

*This is a very, very preliminary draft.  
Comments are solicited, but unlimited deniability  
is reserved.*

## **Revenge of The Given**

Jerry Fodor

### *1 Introduction*

Could there be unconceptualized representations? What would they be like? Where might one look for them? How would you know if you'd found one? I think these questions are, in the long run, largely empirical. But limning the landscape that they occupy is a philosophical enterprise within the meaning of the act, and that I shall presently try to do. I also propose to push at them and prod at them until they're in a form where psychological data can be seen to bear; and to say something about how the relevant data have come out so far.

For present purposes, I assume without argument that there is *conceptualized* representation; indeed, that there is conceptualized *mental* representation; indeed, that there are conceptualized mental representations. These assumptions are the backbone of the 'Representational Theory of Mind', which is, of course, famously tendentious; but I'm too old to worry about what to do if RTM isn't true.

I shall also assume what's maybe less familiar: that *conceptualized representation* is *representation as* and vice versa. So, to represent Mr James as a cat is to represent him as falling under the concept CAT; and to represent Mr James as falling under the concept CAT is to represent him as a cat; and thinking of Mr James as a cat requires applying the concept CAT to Mr James. This is, to be sure, just the sort of 'intellectualist' kind of thinking about thinking of that mid-century logical behaviorists (notably Ryle and Wittgenstein) said that psychology and the philosophy of mind could profitably do without. Well, they were wrong.

If that's all granted, it suggests a first move; we have the option of exchanging 'can there be unconceptualized mental representation' for 'can there be mental *representing* without mental *representing as*?' This serves to locate the present concerns in relation to a family of others that are philosophically familiar: for example, 'represents X' is transparent

to substitution of coextensive singular terms at the X position; but 'represents as X as F' is opaque to substitution of coextensive predicates at the 'F' position. If (a token of) 'that cat' represents that cat, and if that cat is Granny's favorite cat, then it follows that that token of 'that cat' represents Granny's favorite cat. If, however, 'that cat' represents that cat *as* that cat, it doesn't follow that 'that cat' represents that cat as Granny's favorite; not even if that's the cat that Granny prefers to any other. In this respect the semantical distinction between *representating and representing as*, works like the psychologist's distinction between *seeing* and *seeing as* since you can see that cat without seeing it as that cat, but you can't see that cat *as* that cat without seeing it as a cat. That's all implicit in RTM, according to which seeing X only requires mentally representing X somehow or other, but seeing X as F requires applying to X the mental representation that expresses the concept F.

Conversely, according to RTM, if there is seeing without seeing as, then there is unconceptualized mental representation. Piety suggests a more traditional formulation: If there is seeing without seeing as, then there is a 'perceptual given.' In what follows I will often put the matter this second way, but with a caveat: I assume, contrary to a main epistemological tradition, that the given may be both subpersonal and encapsulated; which is to say that it may be neither conscious nor (in Steven Stich's term) 'inferentially promiscuous'. A fortiori, the given in perception might be inaccessible either to voluntary report or to cognitive processes that aren't perceptual (thinking, for example), or to both. This is, to be sure, the thin edge of a complicated tangle of issues. Consider inferential promiscuity. It's arguable that the possibility of certain inferential elaborations figures in the very individuation conditions for certain concepts; in such cases, the connection between the content of the representation and its inferential role would be inalienable and certain kinds of encapsulated mental representation would be impossible in principle. Perhaps you can't alienate DOG from its role in the inference from DOG to ANIMAL. If so, then, arguably, a mental process that has access to DOG must have access to ANIMAL as well. Likewise, it's arguable that there can't be perceptual content, given or otherwise, that is unavailable for report, perhaps because the connection between perceptual content and first person justifications of perceptual judgments is inalienable. In that case, the notion of an unconscious given would be a contradiction in terms.

A word or two, before we proceed, about a case where these sorts of issues arise. Christopher Peacocke considers the suggestion that, all else equal, a thinker who has the concept SQUARE and who is in the presence of a good instance of a square “must find the present-tense demonstrative thought that that object is a square to be primitively compelling (Gunther, 116-117).” If that is so, then, there’s a necessary connection between having SQUARE and having a certain capacity to recognize squares. This suggestion Peacocke rightly rejects. For, consider Mach’s ‘square /diamond’ illusion (Figure 1). “A thinker, taking his experiences at face value and possessing this concept SQUARE, need not find it primitively compelling (without further reflection) that a floor-tile in the diamond orientation is square. But it can still be that... the region of space apparently occupied by the floor-tile is square--- as indeed it will be if his experience is veridical.”

Figure 1 about here

Quite so; someone who has the concept SQUARE can have the kind of experience that seeing a good instance of squareness typically causes, even though he fails to apply the concept SQUARE when he has the experience (eg. when he sees 1B). Moreover, someone can have the kind of experience that seeing a good instance of squareness engenders even though he lacks the concept SQUARE (eg. when he sees 1A). So having the concept SQUARE isn’t sufficient, or even necessary, for having that sort of experience<sup>1</sup>. So, why doesn’t that settle the question whether there can be unconceptualized experiential content?

Well, because it’s one thing to say that you can have the sort of experience that’s typical of seeing squares even though you lack the concept SQUARE; it’s something quite else to say that you can have that sort of experience even if you lack concepts entirely. Peacocke remarks, “the difference between perceiving something as a square and perceiving it as a (regular) diamond is in part a difference in which symmetries are perceived. When something is perceived as a diamond, the perceived symmetry is a symmetry about the bisector of its angles. When something is perceived as a square, the perceived symmetry is a symmetry about the bisector of its sides. But does[n’t] perceiving something as symmetrical require the perceiver to possess and make use of the concept SYMMETRICAL? If so, then

---

<sup>1</sup> As Peacocke also points out, experiences as-of-squares are plausibly what one learns the concept SQUARE from. A fortiori, having the experience can’t require having the concept.

what we will have done is simply to explain (this part of)\_mastery of the concept SQUARE in terms which require mastery of other concepts (118)”

So it would appear. But then, why does Peacocke think that the square/diamond illusion argues for unconceptualized experiential content (as opposed to experiential content that isn't conceptualized by SQUARE?) I think it's because he takes for granted a doctrine I've elsewhere called 'conceptual pragmatism' (Fodor, forthcoming): namely, that having a concept is knowing how to apply it and/or knowing what to infer from its applicability.<sup>2</sup> If that's so, then there's *no such thing* as having a concept that is accessible only to subpersonal cognitive processes; for example, there's no such thing as a concept of symmetry that is accessible to a perceiver's visual system but not to the perceiver. In such a case, seeing symmetrical things as symmetrical wouldn't, per se, manifest possession of the concept SYMMETRY, so my facile identification of seeing as with seeing under a concept would be undermined. The most that would follow from one's ability to see symmetrical things as symmetrical is that *one's visual system* is able to represent symmetry as such.

