

6 *The nonconceptual content of experience*

TIM CRANE

1 *Concepts and perceptual experience*

To what extent do our beliefs about the world affect what we see? Our beliefs certainly affect where we choose to look, but do they affect what we see when we look there?

Some have claimed that people with very different beliefs literally see the world differently. Thus Thomas Kuhn: ‘what a man sees depends both upon what he looks at and also upon what his previous visual–conceptual experience has taught him to see’ (Kuhn 1970, p. 113). This view – call it ‘Perceptual Relativism’ – entails that a scientist and a child may look at a cathode ray tube and, in a sense, the first will see it while the second won’t. The claim is not, of course, that the child’s experience is ‘empty’; but that, unlike the scientist, it does not see the tube *as* a cathode ray tube. One way of supporting this claim is to say that one cannot see something *as* an F unless one has the *concept F*. Since the child plainly lacks the concept of a cathode ray tube, it cannot see it as a cathode ray tube.

Although Perceptual Relativism is hard to believe, this supporting suggestion is not so implausible. After all, when we see (and more generally, perceive) the world, the world is presented to us in a particular way; so how can we see it as being that way unless we have some *idea* or *conception* of the way it is presented?

We need not be committed to a representative theory of perception to think that perceptions in some sense represent the world. We can express this by saying that perceptions have *content*. Now it is a commonplace that the contents of beliefs and the other propositional attitudes involve concepts. The belief that this thing is a cathode ray tube involves, in some sense, the concept *cathode ray tube*. So the line of thought behind Perceptual Relativism may be expressed thus: seeing an F *as an F* is a state with content.

So the content of this perception must involve the concept *F* – for if it did not, why should we say that the experience represents the *F* in question *as* an *F*? Thus Christopher Peacocke:

the representational content [of an experience] is the way the experience presents the world as being, and it can hardly present the world as being that way if the subject is incapable of appreciating what that way is.¹

On the other hand, it is plausible to say that the child and the scientist have something perceptually in common, when (let us suppose) they look at the object from the same position, in the same lighting, with normal eyesight. And this common state need not just be qualitative – not just a matter of their visual sensations, or properties of their visual fields, if there are such things.² Rather, there is surely a sense in which their experiences *present* or *represent* the world in the same sort of way. When the child comes to acquire the concept of a cathode ray tube, there is surely something in common between the ways the world was represented before and after this acquisition. Otherwise, how could it come to recognise, for instance, that *that thing* is a cathode ray tube?

So there are reasons for thinking that the concepts perceivers have do affect what they see, and reasons for thinking they do not. I think we should therefore reject Perceptual Relativism in its most extreme form. We should accept that the experiences of two perceivers in a situation such as the one described will be perceptually similar. If this is granted, then the question is: how should a theory of perception classify the similarities between such experiences?

Let us agree that these similarities should be expressed, in some sense, in terms of the contents of the experiences – in terms of how the experiences represent the world to be. To capture the relevant similarities, we would then have to deal with the above suggestion that one cannot have an experience of something *as* an *F* without that experience involving the concept *F*.

¹Peacocke 1983, p.7; as we shall see, Peacocke has changed his views since this was written.

²For sensational properties of experiences, see Peacocke 1983, chapter 1. Michael Tye argues that there are no such things in ‘Visual qualia and visual content’ (this volume). I should say at this stage that I use the terms ‘experience’ and ‘perception’ more loosely than some writers. The way I use ‘perceives’, ‘X perceives that p’ does not entail p.

There are two reasonable strategies we could employ. One is to say that the experience does involve concepts, but these are concepts that any two such perceivers will share. So, for instance, the infant and the scientist would experience the cathode ray tube as (perhaps) tubular, grey, shiny etc. The common content of their experiences would be given in terms of *observational* concepts that they both possess (see McGinn 1989, pp. 59ff.). The second strategy is to say that what they have in common is (at least) an experience with content, but content that is *nonconceptual*: it does not, in some sense, ‘involve’ concepts. It is this strategy I shall examine in this paper.

The claim that experiences have nonconceptual contents is a familiar one in recent philosophy of perception.³ But what does the claim really mean? Saying that an experience has a content that does not ‘involve’ concepts is, as yet, too vague. We need to know more about *concepts*, *contents* and the notion of a mental state ‘involving’ concepts before we can really understand the claim.

