

6

Preconceptual Representation

6.1 Introduction

It's been a working assumption throughout this book that seeing a thing (to say nothing of thinking about one) requires representing it; hence that *seeing* and *thinking* inherit the constraints on *representing*. But (so far at least) nothing precludes the possibility that some of the *representing* that goes on in seeing/thinking is nonconceptual. Well, then are there unconceptualized mental representations? If so, what might they be like? And where might one look for some? And how would you know if you'd found one?

SNARK. And who cares?

AUTHOR. Well, epistemologists care. For example, it's often said (at least since Sellars 1956) that nothing preconceptual is 'given' in experience; there is no such thing as acquaintance with the world 'as it is in itself', in abstraction from the effects of its being conceptualized; in effect, there is no principled distinction between the perceptual realm and the realm of thought.¹ That being so, what is given in experience can't be what justifies perceptual beliefs. This line of argument is sometimes described as 'neo-Kantian'. For all I know, it may be.

¹ It isn't really clear to me that Sellars did endorse this sort of thing. But lots of philosophers who take themselves to be Sellarsians do, which is good enough for present purposes (see e.g. McDowell 1994).

SNARK. Why should I care whether epistemologists care whether there is an experiential given? It's one thing to deny that anything in the content of experience shows us how the world is when it's swallowed raw. It's quite another thing to deny that there is anything in the content of experience that isn't conceptualized. I am not an epistemological Snark, I am a psychological Snark. I am therefore content to wonder whether there is preconceptual experiential content and let the epistemological chips fall where they may. Even if what is experientially given isn't what *justifies* perceptual beliefs, it might be what *gets conceptualized* in the course of the formation of perceptual beliefs.

AUTHOR. Good for you, Snark; it is indeed Very Bad Practice to run your psychology with epistemological malice aforethought. People who do so generally make the worst of both. Still, questions about justification aside, a psychologist might reasonably wonder what role experience plays in belief formation. And the question what (if any) preconceptual content experience contains would surely be crucial in that inquiry. Let us, therefore, have a look at it.

SNARK. I think it's time for lunch.

AUTHOR. Later.

6.2 Kinds of representations

For reasons I'll presently set out, I think that at least some of the mental representations that are causally implicated in the formation of perceptual beliefs are indeed nonconceptual. The line of argument I'll have on offer goes like this: On the one hand, it's (empirically) plausible that at least some of these representations are 'iconic' (rather than 'discursive'); and, on the other hand, it's in the nature of iconic representations not to be conceptual. That being

my polemical strategy, I had better say something about what I take iconic and discursive representations to be.

To begin with, my usage is idiosyncratic. In the semantics/semiotics literature ‘iconic’ frequently comports with notions like ‘pictorial’ and ‘continuous’. But it’s not always clear just what either of those amounts to, or just what the connections between them are supposed to be. As often as not, they simply take in one another’s wash. I’m going to pretend that the slate is blank and just stipulate. So, then, ‘iconic’ and ‘discursive’ denote 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. Offhand, I can’t think of a good candidate, but it doesn’t matter for my present purposes.) I further assume, for the familiar reasons (see Ch. 2), that all the kinds of representations we’re concerned with 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 semantic content of its constituent parts. According to my usage, the distinction between iconic and discursive representations turns on a difference between the ways that they achieve their compositionality.

6.3 Discursive representation

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 phonological features, a bundle of syntactic features, and a bundle of semantic features.² These are enumerated by the word’s ‘entry’ in the ‘lexicon’ of the language. A discursive

² Except that I don’t really believe in semantic features (see Fodor 1998).

representation in L is syntactically compositional iff its syntactic analysis is exhaustively determined by the grammar of L together with the syntactic analyses of its lexical primitives. A discursive representation is semantically compositional in L iff its semantic interpretation is exhaustively determined by its syntax together with the semantic interpretations of its lexical primitives.³ Consider, for example sentence (1) below. Its syntactic structure is (more or less) as shown in (2), and its semantic interpretation is (more or less) *John loves Mary*. The syntax and semantics

(1) John loves Mary.

(2) (John_{NP}) (loves_v) (Mary_{NP})_{VP}

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.

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. 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’.⁴ 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’;

³ It helps with the exposition to identify the semantic interpretation of a sentence with its truth-conditions, as I will generally do. This is, of course, wildly tendentious, but I don’t propose to trade on it. It is, 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 the other, interpretations qua representations that speakers/hearers of L assign to tokens of its sentences in the course of communication exchanges. Which is which will be apparent from context wherever it matters. I hope.