In short, it might seem that concepts can be classified either according to their content or according to degree of their encapsulation, and that these classifications are orthogonal. Not so, however, for a conceptual pragmatist who holds that the 'use' of a concept, (either for inference, or for application, or for report (or, perhaps, for all three)) is ipso facto inalienable. For such a conceptual pragmatist, a capacity to represent a symmetry as a symmetry needn't manifest a grasp of the concept SYMMETRY. Hence Peacocke's view, which is that seeing the symmetries in 1A and 1B doesn't count as conceptualizing either as symmetrical. Well, talk that way if it pleases you to do so; but I don't think that the decision can reasonably be isolated from such exigencies as the empirical investigations of cognition may impose.<sup>3</sup> I think, for example, that the very same concept (eg SYMMETRICAL) that one mind has in play in processes of visual form recognition, may, in some other mind (or in the same mind at a different developmental stage), be available to mechanisms of conscious problem solving. It isn't *what concepts they have* that distinguishes these minds, it's the uses to which they are able to put them.

---

<sup>2</sup> Note Peacocke's endorsement of "Evans Thesis," p 114 .

<sup>3</sup> That is: the question whether subpersonal concepts belong to the same natural kind as concepts *tout court* depends on the similarity (or otherwise) of their roles in psychological explanations; in particular on the identity (or otherwise) of the empirical laws that subsume them. It is thus not an issue that's available to a priori resolution.

If your view is that concepts are constituted by the uses they are put to, perhaps you might wish to consider revising that view. The question whether something is given in perception is distinct from the question whether what's given in perception is subpersonal and/or encapsulated. Philosophical discussions of unconceptualized content (like, come to think of it, philosophical discussions of practically everything else) routinely take for granted that the content of a mental representation is dissociable from its role in inference. They therefore beg the issue between conceptual pragmatists and anybody who isn't a functionalist about mental content (for example, between conceptual pragmatists and conceptual atomists like me). I'm supposing, to the contrary, that this issue is wide open. In fact, the hidden agenda of this paper to illustrate how important it is, not just in philosophy but also in the psychology of cognition, to keep it in mind that content functionalism may not be true.

## *2. Kinds of representations*

For reasons I'll presently set out, I think that there probably are unconceptualized perceptual representations. The line of argument I'll have on offer goes like this: On one hand, it's (empirically) plausible that some perceptual representation is iconic and, on the other hand, it's in the nature of iconic representation not to be conceptualized. That being the intended polemical strategy, I had better now say something about what I mean by 'iconic representation'

First, then, my usage is idiosyncratic. In the semantics/semiotics literature, 'iconic' frequently comports with notions like, for example, 'pictorial' and 'continuous'. But it's not always clear just what either of these comes to, or just what the connections between them are supposed to be; as often as not, they're made to take in one another's wash. For the moment, I propose to pretend that the slate is blank and just stipulate:

First, 'iconic' and 'discursive' are mutually exclusive modes of representation; that a representation is either entails that it's not the other. I leave it open that some kinds of representation are neither iconic nor discursive. Off hand, I can't think of a good candidate but it doesn't matter for the present purposes.

Second, I assume for familiar reasons that the kinds of representations we're concerned with in theories of cognition are compositional. To a first approximation, a representation is compositional iff its syntactic structure and semantic content are both

determined by the syntactic structure and the semantic content of its parts. Compositionality is required by any serious theory of linguistic and/or mental representations because both thought and language are productive and systematic.<sup>4</sup> I suppose everybody knows this story, so I won't elaborate. (For discussion cf. Fodor and Lepore 19xx passim). According to my usage, the defining distinction between iconic and discursive representations turns on a difference between the way they achieve their compositionality

*Discursive representations:* The sentences of natural languages are the paradigms; here again the outlines are familiar. Every sentence is a finite arrangement of constituents that are themselves either primitive or complex. Each complex constituent is a finite arrangement of 'lexical' primitives (words, near enough). Lexical primitives have their syntactic and semantic properties intrinsically; roughly, a word is a triple consisting of a bundle of orthographic/phonological features, a bundle of syntactic features, and a bundle of semantic features. (These are enumerated by the word's 'entry' in the 'lexicon' of the language). A discursive representation in L is syntactically compositional iff its syntactic analysis is exhaustively determined by the grammar of L together with the syntactic analysis of its lexical primitives. A discursive representation is semantically compositional iff its semantic interpretation<sup>5</sup> is exhaustively determined by its syntax together with the semantic interpretations of its lexical primitives.

Consider, for example sentence (1). Its syntactic structure is (more or less) as shown in (2), and its semantic interpretation is (more or less) *that John loves Mary*

1. John loves Mary
2. (John)<sub>NP</sub> ((loves)<sub>V</sub> ((Mary<sub>NP</sub>))<sub>VP</sub>)<sub>S</sub>

The syntax and semantics of the sentence are determined by such facts as that 'John' is a noun and denotes *John*, that 'loves' is a verb and denotes the relation *X loves Y* and that 'Mary' is a noun and denotes *Mary*. Further details are available upon application at your local department of linguistics

---

<sup>4</sup> However, see Fodor (19xx)

<sup>5</sup> It helps with the exposition to identify the semantic interpretation of a sentence with its truth conditions. This is, of course, wildly tendentious, but I don't propose to trade on it. It's, by the way, important to keep an eye on the ambiguity between, on the one hand, *interpretations qua representations that the grammar of L assigns to its sentence types* and, on other, *interpretations qua representations that speaker/hearers of L assign to tokens of L in the course of communication exchanges*. Which is which will be apparent in context wherever it matters. I hope.