In this paper, I shall try to give enough of an account of these notions to yield a notion of nonconceptual content – one which is at least consistent with the leading theories of nonconceptual content. In Section 2, I shall consider how the term ‘nonconceptual content’ should be defined. In Section 3, I shall give a sketch of what it is for a state to be ‘composed of concepts’ or have conceptual content, and derive from that a sketch of what it is for a state to have nonconceptual content. I shall then argue, in Section 4, that this notion of nonconceptual content does apply to perception.

2 *What is nonconceptual content?*

What is this nonconceptual content supposed to be? I shall begin with three examples of states that philosophers have claimed to have nonconceptual contents.

First, there are perceptual experiences – perceiving that it’s raining, that the cat is on the mat and so on.⁴ Second, there are the states of the so-called ‘subpersonal’ computational systems like the visual system – states which

³See, in particular, Dretske 1981, chapter 6; Evans 1982, pp. 122–129, and pp. 154–160; Peacocke 1986a; Peacocke 1989c; and Peacocke’s ‘Scenarios, concepts and perception’ (this volume); Cussins 1990. The proposal has some affinities with the account of the individuation of experiences in Millar 1991.

⁴See Peacocke 1986a; Peacocke 1989c; and Peacocke, ‘Scenarios, concepts and perception’ (this volume); cf. Dretske 1981, p. 155 fn. 7, and p. 162.

have content, but whose content is not (unlike the first sort of state) phenomenologically salient.⁵ Third, there are the states that, following Dretske, we may describe as carrying ‘information’ in the sense of communication theory. To take a well-known example of Dennis Stampe’s (1977), if a tree has 70 rings, then it is 70 years old. Thus this state of the tree’s has a content – it carries the information that the tree is 70 years old (see Dretske 1981).

What makes these all examples of states with nonconceptual contents? Let’s start by assuming that all these states do actually have contents.⁶ To say that any state has content is just to say that it represents the world as being a certain way. It thus has what Peacocke (this volume) calls a ‘correctness condition’ – the condition under which it represents the world *correctly*. It is in this minimal sense of ‘content’ that perceptions, the states of the visual system, and the tree’s rings all have content.

It is well known that we will need to add more to this minimal definition to capture the kind of content beliefs and other propositional attitudes have. So we may say that the content of a belief is a Russellian Proposition, or a set of possible worlds, or something more fine-grained yet, like a Fregean Thought. But whatever the contents of beliefs are, we should not assume at the outset that just because experiences (and the other states mentioned) have content, they have the same sort of content that beliefs have.

This minimal definition of content does not mention concepts. It does not seem to be part of a state’s having a correctness condition that the state (or its content) ‘involves’ concepts. Yet we tend to assume that beliefs (or their contents) involve concepts. What does this mean?

It might mean that the concepts are literally the constituents of the belief. So when someone believes that the sun is shining, this belief state is literally composed of the concepts *sun* and *shining*. Or it might mean that the concepts are the constituents of the *content* of the belief: it is the content *The*

⁵See Evans 1982, p. 104 fn. 22; Campbell 1986, p. 172; Davies 1986; and Cussins 1990. Colin McGinn labels nonconceptual mental content ‘subpersonal content’: ‘the kind of content routinely attributed by cognitive scientists to information-processing systems of which the subject has no awareness’ (McGinn 1989, p. 163).

⁶Some philosophers dispute whether information in Dretske’s sense is really content (see e.g. Cussins 1990, p. 392). For my purposes, it doesn’t really matter if they are right, since I ultimately want to argue that perceptions have nonconceptual contents, and these are not just states with Dretskean informational content. So the fate of Dretskean information as content does not matter to me – I use it only as an example. Others (e.g. Stich 1983) deny that computational states too have content. I do dispute this; see Crane 1990.

sun is shining, rather than the belief state as such, that is literally composed of the concepts *sun* and *shining* – just as the sentence ‘The sun is shining’ is composed of its constituent words.

If this is what it is for a state to have *conceptual* content, then a state with *nonconceptual* content is one which is not composed of concepts. So if experiences have nonconceptual contents, then they do not have concepts as constituents. The experience of the sun shining is not composed of the concepts *sun* and *shining*. Or perhaps it is the *content* of the experience that is not composed of these concepts. So perhaps the idea is this: just as conceptual content is content that is composed of concepts, so nonconceptual content is content that is not composed of concepts. Conceptual content is ‘structured’ content, and nonconceptual content is ‘unstructured’ content.

But on the face of it, this definition will not explain why the examples I mentioned above are all examples of nonconceptual content. For the mere fact that a state has a structured content (in this sense) seems neither necessary nor sufficient for its being conceptual.

