

This causes problems for Fodor's view because it is arguable that these representations are non-discursive, but they underlie perceptual category judgments. This is out-of-bounds for the reference-fixing view. Fodor calls it a *truism* that recognition requires discursive representation. As he says elsewhere: "If the contents of perceptual judgments weren't conceptualized, they wouldn't be judgments. A judgment that a is F ipso facto conceptualizes a as F" (Fodor, 2003, p. 54). These dimensional representations are non-discursive in virtue of their continuity, and in this sense they meet the picture principle – every part of the dimension represents a part of the morphspace. They also exhibit the lack of determinate content that Fodor takes to be characteristic of images.

However, it is pretty clear that these representations do not have many of the other properties of pictures. They do not represent physical space, for one thing. For another, they are not exhausted by particulars in the way that Fodor's view suggests. That is, they represent *patterns of variation* amongst particulars, rather than the particulars themselves. Lastly, it should be noted that representations of the kind discussed in this section are taken to be widely important– they have been taken to underlie expertise in X-ray reading (Sowden, Davies, & Roling, 2000) and facial recognition (Gauthier & Tarr, 2002), among other things.

To summarize: non-conceptual representation, under the concept-trigger view, is pushed out to the very edges of the mind, such that the kind of inductive process that triggers concepts doesn't *really* challenge the broad scope commitment. Non-conceptual representation is a just a buffer where the linkage between sensory input and conceptualization occurs. The rest of the picture can remain just as before. If non-discursive representations are used in *any* type of active judgment, then this questions the move to quarantine imagistic representations to the very boundaries of the mind. As I will suggest in section 5, this opens up a whole range of options for the possible functionality of these representations.

4.3. Relevance to Current Debates

The picture principle, and the issues that surround it, have become central for a variety of issues in the philosophy of perception. The two most pertinent also happen to be intimately related – one is the *perception-cognition distinction* and the other about *higher-level content*. Both aspects of Fodor's view have been influential, including his definition of format and his notion of shallow contents.

As with the case of motor imagery discussed in section 3.3, in has struck some theorists that the best way of drawing a distinction between perception and cognition is in terms of representational format (Burnston, 2017b), and Fodor's characterization has proven foundational for this debate (Quilty-Dunn, 2016). The form distinction is one candidate out of an array for characterizing perception, but it has some significant advantages. For one, it seems to capture the ways in which perception and imagery are continuous. For another, it captures the intuitive picture-like phenomenology of at least some of perception. As we've seen, both of these were present in Fodor's characterization all the way back in *LOT*. Hence, the picture principle has taken center stage in debates about whether perception is distinguished from cognition in terms of its representational format (Green & Quilty-Dunn, 2017).

However, given that, for Fodor, the picture principle was only supposed to apply to the very borders of perceptual processing, it is somewhat odd – and certainly not mandatory – to think that the form distinction is committed to the picture principle being a full characterization of perception. As I noted in the previous section, while the holistic and dimensional representations produced by perceptual learning do conform to the picture principle in some respects, they do not do so in all respects. Whether these kinds of representations must be read in terms of the picture principle, and what this means for the debates surrounding the form distinction, are open questions. Again, this is not meant to overcome anyone's arguments in the debate, but to show how dialectical options have been structured by Fodor's views, in a way that is not fully exhaustive (cf. Green, forthcoming).

This brings us directly to the discussion of higher-level content. The picture principle is rarely explicitly cited in this debate, but it shapes it in a rather deep way. Many different theorists have, at this point, posited higher-level contents as a way for perception to immediately recognize a wide-array of properties, including kind-properties (Siegel, 2006), the morally right action in a situation (Goldie, 2007), and natural scene properties (Bayne, 2016). However, almost *universally*, these contents are described in discursive terms – one recognizes pine trees, for instance, when one *sees* them *as pine trees*. While sometimes higher-level content theorists will consider the possibility that higher-level contents are non-conceptual, they almost never take the possibility as worth much time to analyze (see, e.g., Silins, 2013). I think this comes from an implicit assumption that non-conceptual contents are most likely limited to the boundaries of perception, again in Fodorian spirit.

Of course, this set of views, if they are really committed to discursive representation of higher-level contents, is strongly committed to *not* construing the perception-cognition border in terms of the form distinction. Again, this set of views is not really discussed in the debate, much to its detriment. Without some distinction in mind, there is no way to tell whether a categorical representation falls on one side or other of the perception/cognition border (Burnston, 2017c; Burnston & Cohen, 2015).[3]

The considerations offered in this section offer another option. I argued above that there are non-conceptual representations underlying perceptual judgments. I've also argued elsewhere that perceptual learning is flexible enough to be decidedly non-modular

[3] One notable exception to this is Mandelbaum's (2017) recent work, which explicitly invokes a Fodorian modularity criterion to argue that higher-level contents are both discursive and perceptual.