⁴ Lexical primitives themselves count as constituents because they have semantic interpretations. Whole sentences are constituents by courtesy.

in fact, neither of them *has* a semantic interpretation in that sentence (though, of course, each of the lexical primitives that they contain does).

I'll say that 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.

6.4 Iconic representation

Pictures are the paradigms (but see 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.⁵

Pictures and the like differ from sentences and the like in that icons don't have *canonical* decompositions into parts; *all* the parts of an icon are *ipso facto* constituents. Take a picture of a person, cut it into parts whichever way you like; still, each picture part pictures a person part.⁶ And the whole that you have if you reassemble all the picture's parts is a picture of the whole person that the parts of the picture are pictures of.

So, then, in everything that follows, a representation that has no canonical decomposition is an icon. I will argue (quite soon now) that iconic representations *ipso facto* lack a number of the characteristic features of conceptualized representations; so, the question we started with ('Are any mental representations unconceptualized?') can be swapped for the question 'Are any mental representations

⁵ I assume, for ease of exposition, that P is of X is the only semantic relation that holds between pictures and things in the world that they are pictures of; a fortiori, that it is the only such relation that is compositional. I don't think it matters to my line of argument whether this is so. I likewise leave it open that some pictures aren't 'of' anything.

⁶ Some parts of an icon are too small to have interpretations; the atoms it's composed of, for example. I'll generally ignore such matters of 'grain', but see below in the text.

iconic?’ And that, finally, is a question on which empirical evidence can be brought to bear. First, however, I digress for caveats.

6.5 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 that curve *ipso facto* represents the distribution of the property in a part of that population (which is to say that it partially represents the distribution of that property in the population). So, graphs are (typically) icons according to my usage. Since icons don’t have to be pictures, it can’t be assumed that icons *ipso facto* resemble what they represent; not even on the (mistaken) assumption that pictures do.

6.6 How icons work

So far: Iconic representations, like discursive representations, are typically *of* this or that. But the former have no canonical decomposition; which is to say that 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 is a constituent, and each constituent gets a semantic interpretation in accordance with the Picture Principle. But neither is true of discursive representations. Only a specifiable subset of the parts of a discursive symbol (namely, its canonical parts) are syntactic or semantic constituents; and it is 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, are clearly uniconic in both these respects. 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. 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 discompose into syntactically and semantically heterogeneous constituents, discursive representations can have logical forms (maybe all discursive representations do 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 truistic. The logical form of a symbol is supposed to make its compositional structure explicit; that is, it's supposed to make explicit the contribution that each of the interpreted parts of the symbol makes to its interpretation. But each part of an iconic symbol is a constituent, and each constituent contributes in the same way to the interpretation of the whole icon: each part pictures part of what the icon pictures.

Because the interpreted parts of an iconic representation are in this sense syntactically and semantically homogeneous, iconic symbols can't represent things that discursive symbols can. For example, icons can't express the distinction between negative propositions and affirmative ones, which turns (*inter alia*) on semantic distinctions among logical constants. Likewise they can't express quantified propositions; or hypothetical propositions; or modal propositions. They can't even express predication, since doing that requires (*inter alia*) distinguishing terms that contribute individuals to semantic interpretations from terms that contribute sets (or properties, or whatever). 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 nor does it tell the truth.⁷

It's implicit in these sorts of considerations that discursive representations typically carry ontological commitments, but iconic representations don't.⁸ In particular, discursive representations do, but iconic representations don't, impose principles of individuation on the domains in which they are interpreted.⁹ 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 census 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 representational system doesn't include quantifiers (or classifiers or something of the sort), there is no right answer to the question 'Which things (how many things?) does this symbol represent?'. (Didn't Quine say something of that sort? I hope he did; I would so like to be in respectable company for a change.) To be sure, a photograph may show three

⁷ This is continuous with the familiar objection to the 'picture theory' of ideas: There is nothing in John's not loving Mary for a picture to resemble.

⁸ It is notoriously moot whether, and under what conditions, the ontological commitments of a system of discursive representations might be unique (for example, whether there's a fact of the matter about what the representations refer to). As far as I can tell, the usual arguments for indeterminacy assume that the data for interpretation(/translation) are exhausted by correlations between the informant's utterances and the situations in which he utters them. The claim is that there is, even in principle, nothing in that sort of data that could distinguish (e.g.) commitment to an ontology of rabbits from commitment to an ontology of undetached rabbit parts. Since, however, I think the correlationist premise is surely false, it strikes me as not terribly important whether the conclusion follows. For discussion see Fodor (1994).