It is not necessary because it seems that some conceptual and nonconceptual states can share contents. When I believe that the sun is shining because I see that it is, then in an obvious sense I believe what I see. What I am believing is surely in some sense the same as what I am experiencing – namely, *that the sun is shining* – the way the world is represented to be, or the content, in the minimal sense of that term mentioned above. But if experiences have nonconceptual contents, and beliefs do not, then how can this apparently obvious fact be explained? How can the belief have a structured content if the experience hasn’t?⁷

It is not sufficient because it seems that the contents of computational states can be structured without being conceptual. According to certain computational theories of vision, light is reflected onto the retina, and the information it contains is processed – by algorithms that compute functions – to form a representation of the scene perceived. The states involved in this process represent visually perceptible properties of the scene: properties such as reflectance, illumination, and the orientation of edges of objects.

Now many think it obvious that the contents processed by the visual system do have constituents. For the theory assigns structure to these states, analysing them – as it may be – in terms of concepts such as that of a *zero-*

⁷At the end of §4 of this essay I shall suggest an answer to this question. The present point is that it is hard to see how it can be correctly answered if we simply define nonconceptual contents as ‘contents that are not composed of concepts’.

crossing (Marr 1982). But if a concept just is a constituent of a content, then the constituents of these computational contents will be concepts by definition: some of the states of a perceiver's visual system will have contents with zero-crossing-representations as constituents, and will thus be composed of concepts, including the concept *zero-crossing*. But these were supposed to be states with nonconceptual contents.

So the mere idea of a content that is not composed of concepts does not help to explain the idea of nonconceptual content. To this it may be said: so much for the idea of nonconceptual content. If the only sense we can make of the idea of a concept is as the constituent of a representational state (or of the content of that state), then why not just define 'concept' thus and have done with it? It is not, after all, as if the everyday notion of a concept is clear enough to be saved for systematic philosophy; and what can be saved applies just as well to these allegedly 'subpersonal' states as to beliefs.

But this reaction is premature. My point so far is only that the idea of nonconceptual content should not be defined simply in terms of having constituents. This does not mean that there is nothing in the idea of nonconceptual content. It's just that it shouldn't be defined that way. What I need to show is how it should be defined.

To get a better grip on the notion of nonconceptual content, take the simplest of our three examples: the tree's rings. Why should anyone think that this state does not 'involve' concepts? The answer is obvious: the state of the tree cannot involve concepts in any sense, since the tree has no mental states, and so *a fortiori* possesses no concepts. When we describe the tree as representing its age, or as carrying the information that it is 70 years old, we do not suppose that the tree possesses the concept of a year – or indeed any other concept. However, we think that when someone believes that they are 70 years old, they presumably cannot believe this unless they possess the concept of a year – *whatever* concepts may be. The believer needs to possess the concepts; the tree doesn't.

To generalise from this: for something, *X*, to believe that *a is F*, *X* must possess the concepts *a* and *F*. But for *X* to merely represent that *a is F*, *X* does not have to possess these concepts. It is in the latter case that *X* is in a state with nonconceptual content. Rather than being defined merely in terms of the content's having constituents, the notion of conceptual content is now defined in terms of whether its constituent concepts need to be *possessed* in order for something to be in that state.

Adrian Cussins has recently outlined and defended this idea (see Cussins 1990, pp. 380-401). It will be instructive to look at the way he sets the issues out. Cussins begins by defining conceptual and nonconceptual *properties*:

A property is a conceptual property if, and only if, it is canonically characterised, relative to a theory, only by means of concepts which are such that an organism *must have* those concepts in order to [instantiate] the property. A property is a nonconceptual property if, and only if, it is canonically characterised, relative to a theory, by means of concepts which are such that an organism *need not have* those concepts in order to [instantiate] the property.⁸

Conceptual and nonconceptual content are then defined as ‘content which consists of conceptual properties’ and ‘content which consists of nonconceptual properties’ respectively.

This definition is not entirely straightforward, and a few comments are in order. First, it is obvious that according to the definition, almost all properties are nonconceptual, since there are very few properties that require an organism to have a concept of the property in order to instantiate it. And that is how things should be.

But it looks initially as if some obviously conceptual properties come out as nonconceptual on this definition. Suppose, for the sake of argument, that the type-identity theory is true: mental properties are identical with brain properties. Then my instantiating the property *thinking about Vienna* is my instantiating a certain brain property – call it ‘B’. But I need have no concept of the property *B* in order to instantiate it – I don’t need to be a neuroscientist in order to think! So *B* is a nonconceptual property. But then so is the property *thinking about Vienna*, since it is identical with *B*. But surely you have to have a concept of Vienna in order to think about it?