(Burnston & Cohen, 2015). But, if the considerations I've invoked here are correct, it may be possible that a distinct combination of non-conceptual, higher-level, *and* perceptual contents exist, which are produced by a non-modular process (Burnston, in submission). This is not a position currently considered in the debate.

5. Conclusion

As should be clear by now, I think that a rich role for imagistic representation in cognition just is incompatible with broad-scope computationalism. And I think that there is at least suggestive evidence that non-conceptual representation does play a rather rich role in the mind. If one thinks both of these things, then one has strong reason to doubt whether broad-scope computationalism is true. I want to conclude by just laying out a bit of the landscape for someone who holds this set of opinions.

As it turns out, this kind of discussion is already at work in the debate surrounding embodied cognition. More modest forms of embodied cognition, occasionally referred to as "grounded" cognition, posit exactly that sensorimotor representations play a role in the kinds of functions that are traditionally the domain of cognition – including reasoning, semantics, action planning, etc. Some have even posited that mathematical skill in part relies on sensorimotor knowledge of how to manipulate symbols (Landy, Allen, & Zednik, 2014).

The staunchest versions of these views says that the *entirety* of these mental functions is done in sensorimotor representation. I think this is too strong; my preferred position is a *hybrid* one, on which there are both sententially structured representations and imagistic ones, and that mentation relies on both (cf. Gauker, 2011). This has been characterized by Machery (2007) as a kind of "narrow-scope neo-empiricism."

On this kind of position, Fodor may largely have been right about *part* of the mind – whichever part encodes rules for reasoning, operates purely on discursive representations etc. But if that is not *all* of the mind, then all kinds of fascinating dialectical possibilities open up. For instance, much of occurrent cognition may be due to the *interaction* of discursive and sensorimotor processes, without one of them being solely responsible for much. Or, it may be that there are distinct forms of reasoning that happen solely within one or the other.

I am partial to the first hypothesis. But even here, there is a huge amount of work to be done in considering further the nature of imagistic representation, what kinds of mental processes (if any) can be done purely within that format, and how far "up" the ladder of cognitive sophistication these processes go. And work is already being done in this respect, about the potential for perceptual representation to generalize (Buckner, 2018), and about the role of imagery in decision making (Briscoe, 2018; Nanay, 2016) and in skill (Burnston, in submission). Moreover, theorizing about the *interaction* of sentential and imagistic representations opens up new possibilities for thinking about cognitive penetration (Burnston, 2017b) and the causal theory of action (Burnston, 2017a). As is often the case with Fodor, he is at his philosophically most fruitful as a foil for developing alternative views.

REFERENCES

Barsalou, L. W. (1999). Perceptual Symbol Systems. *Behavioral and Brain Sciences, 22*(4), 577-609.

Barwich, A.-S. (2014). A sense so rare: Measuring olfactory experiences and making a case for a process perspective on sensory perception. *Biological Theory, 9*(3), 258-268.

Bayne, T. (2016). Gist! *Proceedings of the Aristotelian Society, 116*(2), 107-126. doi:10.1093/arisoc/aow006

Briscoe, R. E. (2018). Superimposed Mental Imagery. In *Perceptual Imagination and Perceptual Memory* (pp. 161).

Brozzo, C. (2017). Motor intentions: How intentions and motor representations come together. *Mind & Language, 32*(2), 231-256.

Buckner, C. (2018). Empiricism without magic: transformational abstraction in deep convolutional neural networks. *Synthese, 195*(12), 5339-5372. doi:10.1007/s11229-018-01949-1

Burnston, D. C. (2016). The evolution of evolutionary psychology. http://cognitionandculture.net/webinars/the-shape-of-thought-book-club/the-evolution-of-evolutionary-psychology-1457180813/

Burnston, D. C. (2017a). Cognitive penetration and the cognition–perception interface. *Synthese, 194*(9), 3645-3668. doi:10.1007/s11229-016-1116-y

Burnston, D. C. (2017b). Interface problems in the explanation of action. *Philosophical Explorations, 20*(2), 242-258.

Burnston, D. C. (2017c). Is aesthetic experience evidence for cognitive penetration? *New Ideas in Psychology, 47*, 145-156.