What matters for us is this: the semantic interpretation of a sentence (*mutatis mutandis*, of any discursive representation) depends exhaustively on the way that properties of its lexical primitives interact with properties of its constituent structure; and *not every part of a discursive representation is ipso facto one of its constituents*.<sup>6</sup> So, for example, ‘John’ ‘Mary’ and ‘loves Mary’ are among the constituents of (1) according to the analysis (2) But ‘John loves’ isn’t, and nor is ‘John... Mary’.<sup>7</sup> This is part and parcel of the fact that neither the semantic interpretation of ‘John loves’ nor the semantic interpretation of ‘John...Mary’ contributes to determining the semantic interpretation of ‘John loves Mary’; in fact, neither of them *has* a semantic interpretation in that sentence (though, of course, each of the lexical primitives they contain does). I’ll say: the constituents of a discursive representation are those of its parts that are recognized by its *canonical decomposition*. According to me, it is *having a canonical decomposition* that distinguishes discursive representations from iconic ones.

*Iconic representations:* Pictures are paradigms (but see the caveats to follow) I suppose that pictures, like sentences, have a compositional semantics. Their principle of compositionality is this:

*Picture Principle:* if P is a picture of X, then parts of P are pictures of parts of X<sup>8</sup>

Pictures and the like differ from sentences and the like in that icons don’t have canonical decompositions; they have *interpretable parts*, but they don’t have *constituents*. Or, if you prefer, all the parts of a picture are ipso facto among its constituents; icons are compositional according to the Picture Principle whichever way you carve them up.<sup>9</sup> Take a picture of a person, cut it into parts however you like; still, each picture-part pictures a

---

<sup>6</sup> Lexical primitives count as constituents qua syntactically and semantically interpreted by their lexical entries. Sentences are constituents by courtesy (*viz.* they are ‘improper’ constituents).

<sup>7</sup> It’s plausible that the constituents of a sentence can be discontinuous; cf ‘look...over’ in ‘John looked the situation over’.

<sup>8</sup> I’ll assume, for ease of exposition, that *P is of X* is the only semantic relation that holds between pictures and ‘things in the world’; a fortiori, it’s the only such relation that’s compositional. I don’t think it matters to my line of argument whether that is actually so, (I likewise leave it open that some pictures aren’t *of anything*).

<sup>9</sup> Presumably, however, some parts of any icon are too small to have interpretations; the atoms that it’s composed of, for example. I’ll usually ignore such matters of ‘grain’, but see below in the text

person-part. And the whole that you have if you reassemble all the pictures parts is a picture of the whole person that the parts are pictures of.<sup>10</sup>

So, then, in everything that follows, a representation that has no canonical decomposition is an icon. I will argue (quite soon now) that having a canonical decomposition is a necessary condition for a representation to be conceptualized; hence that the question we started with ‘are any mental representations unconceptualized’ can be swapped for the question ‘are any mental representations iconic?’ And that, finally, is a question on which empirical evidence can be brought to bear. First, however, I digress for some caveats.

#### *Digression on icons*

I’ve taken pictures as my paradigms of icons, but I don’t mean to suggest that they are the only examples. I suppose that graphs aren’t pictures, but if you draw a curve that represents the distribution of a property in a population, a part of the curve ipso facto represents the distribution of the property in a part of the population (which is to say that it partially represents the distribution of the property in the population.) So graphs are (typically) icons according to my usage. Since such icons aren’t pictures (you can’t *picture* the distribution of a property in a population), it can’t be assumed that icons represent what they resemble; not even on the (mistaken) assumption that pictures do.

Likewise, iconic representations needn’t be “continuous.”. (There is, for example, the boring point about grain that I mentioned in fn10). The Picture Principle can perfectly well apply to a representation that has fewer than continuously many parts, so long as whatever parts it has have interpretations. Nor, conversely, does continuous representation have to be iconic. Conventionalized continuous representations, like the speed lines that grow from cartoon characters in motion, often aren’t. On the other hand, iconicity does offer the possibility of representations with a nondenumerable infinity of interpreted parts;

---

<sup>10</sup> It’s arguable that some ‘things in the world’ (things other than representations) also have canonical decompositions; watches, people and the like. So be it if it wants to be; the Picture Principle holds in any case. On the one hand, any part of a picture of a person is a picture of *some* part of the person, though the part needn’t be canonical. On the other hand, oceans don’t have canonical parts, but any part of a picture of the ocean is a picture of a part of the ocean.



whereas, by definition, discursive representation does not.<sup>11</sup> I mention this because some of the recent philosophical discussion of unconceptualized mental representation has turned on the continuity issue.

Thus McDowell on Evans: (Gunther, 82): “One consideration that impresses Evans is the determinacy of detail that the content of experience can have... Evans suggests that our repertoire of color concepts is coarser in grain than our abilities to discriminate shades, and therefore unable to capture the fine detail of color experience... Color experience can present properties that correspond to something more like lines on the spectrum, with no discernible width.” McDowell is less impressed with this line of thought than I think he should be: “...why should we accept that a person’s ability to embrace color within her conceptual thinking is restricted to concepts like ‘red’ or ‘green’...In the throes of an experience that putatively transcends one’s conceptual powers... one can give linguistic expression to a concept that is exactly as fine-grained as the experience, by uttering a phrase like ‘that shade,’ in which the demonstrative exploits the presence of [a] sample ”

I think this reply rather misses the point. True, you can demonstrate any shade that you may encounter, whether or not you have a name for it.<sup>12</sup> It’s nonetheless very plausible that your capacity for experiencing shades outruns even your capacity for demonstrating them, and that this would be so even if you were to demonstrate forever. Quanta aside, there is a shade between any two that you can distinguish. That one can’t always *see* the intermediate shade has perhaps to do with the acuity of one’s visual system, not with its representational capacities.<sup>13</sup> Evans is surely right that one can mentally represent as many shades as there are “lines on the spectrum, with no discernible width.” But, more to the point, there are as many shades that one can mentally represent as there are lines on the spectrum with no width at all, discernible or otherwise. This isn’t, of course, a paradox; one

---

<sup>11</sup> Because every constituent of a discursive representation is a finite arrangement of discreet primitive parts that are themselves selected from a finite lexicon.