⁹ However, it's of prime importance not to read 'principle of individuation' epistemically. If I have the concept CAT, I know the principle of individuation for cats; namely, that each cat is one cat, and nothing else is any. So, if I have the concept CAT, then I know how to count cats; that is, I know that I must count one for each cat and nothing for anything else. But what does not follow is that if I have the concept CAT I am thereby enabled to count 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 computing the cardinality of a flock).

giraffes in the veld; but it likewise shows a family (*one* 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 veld that's inhabited by any or all of these. No doubt, we can usually agree about how to interpret the ontology of such a photograph; we do so in light of whatever project we happen to have in hand. But what matters to the present issues is that the discursive symbol 'three giraffes in the veld' specifies a scene relative to such concepts as THREE, GIRAFFES, IN, and THE VELD. A fortiori, a mind that lacks these concepts can't read the symbol as representing three giraffes in the veld (or as representing one family of giraffes in the veld, or as representing an odd number of giraffes in the veld, etc.). Contrast iconic representation: you can, of course, see three giraffes in the veld without having the concept GIRAFFE. Nor do you need GIRAFFE to take a picture of three giraffes in the veld; a camera and film will suffice. (Less these days.)

Equivalently (more or less): the context 'iconally represents ...' is like the contexts 'sees ...', '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 things as C and describing things as C, like other acts of conceptualization, operate by subsuming the things seen under the concept C. It's thus entirely in the spirit of RTM that 'conceptualizing X as C' and 'predicating C of X' are two ways of saying much the same thing.¹⁰

6.7 Brief review

We started with conceptualized versus unconceptualized representation, which we then swapped for discursive versus iconic representation. This allowed us to reformulate the question whether

¹⁰ Except, of course, that 'C' is used in the first but only mentioned in the second.

there are unconceptualized mental representations as the question whether any mental representations are iconic. We then suggested that (because they lack logical form) iconic representations don't provide principles of individuation for their domains of interpretation. This led to a final metamorphosis: 'Are there unconceptualized mental representations?' becomes 'Are there (empirical) 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 can look like. But, first, a quick consideration of an *a priori* objection that many philosophers appear to find persuasive; indeed decisive.

'Look', you might say (or, anyhow, Snark might), there just couldn't be phenomena in which representation and individuation are dissociated. For, if a symbol represents a such and such, it must represent it *as* something or other. 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 reading of 'represents...'; that's to say, by ignoring the mode of representation in a specification of a representational content. There's a close analogy to two ways that one might read an assertion that John believes the king of France is bored: According to the transparent reading, it's the ascriber who takes responsibility for the definite description; according to the opaque reading it's John who does. (See e.g. Brandom 2000; Dennett 1982.) That, however, is a distinction between styles of belief ascription rather than a distinction between kinds of representing (or, *a fortiori*, between two kinds of believing). Likewise for the presumptive difference between *representing* and *representing as*. Representing a thing is always representing it as this-or-that. But one has the option, in saying that something is represented, of not bothering to say what it's represented *as*. So, iconic and discursive representations couldn't differ in the way that this discussion has supposed; the difference couldn't be that the

latter but not the former conceptualizes, and hence individuates, whatever it is that it represents.¹¹

That sort of objection has force on the assumption that all representing implies representing as. But assuming that it does would beg precisely the issue we are discussing; namely, the possibility that some mental representation is nonconceptual. That representing requires conceptualization was supposed to be the conclusion of the argument, not its premise; so we can't take for granted an analysis of either that postulates an inalienable connection between the two. Suffice it for present purposes, then, if we can imagine, even roughly, how X might represent Y without representing Y as falling under some concept or other.¹² That would be representation without individuation according to the present line of thought.

6.8 Informational content

I think, in fact, that there's quite a plausible candidate for representation of that kind: X represents Y insofar as X carries information about Y, where 'carries information about...' is read as transparent. If so, then maybe construing what's given as some sort of information would allow for representation that's not 'under a description' and hence for unconceptualized representing. In particular, the claim would be that tokens of symbols typically

¹¹ If in the case of iconic symbols resembling is sufficient for representing (and if resembling isn't itself description-relative) that would explain how there could be representation without individuation. But I don't believe that preconceptual representations resemble what they represent; nor, I suppose, does anyone else who is party to this discussion. Alternatively, demonstrations ('this' and 'that') might be instances of representation without conceptualization, but the claim that they are is highly tendentious.