Cussins’ answer is not that *thinking about Vienna* and *B* are different properties, but that specifying that property as ‘B’ does not, from the point of view of the theory of content, specify it in the right way. It is not what he calls a ‘canonical characterisation’ of the property. A theory canonically characterises something when it describes it in terms of the properties that it treats as essential to that thing. Content, Cussins says, is ‘canonically characterised by a specification which reveals the way in which it presents

⁸Cussins 1990, pp. 382–383. I replace Cussins’ ‘satisfy’ with ‘instantiate’, since the term ‘satisfy’ is more usually applied to predicates rather than properties.

the world' (see Cussins 1990, p. 383 fn. 25). So describing my instantiation of the property as *B* does not reveal the way the world is presented to me – the world is not presented, in that instantiation, as containing instances of *B*. This is why 'B' is not a canonical characterisation of the property.

Second, it is plain from this that although 'B' is not a canonical characterisation of the property for the purposes of the theory of *content*, it will be canonical for, say, neuroscience. This is why Cussins says that canonical characterisations are 'relative to a theory'. He thus leaves open the possibility that the same property could be canonically characterised as conceptual by one theory, and as nonconceptual by another.

The route from conceptual and nonconceptual properties to conceptual and nonconceptual content is straightforward. A state with conceptual content – e.g. a belief – is one such that the subject of that state has to possess the concepts that canonically characterise its content in order to be in that state. Any state with content that does not meet this condition has a nonconceptual content. More strictly, we can use Cussins' definition to construct a definition of nonconceptual content as follows:

For any state with content, *S*, *S* has a nonconceptual content, *P*, iff a subject *X*'s being in *S* does not entail that *X* possesses the concepts that canonically characterise *P*.

This definition can now be applied to the examples I introduced at the beginning of this section.

(1) Experience: *X*'s seeing that the sun is shining does not entail that *X* has the concepts of *the sun*, or *shining*. That is, in order to see that the sun is shining, a subject does not *have* to possess these concepts (though of course he or she *may* possess them).

(2) The visual system: in order for a subject's visual system to compute its solution to the complex equations that take retinal information as input and a 3D description of the scene as output, the subject does not have to possess the concepts that canonically characterise these equations (though again, he or she may possess them). You don't need to know the theory of vision in order to see.

(3) The tree's rings: in order for the tree to represent or indicate that it is 70 years old, it does not have to possess the concepts *70*, *year* and *being old*. This is just as well, since it wouldn't be able to possess these concepts anyway.

Saying that it makes sense to *apply* the definition of nonconceptual content to these cases is not, of course, saying very much. We still have to establish that there are any mental states which have nonconceptual content. To do this, we must probe the definition further.

A state with conceptual content, *C*, may be defined in terms of the conditional: if *X* is in *C*, then *X* possesses the concepts that canonically characterise *C*. The point, therefore, is not that conceptual contents have constituents and nonconceptual contents don't – as we saw above, this won't work. It is rather that *X* needs to 'possess' these constituents in order to be in the state.

But what is it to 'possess' these constituents, these concepts? It cannot simply be a matter of having the literal parts of the state 'written' inside one's head. (After all, the tree has 70 rings; in no sense does this amount to possessing the concept 70.) There must be more to possessing a concept than being in contentful states that have constituents. But what more?

It seems, then, that to understand what concepts are, we need to understand what it is to possess a concept. Indeed, Christopher Peacocke has recently advanced an even stronger thesis, his 'Principle of Dependence': that 'there can be *no more* to the nature of a concept than is determined by a correct account of what it is to possess the concept'.⁹ In the next section, I will try to give an argument for this Principle.

3 *Concepts and contents*

So what is it to possess a concept? It is often said that thought requires the possession of concepts (see Evans 1981, p. 132). This can seem almost tautological, as can its converse: there is no possession of concepts without thought. This latter truism is almost self-evident: it is hard to make sense of a thinker whose mental life just consists of a series of concepts, with no intentional states of which they are 'constituents'. So we may conjecture that a concept, *C*, is possessed by a thinker, *T*, only if *T* is in intentional states in whose contents *C* figures. But is the former truism just as obvious? Why should thought require the possession of concepts?