<sup>12</sup> This amounts, by the way, to quite a lot of shades. It’s estimated that there are around xxx shades of them that one can tell apart. It’s a good rule of thumb that a cognitive capacity that big is likely to be infinite.

<sup>13</sup> Of course, you always have concepts enough to *think about* a color in between (reading ‘think about’ de re). The present question is whether you always have concepts enough to, as it were, present the content of the experience to thought. My view is that presenting contents to thought is par excellence the function of conceptualized representations; hence, par excellence a function that *nonconceptual* representations can’t perform. This is, I think, more or less consonant with the traditional view of the given (though it’s no good for the likes of Hume, who thinks that a thought is itself a kind of pale sensation): What’s given in experience is often *represented* in thought. But what’s given in experience is *given* only in experience.

can represent continuously many lines if one has access to continuously many representations. If you want a drawing that represents each of the points on a line, draw another line.

So maybe there is reason to believe that experience represents colors along a continuous dimension of chromaticity.<sup>14</sup> If I had to bet, I'd bet that way. I'm not, however, convinced that it matters greatly to our main topic, which is whether there is *unconceptualized* mental representation; for, as far as I can see, there's no reason why representation by continuous magnitudes should ipso facto count as not conceptual. It sets the bar pretty high to claim (as McDowell does) that it's a priori that there can't be representation without conceptualization. It strikes me as downright unreasonable to claim a priori that there can't be representation without language (though it seems McDowell wants to do that too). If experiential representations are continuous, then of course they aren't linguistic. But (short of stipulation) their not being linguistic wouldn't entail their not being conceptualized. By contrast (so I'll claim), if experience is iconically represented, then it does follow that it isn't conceptualized (in at least one plausible sense of the term). We return, then, to the main line of the discussion.

*Iconicity and individuation:*

So far: Iconic representations are typically semantically evaluable (they are typically *of this or that*). But they have no canonical decompositions; which is to say, they have no constituent structure; which is to say that, however they are sliced, there's no distinction between their *canonical* parts and their *mere* parts. Here's another way to put this: an icon is a homogeneous kind of symbol from both the syntactic and the semantic point of view. Each of its parts ipso facto gets a semantic interpretation according to the same rule of interpretation that applies to each of the others (*viz.* the Picture Principle).

But none of that is true of discursive representations. Only a specifiable subset of the parts of a discursive symbol are syntactic or semantic constituents; and it's thus far open that the various constituents of a discursive representation may contribute in different ways to determining the semantics of their hosts. Our paradigms, the sentences of a natural language,

---

<sup>14</sup> Cf Peacocke (op cit, 111-112) who, however, describes such continuous representations as "analogue". Terminological consensus isn't widely available in this neck of the woods.

are clearly structurally heterogeneous in this respect. Considered syntactically, they contain: Nouns, Verbs, Adjectives, NPs, VPs, PPs, and so on. Considered semantically, they contain: singular terms, descriptions, predicates, (including complement structures) and an apparatus of logical terms like quantifiers, variables, and connectives, and so on once again. Correspondingly, both the rules that distinguish sentential constituents from mere sentential parts, and the rules that compose the interpretation of sentential expressions from the interpretation of their constituents, turn out to be disconcertingly complex and hard to state; linguists have thus far had only very partial success in formulating either. Compare the unarcane apparatus that sufficed to formulate the Picture Principle.

Because they decompose into syntactically and semantically heterogeneous constituents, discursive representations can have logical forms<sup>15</sup> (maybe all discursive representations have that can express truths). By contrast, because they decompose into syntactically and semantically homogeneous parts, iconic representations don't have logical forms. I take that to be truisitic. The logical form of a symbol is supposed to make its compositional structure explicit; viz to make explicit the contribution that each of the interpreted parts contributes to its content. But *every* part of an iconic symbol is interpreted, and each part contributes in the same way to determining the content of the whole icon that it's part of: each part pictures part of what the whole icon pictures.

The heterogeneous structure of their parts matters to the representational capacity of discursive symbols; it means, in particular, that they can represent things that icons are unable to. The various constituents of a discursive representation generally contribute in different ways to determining the content of their hosts; singular terms contribute in one way, predicates contribute in quite another way, and logical constants in still another. But the picture Principle says that every part of an icon contributes to its interpretation in the same way: it pictures part of what the icon does. In consequence, icons can't express (eg) the distinction between negative propositions and affirmative ones, since that turns (inter alia) on distinctions among logical constants. For similar reasons, they can't express quantified

---

<sup>15</sup> It would do well enough, for our purposes, to identify the logical form of a sentence with its canonical decomposition; namely with a representation that specifies each of its semantically interpreted constituents. However this is just heuristic since sentences are generally supposed to have semantically interpreted parts that don't occur as constituents (that is, as constituents of their 'surface' structure): thus the "understood" subject of 'to run' in 'John likes to run'. Perhaps the logical form of a sentence should be identified with its constituent structure at some relatively abstract level of grammatical description ('LF' as one says). Having mentioned this, I shall now proceed to ignore it.

propositions; or hypothetical propositions; or modal propositions.<sup>16</sup> They can't even express predication, since that requires (inter alia) distinguishing terms that contribute individuals from terms that contribute sets (or properties, or whatever).<sup>17</sup>

For reasons that are quite closely related, whereas discursive representations typically carry ontological commitments,<sup>18</sup> iconic representations don't. In particular, discursive representations do, but iconic ones do not, impose principles of individuation on the domains in which they are interpreted. I don't want to talk about this at length because I'm scared to. So it would help enormously if you'll just let me assume that what individuals a system of representation is ontologically committed to depends on the apparatus of quantifiers, variables, singular terms, and sortal predicates to which it has access. To a first approximation, systems of representation are committed to the individuals over which they quantify; conversely, if the available representations don't include quantifiers (or classifiers or something of the sort), then there won't be principles of individuation for whatever it is that the representations are of. Lacking that sort of apparatus, there is no right answer to the question 'which things (how many things?) does this symbol represent?'<sup>19</sup> (Didn't Quine say something of that sort? I hope he did; I would so like to be in respectable company for a change).