¹² SNARK. What on earth are you talking about?

AUTHOR. The thesis up for discussion is that representation per se needn't presuppose principles for the individuation of what is represented; roughly, discursive representation does but iconic representation doesn't. So, the polemical situation is that if you want to argue against the thesis, you mustn't take for granted that representations are *ipso facto* conceptualized; that would beg the question at issue. Does that help?

SNARK. Did you really think it would?

carry all sorts of what I'll call 'Dretskeian' information. (Come to think of it, so too does practically everything else.) And 'carries information about...', unlike 'represents as...' is transparent to the substitution of coextensives at the '...' position. On the other hand, what information an interpreter can recover from a tokening of a symbol depends on what concepts the interpreter has available. Even though a picture carries information about giraffes, only an interpreter who has the concept GIRAFFE can see that it does (i.e. can see it as doing so).¹³ What information a representation contains doesn't depend on how it is interpreted. But what information you can recover from a representation depends (not just on what information it contains, but also) on what concepts you bring to it. The given is unconceptualized representation that is awaiting conceptualization. This differs very much from the traditional empiricist account according to which the concept X is acquired from experiences as-of-Xs (presumably by some sort of process of abstraction). For reasons that connect with the discussion in Chapter 5, that way of thinking about concept acquisition seems to me to be inevitably circular, so it can't possibly be right.

In any event, it can't be right if, as I'm now suggesting, experiential content is the same thing as informational content. You can't abstract DOMESTIC FELINE from experiences-as-of-domestic-felines unless you *take them to be* experiences-as-of-domestic-felines; and you can't do that unless you have DOMESTIC FELINE in

¹³ Information-construed-as-causal-correlation was introduced into the philosophy of language in Dretske (1981). I think that was a major advance. But, unlike Dretske, I don't suppose 'carries information' is the basic notion in this area; in particular, I doubt that concept possession can be understood in terms of it. Rather, the analysis ought to run the other way around: What concepts are available to the interpreter sets an upper bound on what information he may be able to recover from the tokening. What's crucial is that if a tokening of a representation carries information about F, and only if it does, an interpreter who has the concept F may be able to recover that information. ('He may be' not 'he must be'; since all forms of pragmatism are to be shunned, it mustn't be assumed that to understand a representation type implies being able to interpret its tokens: see Ch. 2.) What's crucial is that if a representation carries information about F, then an interpreter who has the concept F may be able to recover that information from the representation. But if he doesn't, then he can't.

your conceptual repertoire. Taking-X-to-be F and applying-the-concept F to X are two ways of saying the same thing.

I'm endorsing a quite different way of thinking about the relation between concepts and experiences.

The idea is that the role of concepts in the perceptual analysis of experience is to recover from experience information that it contains. (I think this is a Kantian kind of idea. I think Kant would put it that the function of concepts in the perceptual analysis of an experience is to afford a 'rule for the synthesis' of the experience. But I may be wrong to think this. What do I know about Kant?) At very least, this idea is easy to illustrate. What's this a picture of?

Give up? It's a picture of a bear climbing a tree (as seen from the other side of the tree).¹⁴ If you have the concept BEAR (and the concept TREE and so forth) you can, as it were, pull the picture together under this description. But if you don't, you can't. Clearly, you couldn't 'abstract' the concept BEAR from encounters with this sort of bear-picture (or, *mutatis mutandis*, from encounters with the corresponding bear-experience).

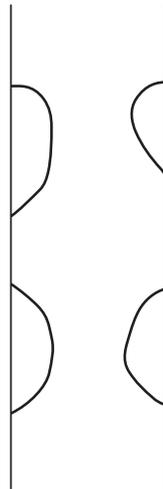


Figure 6.1 *What's this?*

¹⁴ I'm indebted to Lila Gleitman for suggesting this very un-Kantian way to illustrate (what I'm taking to be) a Kantian point. There aren't, if memory serves me, many jokes in Kant.

Well, so there are some interesting implications of thinking about experiential content as informational content, quite aside from the one that's of most present concern: that the notion of carrying information about X seems to offer a way of representing X without representing it as anything; hence a way of meeting what I took to be the Dennett/Brandom objection to a preconceptual given. But I haven't got a worked-out elaboration of this kind of treatment, nor can I prove that one is possible. Whether it is is another of those questions that scare me. I proceed, therefore, to conditionalize: Assuming that an informational construal of the given allows for representation *of* that isn't representation *as* (hence for representation without conceptualization, hence for representation without individuation), the empirical issue is whether there is empirical evidence for such dissociations in the psychology of cognition. I think there's lots of it, and of several kinds. Selected examples to follow.