Consider my possession of the concept *cheese*. According to the above conjecture, I only possess this concept if I am in intentional states in which the concept *cheese* figures. Suppose, for example, I believe that cheese is

⁹Peacocke 1989c, p. 2 (my emphasis); cf. also Peacocke 1989b, p. 51, where Peacocke says that 'the idea that a concept is individuated by a correct account of its possession' is 'a master key' to the theory of concepts.

nutritious. It will not be possible for me to have only *this* belief about cheese, since I cannot believe that cheese is nutritious (call it '*P*') without having certain beliefs that are the obvious logical consequences of this belief – the belief that $\text{not}(P \ \& \ \text{not-}P)$ for example.¹⁰ These beliefs are uninteresting consequences of *P*, but there are more interesting beliefs I must have too, that are not related to *P* by logic alone. For example, if I believe that cheese is nutritious, I must believe that cheese is edible, and arguably, if my belief is genuinely to be about cheese, I must believe that cheese is made of milk.

This list of beliefs could be extended; but there is no need to do so to underline the familiar point that having one belief entails having a lot. While there is, arguably, nothing incoherent in supposing that a mental life could consist of just one sensation, it seems plainly incoherent to suppose that it could contain just one intentional state. If a thinker has a belief, then he or she must also have many others. These beliefs will be either those related logically to the original belief (*P* or *not-P*) or those related by what we can call the 'semantic' properties of their contents (the belief that cheese is nutritious and the belief that cheese is edible).

So if I can't have the *concept* of cheese unless I have beliefs in which the concept figures – whatever concepts are – then I can't have the concept if I only have one belief about cheese. For I can't have only one belief about cheese. Since intentional states come not in single spies but in whole battalions, then since possession of concepts needs intentional states, it needs a multiplicity of them.

Moreover, the reason why there has to be this multiplicity is that beliefs have to stand in these relations to other relevant intentional states in order for them to have the contents they do. The upshot of the 'cheese' example is that the belief just wouldn't be *about cheese* unless the thinker had the other relevant beliefs. The much-discussed 'holism of the intentional' resides not simply in the fact that there must be a multiplicity of intentional states if there are to be any, but in the fact that the content of any one intentional state depends, to some extent, on the contents of the others. Indeed, it is this latter fact that *explains* the former.

But a thinker's beliefs about the world are also sensitive to perceptual evidence, and in some cases, their contents are partly defined by the perceptual evidence that a thinker would take as counting in their favour.

¹⁰People do not *always* believe the logical consequences of what they believe, of course. I am here skating over a number of complications, in order to give the general picture.

Someone could not believe that it's raining unless they were disposed to regard a perception of drops of water falling from the sky as evidence for this belief. Intentional states stand not only in logical and semantic relations to one another, but also in evidential relations to perceptions. All these three kinds of relations help to fix the content of a given intentional state.

But what does this tell us about concepts and their possession? These uncontroversial points about holism do not answer our initial question about why thought requires the possession of concepts. To answer this question, we next need to ask why a theory of the mind needs the notion of a concept at all.

When looking for the notion of the *content* of an attitude, there are straightforward dissections we can make in the anatomy of the mind – we can dissect people into their mental properties and their other properties; then we can dissect their mental properties into their intentional states and their nonintentional mental states; and we can dissect their intentional states into their attitudes and the contents of these attitudes. That gets us down to whole contents, the bearers of truth values. But where do concepts fit in? What do we add to the claim that people have beliefs, desires, intentions (and so on) about cheese by saying that they have the *concept* of cheese? Why talk about concepts at all? (See Hart 1983.)

To answer this, we need to know what the idea of a concept is meant to explain. Frege writes:

The task of our vernacular language is essentially fulfilled if people engaged in communication with one another connect the same thought, or approximately the same thought, with the same proposition. For this it is not at all necessary that the individual words should have a sense and meaning of their own, provided only that the whole proposition has a sense. Where inferences are to be drawn the case is different: for this it is essential that the same expression should occur in two propositions and should have exactly the same meaning in both cases. It must therefore have a meaning of its own, independent of the other parts of the proposition.¹¹

¹¹The quotation is from a letter from Frege to Peano (29 Sept 1896) in Frege 1980, p. 115. I am grateful to Bill Hart for drawing my attention to this passage. Notice, of course, that the term 'proposition' in this context means what I mean by 'sentence', and the term 'thought' means what I call 'content'.

Frege was concerned here with language, but a parallel point can be made for thought. If we simply wanted to represent facts, then our beliefs would only need to have ‘whole’ contents. All that would matter would be whether a content was true or false. The *fact* might have constituents (particulars and properties) but they would have no reflection in the content, since (to echo Frege) they would as it were have no role, no ‘meaning of their own’.