---

<sup>16</sup> This is continuous with the standard objection to 'resemblance' theories of mental content (which, of course, presupposes that 'ideas' are kinds of icons.) There is nothing in John's not loving Mary for a part of a picture to resemble.

<sup>17</sup> For very closely related reasons, pictures don't have truth conditions. In the root case, for a symbol to be true it has to pick out an individual and a property and predicate the latter of the former, but iconic representations have no way to do either. So the camera doesn't lie, but neither does it tell the truth.

<sup>18</sup> It's notoriously moot whether, and under what conditions, the ontological commitments of a system of discursive representation might be unique (for example, whether there's a fact of the matter about what the representations refer to). I propose not even to mention such matters; but, for what it may be worth, I do have views about them. The putative argument for referential undetermination is that there is, even in principle, nothing 'in the linguistic data' that could distinguish (eg) commitment to an ontology of rabbits from commitment to an ontology of undetached rabbit parts. I believe that this claim is false. For discussion, see Fodor 19xx.

<sup>19</sup> It is, however, of prime importance *not* to read 'having a principle of individuation' epistemically. If I have the concept CAT, I know the principle of individuation for cats; viz. that *each cat is one cat, and nothing else is*. So, it's true, in the following sense, that if I have the concept CAT, then I know how to count cats: I must count one for each cat and nothing for anything else; I must count each cat and nothing else *as* a cat. But what does *not* follow is that, if I have the concept, I am thereby able accurately to assign cardinalities to arbitrary aggregations of cats. Likewise, to know that each sheep is one sheep is to command a metaphysical truth of some depth. But it doesn't imply a procedure (an algorithm; a 'criterion') for estimating the size of flocks. Confusing metaphysics with epistemology (and/or semantics) is the defining philosophical disease of our time, and it seems that the psychologists have caught it too. See, for example, Xu (19xx) and her commentators.

To be sure, a photograph may show three giraffes in the veldt; but it likewise shows a family of giraffes; and an odd number of Granny's favorite creatures; and a number of Granny's favorite odd creatures; and a piece of veldt that's inhabited by any or all of these. No doubt, we usually can agree about how to understand the ontology of such a photograph; we do so in light of whatever project we happen to have in hand. But that's not the relevant consideration; what matters is that the discursive symbol 'three giraffes in the veldt' specifies a scene relative to such concepts as THREE, GIRAFFES, IN and THE VELDT. A fortiori, a mind that lacks these concepts, can't use that symbol to represent the scene. Contrast iconic representation: you can, of course, *see* three giraffes in the veldt without having GIRAFFE, etc.. Nor do you need them to take a picture of three giraffes in the veldt; a camera and film will suffice.

Equivalently (more or less): the context 'iconally represents...' is like the contexts 'sees...', and 'describes...' 'points at...'; 'and 'photographs...'. they are all transparent to substitution of coextensive descriptions. But 'discursively represents...' is like 'sees as', and 'describes as...'; it always has an opaque reading (which, in fact, it usually prefers). According to RTM, that's because *seeing as...* and *describing as...*, like other acts of conceptualization, operate by subsuming distal things under the concept that is expressed by the predicate of some mental representation. It's entirely in the spirit of RTM that 'conceptualizing' and 'predicating' are two ways of talking about the same thing.

Brief review: We started with conceptualized v. unconceptualized representations. We swapped that for *representing as* v. *representing tout court*, which we then swapped for iconic v. discursive representation. This allowed us to replace the question whether there are conceptualized mental representations (in particular, whether there's a perceptual given) with the question whether any mental representation is iconic. I then suggested that (because they lack logical form) iconic representations don't provide principles of individuation for their domains of interpretation. This suggests a final metamorphosis: 'Are there unconceptualized representations?' becomes 'are there mental phenomena in which representation and individuation are dissociated?' If there are, then that's prima facie evidence of nonconceptual mental representation.

We're just about to see what such evidence looks like. But first, a quick consideration of an a priori objection that many philosophers appear to find persuasive.

‘Look,’ you might say, ‘there just couldn’t be any such data. For, if a symbol represents a such and such, it *must* represent it as a so and so. There aren’t, as it were, two kinds of representing, one of which ignores this maxim and is thus transparent, and the other of which obeys it and is thus opaque. Rather, we get the ‘transparent’ reading of ‘represents...’ by abstracting from some or other opaque one; viz by ignoring the *mode* of representation in certain specifications of a representational *content*. There are, to be sure, two ways one might read an assertion that John believes the king of France is bored: according to the opaque construal, the speaker takes responsibility for the definite description, according to the transparent construal he doesn’t. (See eg. Brandom, 19xx; Dennett, 19xx). But that’s a distinction between two styles of belief *ascription*, not between *two styles of believing* or two ways of mentally representing. Likewise, *mutatis mutandis*, for the difference between representing and representing as; representing a thing is always representing it as this-or-that; though one always has the option, in saying that something is *represented*, of not bothering to say *as what*.’

That sort of objection has force on the assumption that *representing* is indeed an abstraction from *representing as*. But that assumption is not available in the present polemical context; that representing requires representing as is supposed to be the conclusion of the argument, not its premise. Suffice it for present purposes then, if we can imagine, even roughly, how X might represent Y without representing it under some description or other<sup>20</sup> I think, in fact, that there’s a plausible candidate: X represents Y insofar as X carries information about Y, where ‘carries information about...’ is itself read as transparent. (Information is (more or less) correlation according to standard accounts<sup>21</sup> and I don’t suppose anybody thinks that correlation is description-relative). Maybe construing content as a kind of information in the case of the given, would allow for *representing* without *representing as*, and hence allow for unconceptualized representing.

Well then, is there a serious possibility of thinking of the content of the given in some informational sort of way? This is another of those questions that scare me, so I don’t

---

<sup>20</sup> As, by the way, ‘bare’ demonstratives are sometimes supposed to do.