6.9 The 'items effect'

To test a theory, you need (what used to be called) 'correlating definitions'. Here's one: all else equal, the 'psychological complexity' of a discursive representation (for example, the amount of memory it takes to store it or to process it) is a function of the number of individuals whose properties it independently specifies.¹⁵ I shall call this the 'items 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. All sorts of things follow: the phone books of big cities are generally bigger than the phone books of small cities; they weigh more and they take

¹⁵ That is, the individuals and properties that are enumerated (rather than merely quantified over). De facto, 'All men are mortal' predicates mortality of more men than 'Two men are mortal'. But the relative complexity of these representations doesn't differ in consequence.

up more shelf space; 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 (these truths being, of course, purely contingent). This is all because the representations in phone books are discursive, hence conceptualized; accordingly, their interpretation presupposes the possession and application of such concepts as X'S NAME IS 'X' and PHONE P HAS THE NUMBER N. Lists, like sentences, are paradigms of discursive representation: on the one hand, 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, on the other hand, they 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 (since very old ones are generally more degraded than very new ones, token photographs can carry information about their age). But they aren't item sensitive. Photographs of many Xs aren't, in general, more complex than photos of a few Xs; the former don't, for example, have more constituents than the latter. This is hardly surprising in light of the preceding discussion: iconic representations don't individuate; they don't represent individuals as individuals. A fortiori, nothing about them depends on the number of individuals that they represent.

Still, I do want to emphasize that failing to find items effects is not, all by itself, a knock-down test for iconic representation; as usual with 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 is no items effect, one sometimes can't tell whether that's because it's a search of an iconic representation or because it's a parallel search.

On the other hand, the discursive/iconic distinction is orthogonal to the parallel/serial distinction: Some parallel searches involve conceptualizing the items in an array, and others don't, and empirical evidence can distinguish between the two. For example, the former, but not the latter, requires identifying negative instances as well as positive instances.

A toy example will serve to illustrate the principle. Imagine a page divided into squares each of which contains a letter that's assigned randomly (an 'A' or a 'V' or a 'W' or whatever). Suppose you want to find all the Ks on the page. The search can be carried out either serially or in parallel, but if you find the Ks by a serial search, you ought to acquire, in passing, lots of 'incidental' information about the negative items, and (all else equal) you should be able to identify at least some of them. ('Were there any Ms on the page?'. 'Yes, I think so.'). Compare a parallel search. Take a transparency the same size as the stimulus array and ruled into boxes in the same way. Put a 'K' in each box. Place the transparency over the stimulus page and read off every identifiable letter. All and only the items you can read off are Ks. Since the identities of the negative instances aren't registered in this kind of parallel search, the subject can't tell you what negative items there were. ('Were there any Zs?' '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 mightn't decide the issue beyond a reasonable empirical doubt if they are taken together. If (to return to an earlier example) you are counting giraffes in parallel, there's no reason why it shouldn't take you longer to count six of them than to count sixty. But, since it usually does, it's a good bet that you don't usually count giraffes in parallel.

6.10 Some facts 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 on the cards that such representations are unconceptualized, hence that there is a perceptual given. In fact, relevant examples are the stock in trade of introductory-level cognitive-science texts. The basic idea is that perceptual information undergoes several sorts of process (typically in more or less serial order) in the course of its progress from representation on the surface of a transducer (e.g. on the retina) to its representation in long-term memory. Some of the earliest of these processes operate on representations that are stored in an 'echoic buffer' (EB),¹⁶ and these representations are widely believed to be iconic. Two consequences of their presumed iconicity should be stressed, since both suggest possible experimental investigations.

First, since iconic representations are unconceptualized, they don't individuate items in the domains they represent; so representations in EB oughtn't to produce item effects. Second, qua unconceptualized, iconic representations don't express properties whose recognition requires perceptual inference. Inferring is in the same basket as saying and thinking; they all presuppose conceptualization. So, 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

¹⁶ Not to be confused with 'short-term memory' (which is supposed to be conceptualized and hence item limited except when rehearsal is allowed). It's STM, rather than EB, to which George Miller's famous 'seven items plus or minus two' is supposed to apply (Miller 1956). The expository problem is partly that there's no settled terminology in the psychological literature, so one just has to muddle along. But there's a substantive empirical issue about whether EB and STM really are identical psychological mechanisms. I don't suppose this issue actually to be settled; but for the present purposes I shall assume that they're not.