But once we consider the role our beliefs play in reasoning, then it starts to become clear why their contents need constituents. A thinker who believes that *a is F*, and that *b is F*, and that *a is not b* will be disposed to believe that *at least two things are F*. Surely the states in this inference cannot just have unstructured contents, or or we would not be able to explain its validity. And if we cannot explain this, we cannot explain why the constituent beliefs have the contents they do – according to the holistic proposal just outlined.

To recall Frege’s remark: it is essential that the same (type of) ‘part’ of the content should occur in the two states, and that ‘it should have exactly the same meaning in both cases’ – i.e. the parts should be tokens of the same *semantic* type. So in the simple example above, the states must both contain *F* as a part. I say that it is only a terminological variant of this to say that they must contain the *concept F*.

To account for the inferential powers of these beliefs, then, the thinker’s beliefs should contain an element common to their contents: the concept *F*. What I want to suggest, then, is that concepts are the *inferentially relevant constituents of intentional states*. To discern conceptual structure in a thinker’s thoughts, we need to look (as it were) not ‘down’ into the propositional content, but ‘up’ into the nature of the intentional states themselves. The idea is to derive the psychological notion of a concept from facts about the inferential relations among beliefs.¹²

The notion of *possessing* a concept is then naturally explained as follows. To possess a concept is to be in intentional states whose inferential relations are an appropriate function of their contents. The elements in a thinker’s network of intentional states are essentially inferentially related to one another. Concepts are the constituents required to explain these inferential

¹²The idea is not new, of course. Compare Gareth Evans’ remark that ‘behind the idea of a system of beliefs lies that of a system of concepts’ whose structure ‘determines the inferential properties’ of the beliefs (Evans 1981, p. 132). John Campbell (1986) has argued from inference to structure, but not in the very general way suggested here. Campbell’s strategy is to examine certain primitive patterns of spatial reasoning, and show how they thereby exploit the conceptual structure of the thoughts involved.

relations. So a thinker could not be in the relevant intentional states unless they contain concepts. Since possessing concepts entails that one is disposed to make certain inferences, then possessing concepts entails that one's intentional states are 'composed' of concepts. This establishes the link we sought at the end of Section 2, between a state's being *composed* of concepts and its subject *possessing* concepts.

To say that beliefs have conceptual structure because of their inferential relations is not yet to say anything about the mechanism that underpins this structure. The structure of this mechanism need not mirror the inferential structure of beliefs. The mechanism could be a giant 'look-up table' that 'infers' the 'belief' that *at least two things are F* from the inputs *a is F*, *b is F* and *a is not b* simply by having the first 'belief' written, as it were, at the intersection of the others.¹³

But I do not need to be committed to any *a priori* claims about mechanisms. My definition of a concept is a claim about the relations between beliefs, not a claim about their underlying mechanisms. This is why the definition does not commit me to the Language of Thought (LOT) hypothesis (Fodor 1987, appendix), for that is precisely the claim that the structure of the mechanism of thought mirrors its inferential structure. To say that beliefs are inferentially related to one another, and that they therefore have a structure that makes this possible, does not entail that this structure is underpinned by anything *syntactic* – in any interesting sense of 'syntactic'. Maybe the underpinning mechanism is that of a connectionist network. The LOT hypothesis is independent from the present claim about concepts.¹⁴

This proposal about concepts is not *ad hoc*, for at least two reasons. First, it explains why the claim that there is no thought without possession of concepts is not just a vacuous slogan. We should distinguish between the content and having an attitude with that content – in Frege's terminology, for example, between the Thought and its 'grasping'. I claim that if contents are to be grasped at all, their inferential relations must be grasped. That is, according to my definition of possessing a concept, contents must be grasped

¹³For a giant look-up table, see e.g. Clark 1989, p. 56. I have learned much from Barry Smith on this issue.

¹⁴This paragraph merely glances at a number of complex issues. For the idea of syntax as used in the LOT hypothesis, see Crane 1990. Martin Davies (forthcoming) has argued, ingeniously, that the conceptual structure of thought needs the truth of the LOT hypothesis. But I agree with Fodor and others that the LOT hypothesis should rather be thought of as the best empirical explanation of conceptual structure.

by means of the thinker possessing certain concepts. So there is no thought without possession of concepts because states can have the contents they do only if they have inferentially relevant structure.

Second, the proposal offers us an independent motivation for (something like) Peacocke's Principle of Dependence, that 'there can be no more to the nature of a concept than is determined by a correct account of what it is to possess the concept'.¹⁵ A correct account of possessing a given concept, *C*, would on my proposal involve a specification of the beliefs one has to have in order to have any belief about *C*s. The account would have to specify what those beliefs are, and how they are related.¹⁶ The constituent, *C*, of these relevant beliefs is invoked only to explain the inferential relations between them. Since that is what the concept *is*, then of course there is no more to its nature than is given in an account of its possession.