<sup>21</sup> See, eg, Dretske, 19xx; Fodor 19xx etc.

<sup>22</sup> If there is a correlation then, presumably, there will be a covering law which explains it, and there will be counterfactuals that the law supports. It is thus important to recognize that ‘correlates with’ and ‘carries information about...’ can be transparent even if ‘nominally explains’ contexts and subjunctive contexts are not. According to the usual construal, events typically carry information about the events that caused them, but *not* about what it was about the events that caused them in virtue of which they did so.

propose to undertake it. Suffice it that the crux is often supposed to be whether there could be an informational construal of *misrepresentation*; and, I'm not at all certain that there can't be. (For discussion, see Fodor, xxxx) I proceed, therefore, to conditionalize: Assuming that an informational construal of content allows, at least in principle, for *representing* without *representating as*, the empirical issue is whether there is experimental evidence for such a dissociation in the psychology of perception. Well, I think there is; in fact, I think there's lots of it of several different kinds. Selected examples to follow.

To begin with, if you want to test a theory, you need (what used to be called) 'correlating definitions' Here's one: it's a rule of thumb that, all else equal, the 'psychological complexity' of a discursive representation (for example, the amount of memory it takes to store it) is a function of the number of individuals whose properties it independently specifies. I shall call this the 'item effect.'

Consider, as it might be, phone books. They specify properties of individuals (their numbers and addresses), and they are explicit as to both the individuals and the properties.<sup>22</sup> All sorts of things follow: the phone books of big cities are generally bigger than the phone books of small cities; and they take up more shelf space; and it takes longer to look up an arbitrary number in a big phone book than in a small one; and it's harder to memorize (or even to copy) the contents of a big book than that of a small one; ... and so forth. This is all because the representations in phone books are discursive, hence conceptualized; they presuppose the possession and employment of such concepts as X'S NAME IS 'Y' and PHONE P HAS THE NUMBER N Lists, like sentences, are paradigms of discursive representation: they exhibit effects of their content (it's the number that's listed for John that you proceed to dial in consequence of looking his number up); and they also exhibit an effect of the number of items they contain.

Compare photographs: A photograph of 60 giraffes takes no more space in your album (or on the screen) than a photograph of 6 giraffes. For that matter, it takes no more space than a photograph of no giraffes (the one that you made when you forgot to take the lens cap off). Photographs are time sensitive (very old ones are generally more degraded

---

<sup>22</sup> That is, the individuals and properties are enumerated, not merely quantified over. 'All men are mortal' assigns mortality to more men than 'two men are mortal', but the complexity of these representations doesn't differ in consequence. To repeat the text, the magnitude of the 'items effect' is roughly proportional to the number of individuals whose properties are *independently* represented.

than very new ones) but they aren't item sensitive. This is hardly surprising in light of the preceding: iconic representations don't individuate; they don't represent individuals *as* individuals. A fortiori, nothing about their behavior depends on the number of individuals that they represent.

Still, I do want to emphasize that failing to find item effects is not, all by itself, a knock down test for iconic representation; as usual in empirical inferences, alternative explanations have to be ruled out. For example, tasks that involve searching for an item in a stimulus array ('Where's Waldo?') may be insensitive to the size of the array if the search is carried out 'in parallel'. So, given a search in which there's no item effect, one sometimes can't tell whether that's because iconic representation is involved or because there's parallel processing.

On the other hand, sometimes one can. Roughly, the discursive/iconic distinction is orthogonal to the parallel/serial distinction. Some parallel searches do involve conceptualizing the items in the array, and others don't, and it is (in principle) possible to distinguish between the two on empirical grounds. For example, the former, but not the latter, require identifying the *negative* instances as well as the positives. A toy model will illustrate the point. Imagine a page divided into squares each of which contains a randomly assigned letter (an 'A' or a 'V' or a 'W' or whatever). Suppose you want to find all the Xs on that page. One way to do so is to identify each letter, rejecting everything but X-tokens. This search can be carried out either serially or in parallel; if the latter, there need be no correlation between the search time and the number of letters in the array. However, if you do find the Xs this way, you ought in passing to acquire lots of 'incidental' information about the negative items. In fact, all else equal, you should be able to identify at least some of them ('Were there any Ls on the page?'; 'Yes, I think so.'). Compare an unconceptualized parallel search. Take a transparency the same size as the stimulus array and ruled into boxes in the same way. Put an 'X' in each box. Place the transparency over the stimulus page and read off every identifiable letter. Except for accidents (where the shape produced by laying an X on a letter is itself a letter), all and only the items you can read off are Xs. In this kind of parallel search the identities of the negative instances aren't ascertained; in consequence, the subject can't tell you what negative items there were. ('Were there any Ls?' 'I haven't a clue. ') The moral, then, isn't that the item-effect is a litmus for iconic representation; rather



it's that there is a galaxy of related indicators of iconicity, and there's no principled reason why, in a given case, they might not decide the issue beyond a reasonable empirical doubt.

*Some data at last*

Can we find, in the perceptual psychology literature, indications of a mode of representation that exhibits typical effects of iconicity? If we can, then it's in the cards that such representations are unconceptualized, hence that there is a perceptual given.

In fact, relevant examples are the stock in trade of intro-level cog sci texts. The basic idea is that perceptual information undergoes several sorts of processes (in more or less serial order) in the course of its progress from representation on the surface of a transducer (eg on the retina) to representation in long term memory. Some of the earliest of these processes operate on representations that are stored in an 'Echoic' Buffer (EB)<sup>23</sup> and these representations are widely believed to be iconic.