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.

Correspondingly, in the case of auditory perception, icons in the echoic buffer should specify the sorts of properties that show in a spectrogram: the frequency, amplitude, and duration of a sound, but not whether the sound is a rendering of 'Lilibulero'.

You may have noticed that I stuck in 'register' without telling you what it means. Well, roughly, for a representation to register *a*'s being *F* is for it to contain Dretsian information about *a*'s being *X*. In particular, 'registers', like 'contains Dretsian information about' is transparent to the substitution of identicals. If *a* contains Dretsian information about *e*₁, and *e*₁ = *e*₂, then *a* contains information about *e*₂. Likewise, if *a* registers information about *e*₁, and *e*₁ = *e*₂, then *a* registers *e*₂.¹⁷ Also, it's possible to register the Dretsian information that *a* is *F* even if you don't have the concept *F*. (You get a plausible instance of the registration of Dretsian information if you assume that you can have an experience as of red without having the concept RED.)

SNARK. Grrr!

AUTHOR. Be that as it may.

There are two other differences between registering and perceiving. First, a representation token can carry information about anything, but there are constraints on which of its properties can be registered; only 'sensory' or 'transducer detectible' properties can (see above); a mind can register only such properties as it has mechanisms for transducing.

SNARK. And, what, pray, is a transducer?

AUTHOR. I thought you'd ask. Here's one way to think about it:

¹⁷ I persist in my usual serene impartiality as to what the variables range over. For expository purposes, let it be events; but you aren't allowed to ask me what events are.

Computation (like, for example, thinking) takes mental representations onto other mental representations. Transduction (like, for example, registering impinging redness) takes ambient energy onto mental representations. In the usual case (barring hallucinations and such) perceptual integration starts with the registration of sensory information. Were there no transducers, perception couldn't get started. Will that do?

SNARK. No.

AUTHOR. Very well then; I return to the main line of exposition.

The second difference between perception and registration: Qua unconceptualized, registrations can't express properties whose detection requires inferences. (Inferences need premises and conclusions, and premises and conclusions are *ipso facto* conceptualized; which, by assumption, registrations aren't.)

Bearing all that in mind, we may now return to the question whether there is evidence for unconceptualized representation in perception (that is, evidence of the registration of stimulus properties). 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 ones that I hadn't previously noticed*. Strikingly (so, anyhow, the phenomenology goes), it's not that I say to myself 'There have been three chimes so far, therefore what I'm now hearing is chime four'; rather, it's that I count chimes that were previously not attended to: 'One chime, two chimes, three chimes', I say to myself, thereby subsuming each chime under the sortal

concept CHIME. Four more chimes follow and I duly add them to get the total. I think: ‘It must be 6: 30’ (the clock in the hall runs half an hour fast).¹⁸ Given such observations a psychologist might well want to consider, as a working hypothesis, that there is a brief interval during which an unconceptualized (presumably iconic) representation of the chiming is held in the EB. 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 unattended chimes that you heard yesterday. Within the critical 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 the psychologist might well be right to conclude all that as a working hypothesis.

Prima facie objection: ‘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 three seconds and it takes the clock four seconds to chime four times. You will then “lose” the first chime in the sequence before you register the fourth. This is not, however, an effect of the number of stimulus items that can be registered in EB; it’s just an interaction between the temporal duration of the input and the temporal capacities of the buffer’. Second reply (more interesting): ‘It’s not because the buffer is item-limited that you can’t count seventeen chimes; it’s because counting involves individuation and individuation requires conceptualization, and it’s independently plausible that conceptualizing is expensive’.

There actually are data that suggest that the second is the right diagnosis. They come from a deservedly famous series of experiments by George Sperling (1960). Sperling’s findings are richer

¹⁸ I’m pleased to report that there is (anecdotal) evidence that mine is not the only head that works this way: ‘[Molly] found it hard work to attend to kind Miss Phoebe’s ceaseless patter. She came to a point, however, when the voice ceased; and could recall, in a mechanical manner, the echo of the last words which ... from the dying accent that lingered in Molly’s ear, she perceived to be a question’ (Elizabeth Gaskell, *Wives and Daughters*, 1866).