We can now return to nonconceptual content. In Section 2, I said that a state, *S*, has conceptual content iff *X*'s being in *S* entails that *X* possesses the concepts that characterise *S*. In this section I have argued that a concept is an inferentially relevant constituent of an intentional state; and that possessing a concept is therefore being in states with inferentially relevant constituents. So *S* has a conceptual content iff *X*'s being in *S* entails that *S* has inferentially relevant constituents, and this requires that *X* is in other states which are inferentially related to *S*.

Similarly for nonconceptual content. *X* is in a state with nonconceptual content iff *X* does not have to possess the concepts that characterise its content in order to be in that state. Since possessing a concept is being in intentional states whose contents are appropriately inferentially related, then a state with nonconceptual content is one whose contents are not so related. So in order to be in such a state, one does not have to be in other inferentially related states of the kind that give the contents of beliefs their conceptual structure.

4 *Belief, experience and nonconceptual content*

What has all this got to do with experience? I think that this proposal about concepts provides us with a two-stage argument for the nonconceptual content of perceptual experience. The first stage is to show why, contrary to

¹⁵Peacocke argues that inferential role is necessary for the possession of certain concepts – see his remarks about the concept of conjunction in Peacocke 1989b, pp. 51–52.

¹⁶Which beliefs these are for any given concept would thus determine whether 'prototype' or 'definition' theories of that concept are correct. See Putnam 1975 and Fodor 1981.




Figure 1 – The Müller-Lyer Illusion

what some philosophers think,¹⁷ perceptions are not beliefs. The second will be to show how this entails that perceptions have nonconceptual contents.

Any theory of perception must explain how perceptions can give rise to beliefs – what makes it the case that we can believe what we see. The belief theory of perception gives a simple answer to this: perceptions just *are* beliefs, acquired in a certain way.

One standard objection to this is that we often do not, and cannot, believe what we see. Consider the famous Müller-Lyer Illusion, where a perceiver, *P*, is presented with two lines of equal length, one with arrows going out (*L1*), the other with arrows going in (*L2*), and suppose *P* does believe that the lines are the same length. No matter how strong this belief of *P*'s is, *P* cannot help but see *L1* as longer than *L2*. A theory that holds a straightforward equation of perception with belief has to say, on the face of it, that the perceivers have contradictory beliefs. But this is implausible. Surely it's better to say that they are in two states with different contents, perceiving that *L1* is longer than *L2*, and believing that *L1* and *L2* are the same length.¹⁸

It will not do to explain the illusions in terms of a 'prima facie but suppressed' inclination to believe that *L1 is longer than L2* – in Armstrong's words, in terms of 'a state which would be a belief state but for the inhibiting effect of other, contrary beliefs' (Armstrong 1968, p. 140). For there would still be an unexplained component – the difference between this state and

¹⁷See especially Armstrong 1968, chapter 10, and Pitcher 1971. I would also count as belief theories those theories which employ notions like 'judgement' (Craig 1976) or 'hypothesis' (Gregory 1970) – notions which must be defined in terms of belief. Many of the arguments for the belief theory are disposed of by Jackson 1977.

¹⁸Compare the curious 'Waterfall Illusion': see Frisby 1980, p. 100, and Crane 1988a.

other, *nonperceptual*, inclinations to believe. When conclusive evidence is presented against my nonperceptual inclination to believe that p , this inclination should vanish. But when conclusive evidence is presented against my (alleged) inclination to believe that $L1$ is longer than $L2$, this alleged inclination remains. (Happy, therefore, are those who believe and have not seen.) Perceptions are resilient to conclusive counter-evidence in such a way that precludes their definition simply in terms of inclinations to believe (see Crane 1988b).

None of this, of course, prevents us being *normally* disposed to believe what we perceive. Perceptions (unlike desires, for instance) are representations of how the world *is*, and thus ‘made true’ by the facts. In Searle’s terminology, they have ‘mind-world direction of fit’ (Searle 1983, p. 53). Perceptions seem to ‘aim’ at truth in something like the way beliefs do. But the way perceptions aim at truth is not the same as the way beliefs do. For part of what it is for belief to aim at truth is shown by Moore’s so-called ‘paradox’: the absurdity, for all p , of asserting ‘I believe that p but not p ’. Yet as the Müller-Lyer Illusion shows, there are values of p for which asserting ‘I perceive that p but not p ’ is perfectly coherent. There is thus no Moorean ‘paradox’ of perception.