Two consequences of the presumed iconicity of echoic representations should be stressed, since both suggest possible experimental investigations. First, since iconic representations are unconceptualized, they do not individuate items in the stimulus domains; so representations in EB oughtn't produce item effects. Second, qua unconceptualized, iconic representations can't express properties whose recognition requires perceptual inferences; inference is about *truth* preservation, so only conceptualized representations can be premises or conclusions. So, in the case of vision, icons register<sup>24</sup> 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*); Correspondingly, in the case of auditory perception, icons should specify the sorts of

---

<sup>23</sup> Not to be confused with 'Short Term Memory' (which is supposed to be conceptualized hence item limited except when rehearsal is allowed). It's STM, rather than the buffer, to which George Miller's famous 'seven items plus or minus two' is supposed to apply (see note xxx). The problem is partly that there's no settled terminology in the psychological literature and one just has to muddle along. But, more important, there's a substantive empirical issue about whether EB and STM really are distinct psychological mechanisms. I don't suppose this issue actually to be settled; but, for the present expository purposes, I shall assume that the 'Aye's have it.

<sup>24</sup> I'm taking it as a terminological point that 'registering' a property doesn't require a symbol (mental or otherwise) that expresses it. So, suppose you can have an experience *as of red* without having the concept RED (or THAT SHADE OF RED), and suppose that RTM is right that the content of one's experience must ipso facto be mentally represented somehow or other. I'll reserve 'register' for that kind of unconceptualized representation. (Likewise, mutatis mutandis, for the relation between a photo and the shape, color, etc of what it's a photo of.)

properties that show in a spectrogram (frequency, amplitude, duration), but not (eg.) whether the distal sound is a rendition of ‘Lilibularo’).<sup>25</sup>

Let’s start with an anecdote just by way of building intuitions. So: here I am, seated at the keyboard, working hard on a piece for *Mind And Language* (or whatever); at the moment, I vacillate between a semicolon and a comma. A clock begins to chime. ‘Chime, Chime, Chime’ the clock says. At first I ignore this, but then it seizes my attention ‘I wonder what it may be o’clock,’ I say to myself (it being my habit to address myself in a sort of Pig-Georgian ). What happens next is the point of interest: I commence to count the chimes, *including the ones that I hadn’t previously noticed*. Strikingly (so, anyhow, the phenomenology goes) it’s not that I just say to my self ‘there have been three chimes so far’; rather it’s that I *count* the unattended chimes: ‘one chime, two chimes, three chimes’ I say to myself, thereby subsuming each chime under the sortal concept A CHIME. Four more chimes follow and I duly add them to the total. And then I think *it’s half passed six* (the hall clock runs half an hour fast.) Notice that one’s ability to do this trick is time bound; it lasts only for perhaps a second or two, so you can’t count the unattended chimes that you heard yesterday. A psychologist would conclude: *There’s a brief interval during which an iconic (hence unconceptualized) representation of the chiming is held in EB. Within this interval, you can conceptualize (hence individuate, hence count) the chimes more or less at will After that the trace decays and you’ve lost your chance*. I think he’d be right to conclude all that.

*Prima facie objection*: But clearly there is an item limit on the buffer. You may be able to count two or three chimes retrospectively, but I’ll bet you can’t do seventeen.

*First reply (in passing)*: Temporal effects can mimic items effects so they must be controlled for. Suppose representations in EB last 2 seconds and it takes the clock 30 seconds to chime 17 times. You will ‘lose’ the last  $n$  chimes in such a sequence. This is not, however an effect of the number of stimulus items that can be stored in EB; it’s just an interaction between the temporal duration of the input and the temporal capacities of the buffer.

---

<sup>25</sup> This is to be sure, the entrance to a minefield; it is, as everybody knows, very hard to give a general principle for distinguishing inferential/perceptual properties from sensory/transducible properties. I’m pretty sure that the issue is empirical (that again!) and hence must be decided post facto. For example, it would be reasonable (and uncircular) to identify the sensory properties of a stimulus with the ones that are registered by the iconic representations that it evokes. For, see Fodor and Pylyshyn 19xx, Antony and Levine 19xx, Fodor, 19xx.

*Second reply (more interesting):* It's not because the buffer is item limited that you can't count up to seventeen retrospectively. Rather (once you control for temporal decay) it seems that how much you can put in the buffer is relatively unconstrained. It's not *representing* many chimes that's hard, it's *counting* them. For: counting requires individuation; and individuation requires conceptualization, and it's independently plausible that conceptualizing costs.<sup>26</sup>

There actually are data that suggest that this is the right diagnosis. They come from a deservedly famous series of experiments by George Sperling (1960). The findings are richer than I have space to summarize, but they support a pervasive phenomenological intuition: “when complex stimuli consisting of a number of letters are tachistoscopically presented, observers enigmatically insist that they have seen more than they can remember afterwards, that is, [more than they] report afterwards.(1)”;<sup>27</sup> in the experiment, “the observer behaves as though the physical stimulus were still present when it is not (that is, after it has been removed) and... his behavior in the absence of the stimulus remains a function of the same variables of visual stimulation as it is in its presence (21)”. That is, there appears to be a kind of very short term memory for visual stimuli on which the item bound is, at a minimum, considerably greater than what is available for short-term report;<sup>28</sup> and it appears that the representational format in this memory is some sort of visual icon.<sup>29</sup>

---

<sup>26</sup> Or, perhaps, it's 'reading out' (reporting) the conceptualized representation that costs. Or perhaps both do. The issue is once again empirical, but any of these interpretations will do for the purposes at hand.

<sup>27</sup> In the basic 'partial report' paradigm, S receives a brief visual exposure to a matrix of letters (and/or numbers). At a controlled interval after this presentation S is cued as to the location of the items to be reported ('top row,' 'middle row,' etc.) In general, S is able to report any of the cued item from a matrix of at least 12 stimuli. This is considerably larger than S's capacity to report 'each of the items he can remember.'

It's of special interest, given our present concerns, that the partial report effect is *not* found when the items to be recalled are cued by category rather than location ('report the numbers but ignore the letters'.) This strongly suggests that the iconic representation is indeed preconceptual since, of course; a photograph can't but register the (relative, two dimensional) location of the objects photographed. By contrast, representing the stimulus as a letter or a numeral requires a perceptual inference from (eg) their shapes.