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 can] report afterwards’ (p. 1). 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’ (p. 2). Performance in the experiment showed that though, when queried after the stimulus is turned off, S can report only 3 of the letters he’s seen, *he can report any 3*. So there appears to be a kind of memory, available for the brief registration of visual stimuli, the item capacity of which is, at a minimum, considerably bigger than what is available for the subject to report. Notice that the items in this memory must have some or other sort of content; that’s required to explain why S’s report is accurate more often than chance.¹⁹

One more observation about Sperling’s results is especially interesting given that we are committed to representations in EB not being conceptualized: Sperling’s ‘partial report’ effect is *not* found when the items to be recalled are cued by category (‘Report the numbers but ignore the letters’). This strongly suggests that representation in EB is indeed preconceptual. You can only report an ‘A’ token as a letter token if you have categorized it as a letter token. So, if the ‘given’ is what has content but isn’t conceptualized, it’s thus far plausible that the iconic representations in EB qualify as given. But I do want to emphasize the ‘thus far’ part. The argument just set out is empirical; it suggests that there is iconic representation in perception, but it certainly doesn’t demonstrate

¹⁹ In the basic ‘partial report’ experiment S receives a brief visual exposure to an alphanumeric matrix. After a controlled interval S is cued as to the location of the items to be reported (‘top row’, ‘middle row’, etc). On average S is able to report any 3 of the cued items from a matrix of at least 12 stimuli. 12 stimuli is considerably more than S is able to report when asked to recall all of the items he can remember.

that there is. Demonstrations are ever so much nicer than suggestions; their level of confidence is so much higher. But there isn't one, either pro or con the given, nor will there be. There is, I think, less apriority in heaven and earth (or anywhere else) than philosophers have dreamed of. On the other hand, I also want to emphasize that Sperling's study, though particularly elegant, is only one of a plethora of straws in the wind. Effects of content without items effects are actually easy to find when you know where to look.²⁰ I'll mention one other example because it makes the point dramatically.

Béla 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 dots, which were identical except that some of the dots on one are slightly shifted from their location on 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. In particular, the subject has no relevant background beliefs about the stimuli of the kind that a 'top-down' account of the depth illusion might appeal to. (Testing top-down accounts of depth perception was, in fact, Julesz's main interest in experimenting with random dot stimuli.) So, if S's experience failed to preserve the information that some of the dots have been displaced, there could be no illusion of stereopsis. To effect the illusion, the visual system must compare a

²⁰ Sperling is very conservative in estimating how much information a visual icon can contain. 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 (1960: 27).'

representation of the left-eye stimulus array with a representation of the right-eye array and somehow determine which dots have been moved. These to-ings and fro-ings are, it goes without saying, entirely subpersonal.

I won't bother you with the received account of how the Julesz illusion works; it's complicated, and I only think I understand it on alternate Tuesdays. What's important is that the mechanism that detects the displacement couldn't possibly have access to a list of the dots with their positions in each array. Once again, the lack of an items effect is a relevant consideration. You can get the stereo illusion with arrays of thousands of dots. The amount of information that would need to be registered and processed to make the estimates of displacement by comparing such lists would thus be orders of magnitude too large to be feasible. (In fact, within a large range, the fewer dots there are, the *harder* it is to obtain the illusion; reducing the number of dots below a certain threshold makes the effect go away.) What happens is apparently that the displacement is computed over iconic representations from each of the two eyes. It is presumably because these representations are iconic rather than discursive that the optical relation between the retinal images is critical in producing the illusion.²¹

There are lots of other experimental results in the literature that point to much the same kind of conclusions I've been drawing here; but perhaps the ones I've cited will do to give a sense of the thing.

6.11 Conclusion

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

²¹ You can get the stereoptic effect from Julesz patterns if you look at them with your eyes crossed just right. Why on earth should that be so if what is being compared are lists?

a priori that there can't be a given because it's a priori that all content must be conceptualized. Such philosophers are now required to sketch alternative explanations of the sorts of empirical results towards which I've been gesturing. I am not holding my breath.

But does whether there is a given matter philosophically in any other way? Does it, in particular, matter to epistemology? I offer two epistemological reflections.