Burnston, D. C. (in submission). There are higher-level perceptual contents, and they're non-conceptual.

Burnston, D. C., & Cohen, J. (2015). Perceptual integration, modularity, and cognitive penetration. In *Cognitive Influences on Perception: Implications for Philosophy of Mind, Epistemology, and Philosophy of Action*. Oxford: Oxford University Press.

Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research, 88*(1), 119-145.

Carruthers, P. (2015). *The centered mind: What the science of working memory shows us about the nature of human thought*: OUP Oxford.

Ferretti, G., & Caiani, S. Z. (2018). Solving the interface problem without translation: The same format thesis. *Pacific Philosophical Quarterly*.

Fodor, J. A. (2007). The revenge of the given. In J. Cohen & B. P. McLaughlin (Eds.), *Contemporary debates in philosophy of mind* (pp. 105-116). Malden, MA: Blackwell.

Fodor, J. A. (1968). The appeal to tacit knowledge in psychological explanation. *the journal of philosophy, 65*(20), 627-640.

Fodor, J. A. (1975). *The language of thought* (Vol. 5): Harvard University Press.

Fodor, J. A. (1983). *The modularity of mind*: MIT press.

Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind* (Vol. 2): MIT press.

Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*: Clarendon Press/Oxford University Press.

Fodor, J. A. (2001). *The mind doesn't work that way: The scope and limits of computational psychology*: MIT press.

Fodor, J. A. (2003). Hume variations.

Fodor, J. A. (2008). *LOT 2 : the language of thought revisited*. Oxford; New York: Clarendon Press ; Oxford University Press.

Fodor, J. A., & Lepore, E. (1992). Holism: A shopper's guide.

Folstein, J. R., Gauthier, I., & Palmeri, T. J. (2012). How category learning affects object, representations: Not all morph spaces stretch alike. *Journal of Experimental, Learning, Memory & Cognition, 38*(4), 807-820.

Gauker, C. (2011). *Words and images: An essay on the origin of ideas*: Oxford University Press.

Gauthier, I., & Tarr, M. J. (2002). Unraveling mechanisms for expert object recognition: bridging brain activity and behavior. *Journal of Experimental Psychology: Human Perception and Performance, 28*(2), 431.

Goldie, P. (2007). Seeing what is the kind thing to do: Perception and emotion in morality. *Dialectica, 61*(3), 347-361.

Green, E. J., (forthcoming). The perception-cognition border: A case for architectural division. *The Philosophical Review*.

Green, E. J., & Quilty-Dunn, J. (2017). what is an object file? *The British Journal for the Philosophy of Science*.

Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General, 130*(1), 116.

Gureckis, T. M., & Goldstone, R. L. (2008). *The effect of the internal structure of categories on perception.* Paper presented at the Proceedings of the 30th Annual Conference of the Cognitive Science Society.

Landy, D., Allen, C., & Zednik, C. (2014). A perceptual account of symbolic reasoning. *Front Psychol, 5*, 275.

Larkin, J. H., & Simon, H. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive science, 11*, 65-99.

Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition, 104*(1), 19-46.

Mandelbaum, E. (2017). Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception. *Philosophy and Phenomenological Research*.
Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology, 8*(2), 317-336.
Mylopoulos, M., & Pacherie, E. (2018). Intentions: The dynamic hierarchical model revisited. *Wiley Interdisciplinary Reviews: Cognitive Science*, e1481.
Nanay, B. (2016). The Role of Imagination in Decision-Making. *Mind & Language, 31*(1), 127-143.
O'Callaghan, C. (2012). Perception and Multimodality. In E. Margolis, R. Samuels, & S. Stich (Eds.), *The Oxford Handbook of Philosophy of Cognitive Science* (pp. 92-117). Oxford: Oxford University Press.
Quilty-Dunn, J. (2016). Iconicity and the format of perception. *Journal of Consciousness Studies, 23*(3-4), 255-263.
Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*(3), 681.
Shepherd, J. (2017). Double life of intention. *Philosophy and Phenomenological Research*.
Siegel, S. (2006). Which properties are represented in perception. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 481-503). New York: Oxford.
Silins, N. (2013). The significance of high-level content. *Philosophical Studies, 162*(1), 13-33.
Sowden, P. T., Davies, I. R. L., & Roling, P. (2000). Perceptual learning of the detection of features in X-ray images: A functional role for improvements in adults' visual sensitivity? *Journal of Experimental Psychology: Human Perception and Performance, 26*(1), 379-390. doi:10.1037/0096-1523.26.1.379