<sup>28</sup> Sperling is very conservative in estimating how much information a visual icon can contain (compare the indeterminacy of the question how much information an (undigitalized) photograph contains; see text above). But he does remark that "... it seems probable that the 40-bit information capacity observed in these experiments was limited by the small amount of information in the stimuli rather than by a capacity of the observers. (27)''

<sup>29</sup> For analogous studies with auditory phenomena, see the psychological literature on the 'suffix' effect. It seems that there is a very brief period just after the presentation of a list of auditory stimuli in which the interference of the items in a recall task is affected primarily by the *acoustic* similarities among them. Presentation of an acoustically similar item interferences with list recall *even if the item is redundant*; i.e. even if the subject knows antecedently what the item will be. (Here too, however, there is a rather less than perfect

The Sperling results argue that the content of representations in EB is unconceptualized.<sup>30</sup> But, of course, the representations in EB must *have content*; they must carry information about the stimulus. That they do is crucial to explaining why the subject in the ‘partial report’ condition is accurate more often than chance (see fn. 27).

So, if ‘a given’ is what’s unconceptualized but nonetheless semantically evaluable, then it’s thus far plausible that the representation of experience in EB qualifies as given. But I do want to emphasize the ‘thus far’ part. The argument just set out is (at best) empirical; it suggests that there is iconic representation in perception, but it certainly doesn’t demonstrate that there is. A demonstration would be ever so much nicer; their level of confidence is so high. But there isn’t one either pro or con. Nor will there be.

I also want to emphasize that Sperling’s study, though particularly elegant, is only one of a plethora of straws in the wind. In particular, effects of content without items effects are quite easy to find when you know where to look. I’ll mention one that makes the dramatically.

Bella Julesz and his colleagues studied the perception of computer generated displays of matched pairs of visual stimuli, each of which consists of an array of many randomly positioned dots (for a review, see Julesz, 1971). The two arrays in a pair look identical to casual inspection; but, in fact, the location of some of the dots is slightly shifted from one to the other. Under conditions of stereoscopic presentation (one member of a pair is presented to each eye), such stimuli produce a powerful illusion of three dimensionality. The area containing the displaced dots appears to emerge from a shared background.

From our point of view, several considerations are germane. First, the displacement of the dots must somehow be registered by the subject’s sensory representation of the stimulus. After all, the sensory representation is the only information about the stimulus that’s available to affect what the subject sees;<sup>31</sup> so, if it failed to preserve the information

---

consensus on what the data are or what they mean. Those of you who follow the psychological literature will find this unsurprising.) For a review and some skeptical discussion, see Neath et al (1993).

<sup>30</sup> With, however, the usual caveat: there is an extensive and acrimonious literature as to how the partial report effect ought to be interpreted. In particular, there’s disagreement about just how strongly they suggest the mental representations involved in the effect are icons. For relatively recent discussion, see Haber, 1983; Chow, 1986, .

<sup>31</sup> For example, the subject has no relevant ‘background beliefs’ about the stimuli of the kind that a ‘top down’ account of the depth effect might appeal to. This was, in fact, the main reason for Julesz’ interest in experiments with random dot stimuli.

that some of the dots have been displaced, there could be no illusion of stereopsis.<sup>32</sup>

Moreover, the information so preserved must be accessible to the mechanisms of perceptual analysis. In effect, the visual system must compare a representation of the left-eye stimulus array with a representation of the right-eye array in order somehow to determine which dots have moved.

But there is every reason to doubt that the representations that are used to make this comparison could be conceptualized (ie that what's in EB represents each displaced dot *as* a displaced dot.) It is out of the question, for example, that displacement detection has access to a list of the dots with their positions in each array. Since the depth illusion is instantaneous and can be produced by stimulus arrays with thousands of dots, the amount of information that would need to be registered and processed to make the relevant estimates would be orders of magnitude too large to be feasible. And ---a more important consideration in the present context--- if detecting the dot displacements required representing each dot and its position, one would predict a consequent item effect. But, in fact, there isn't one; it's not the case that the more dots there are, the harder it is to obtain the illusion. What happens is apparently that iconic, unsegmented impressions from each of the two eyes are superimposed somewhere in the visual system.<sup>33</sup> Whenever two dots fail to 'line up,' one of them has been displaced.

So, it would seem that stereo depth perception exhibits an effect of content without an items effect; so (plausibly) there's an effect of unconceptualized content on stereo depth perception; so there is quite likely a perceptual given. QED.

*Conclusion:*

I think there is quite likely a perceptual given. In any case, it would seem that the issue is empirical; whether there's a given is no philosopher's business. On the other hand, if there is given, that should be of professional concern to philosophers who argue a priori that there

---

<sup>32</sup> In fact, information carried by the sensory representation must also register the *magnitude* of the displacement, since how far the dots are moved affects the strength of the illusion

<sup>33</sup> That's why you can get a depth illusion from such stimuli without using a stereoscope if you learn to cross your eyes just right

can't be because content has to be conceptualized. This philosopher is now required to sketch an alternative explanation of the sorts of empirical findings I've been illustrating.<sup>34</sup> I am not holding my breath.

But does it matter philosophically in any other way? Does it, in particular, matter to epistemology? I have one very brief remark:

I don't see that the epistemology of perception can simply ignore the empirical question how perception works. Quite generally, justifying a belief can't require a thinker to do such-and-such unless the thinker has the kind of mind that *can* do such-and-such. If, for example, my only justified beliefs are the ones of which I am able actually to cite the justifications, then I have very few justified perceptual beliefs; possibly none.<sup>35</sup> I've heard it said that doesn't matter because epistemology is not a descriptive but a normative discipline. But how could one be bound by norms that one is, in point of nomological necessity, unable to satisfy? And what's the interest of norms that don't bind us?

---

<sup>34</sup> Or he may argue that philosophers are ipso facto licensed to ignore empirical findings, psychological or otherwise. I wonder who issued the license.

<sup>35</sup> Perhaps there's at least this: In justifying one's perceptual claim that P, one can always cite its seeming (its having seemed) to one that P. This doctrine, though venerable, strikes me as confused; in particular, as confusing *offering a justification for making a claim* with *offering a justification for the claim*. Compare belief: that I sincerely believe that P generally justifies my claiming that P; but it's not a reason to believe that P is true (or, anyhow, it's not much of one. Surely it can't be *my* reason for believing that P.) Why suppose that the epistemology of perception differs, in this respect, from the epistemology of belief?