6.10.1 First epistemological reflection

If what is given is supposed to ground perceptual judgements, as in foundationalist versions of epistemology, then it must be both noninferential and introspectible; the former in order to avoid a regress of grounding, the latter in order that the content of one's experience should be available for the justification of one's perceptual beliefs. But the empirical evidence is very strong that there is no psychologically real level of mental representation that meets both these conditions. Rather, it appears that what can be introspected is always the product of subpersonal and encapsulated inferences and, conversely, what is a plausible candidate for being uninferred (e.g. the representations in EB) is almost never available to introspection. So, for example, it holds without exception as far as I know that the deliverances of perception become accessible to the perceiver only after the operation of the perceptual constancies; hence the elliptical plate that looks round, the 'correction' of perceived color for changes in the ambient light, the failure of retinal size to predict apparent size, and so forth through a very long list of familiar examples. And (I think again without exception) all varieties of CTM treat the perceptual constancies as paradigm examples of the products of subpersonal inferences; that is, they imply that, invariably, mental representations that exhibit the effects of constancy are inferred. It is, in short, empirically implausible that what is given could be what grounds perceptual beliefs. That being so, a foundationalist epistemologist has two options: either to ignore

the psychology of perception or to stop being a foundationalist epistemologist.

I sometimes wonder whether our epistemology has quite caught up with the Freudian revolution in psychology. There is every sort of evidence that a great deal of the reasoning involved in the causal fixation of quotidian perceptual beliefs is unconscious and hence unavailable for report by the reasoner; in particular, either what justifies our perceptual judgments isn't introspectively available or most of our perceptual beliefs aren't justified. I'm unclear that it matters much (from, as one used to say, a world-historical point of view) which of these epistemology chooses.

6.10.2 *Second epistemological reflection*

It's often suggested, especially by philosophers in the Sellars tradition (like Brandom, McDowell, and Davidson in some of his moods) that the given can't be what grounds perceptual judgments because justification is a relation among *contents* and whatever isn't conceptualized *ipso facto* has no content. The horrific consequence is the isolation of the psychology from the epistemology of justification (as McDowell puts it, the separation of the 'realm of reasons' from 'the realm of causes'). In particular, the kind of causal explanations of perceptual judgments that psychologists seek can at best provide 'exculpations where we wanted justifications'. This is notoriously a long question; but I do hate a priori arguments that such and such a kind of discourse can't be naturalized; and 'realm' talk makes my skin crawl. So, I can't resist a couple of points under this head.

The first is that discussions of the justification of perceptual belief simply mustn't take for granted that all the perceptual content that's available to perceptual justification is *ipso facto* conceptualized. I suppose that it would be safe to assume so if *belief* content were the only kind of content there is; since it's plausible that beliefs are *ipso facto* conceptualized. But if, as I've been arguing, there's a plausible case for preconceptual, iconic representation, the truism

that justification is a relation among the contents of representations does not entail that justification is a relation among conceptualized representations.

Accordingly, the epistemological question that needs settling is whether the content of an unconceptualized representation might be the datum that grounds (e.g. makes rational) a perceptual judgment. Well, I'm damned if I see why it can't be. A picture of three giraffes in the veld carries information about there being three giraffes in the veld. (It carries all sorts of other information too, of course. But so what?) Somebody who has the concepts GIRAFFE, THREE, VELD, and so on (and only somebody who does) is to that extent in a position to recover that information from the picture; and his reason for believing that there are three giraffes may well be that the picture shows three of them and that he sees that it does. His reason, notice; not his mere exculpation. As far as I can see, none of that is under threat from the consideration that judgment requires conceptualization.

The question how (for example, by what computational processes) unconceptualized iconic representations might get 'collected under a concept' is, of course, very hard; and the answer is unknown for practically any of the interesting cases. On the way of looking at things I've been trying to persuade you of, it's a large part of what the psychology of perception is about. But, as far as I can see, there's nothing to suggest that it's unanswerable in principle. Pragmatism is, of course, always and everywhere false, so if there are rules for the conceptualization of experience, they can't be a priori (or, anyhow, they can't be semantic). There aren't, for example, any 'criteria' for recognizing giraffes. What I'm after is therefore a version of Kant's view that concepts synthesize what's given in perception that does without adopting Kant's pragmatism (or, indeed, anybody else's).

I end with a short methodological homily. I don't see how the epistemology of perception can simply ignore the empirical question of 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. It can't require him to introspectively access the noninferential, preconceptual grounds of his beliefs unless he has the kind of mind that has introspective access to the noninferential, preconceptual grounds of his beliefs. I've heard it said that how perception actually works doesn't matter to epistemologists because theirs is a normative not a descriptive enterprise. But how could one be bound by norms that one is, in point of nomological necessity, unable to satisfy? And what's the conceivable interest, even to epistemologists, of norms that don't bind us?