So one clear reason perceptions are not beliefs – and there are others – is that they are not revisable in the light of either other perceptions or beliefs that the perceiver has. Although perceptions do normally involve inclinations to believe, they cannot be identified with these inclinations, since unlike inclinations to believe, they are resilient to conclusive counter-evidence.

In Section 3 I claimed that beliefs were holistically related to one another by at least three kinds of inferential relations. The first were the logical or deductive relations. The second were what I called the ‘semantic’ relations. And the third were the evidential relations. I claimed that since beliefs would not have the contents they do if they did not stand in these relations, then standing in these relations was essential to beliefs. Moreover, it is because beliefs stand in these relations that they have conceptual structure.

Now take the case of evidential relations. Certain beliefs are partly characterised by the evidence – perceptions and other beliefs – that believers count for or against them. They are therefore revisable on the basis of such evidence. But as we have seen, perceptions are not like this. While they may *be* pieces of evidence, they are not *revisable* on the basis of other evidence – whether that evidence is another belief or another perception. Moreover, if conceptual structure is *only* imposed by these evidential relations and the

other inferential relations, then perceptions will not have conceptual structure. This is why their contents will not have inferentially relevant constituents: they will not be composed of concepts.

Although perceptions are unrevisable, the conclusion that they have nonconceptual contents will not follow from this fact alone. For they may yet stand in some of the other inferential relations that are essential to beliefs. Do these other relations hold between experiences?

Take the deductive relations first. On the face of it, it seems plain that there is no such thing as deductive inference *between* perceptions. If I perceive that *a is F*, and I perceive that *a is G*, there is no such thing as inferring the perception that *a is F and G*. If *a's F-ness* and its *G-ness* are both perceptible at the time of the two perceptions, then the content of my perception is surely *a is F and G* in the first place. But if *a's F-ness* and its *G-ness* are not both perceptible on the same occasion, then I can only infer the *belief* that *a is F and G*. That is, if I perceive the fact that *a is F and G*, then there is one perception; but if I don't, there are two perceptions and one belief. Either way, there seems to be no deductive inference between perceptions.

Or consider the case of contents of the form *P & not-P*. One of the deductive constraints on beliefs is that we cannot have explicitly contradictory conscious beliefs. But, in the case of certain striking visual illusions – notably the 'Waterfall Illusion' – we can have contradictory perceptions: a conscious perception with an explicitly contradictory content (Crane 1988a). So once again, perceptions lack certain individuating features of beliefs.

However, matters are not quite as straightforward as this.¹⁹ Can I see that the table is brown and rectangular without also simultaneously seeing that it is brown? Surely not. So why doesn't this mean, on my account, that the content of the perception must have constituents of some sort – a constituent that 'means' *rectangular* and one that 'means' *brown*?

To say that perceptions lack inferential structure of the kind typical of beliefs does not mean that they are entirely unstructured. But reflection on this example shows that there is an important difference between the *kinds* of structure involved in the cases of perception and belief. The difference is that in the case of perception, there is strictly speaking no such thing as the perception of the brownness of the table, in isolation from the perception of

¹⁹I am indebted to Mark Sainsbury for pressing this objection.

its rectangularity. It is simply not possible to perceive that the table is brown without also perceiving its actual shape.

Of course, it also isn't possible to *believe* that the table is brown without believing that it has a shape. But the difference is that whereas the perception is of the table's *actual* shape, the belief is merely that the table *has* a shape. It would be possible for someone to believe that there is a brown table in a certain room without having any beliefs about its definite size or shape. It would be enough to believe simply that it *has* a definite size – somewhere between the size of the room it's in, and the size of a shoebox, perhaps; and a definite shape – round, rectangular, or hexagonal etc. But it would not be possible for a normal perceiver to see a brown table in the room without also seeing its actual size and shape.

Dretske expresses this distinction by saying that the perception of the table contains 'more specific, more determinate' information than the belief about the table.²⁰ The mere belief that there is a brown table in the room is less specific and determinate than the perception, in the sense that its correctness is compatible with more possibilities. The correctness of this belief is compatible with the table being 6 feet long, or 4 feet 6 inches – and so on. Within certain obvious limits, the belief about its colour and location doesn't rule any of these out. The perception that the table is brown, on the other hand, rules out many more possibilities. There are far fewer ways the length of the table can be that are consistent with the correctness of the perception; so the content of the perception is more specific in that it rules out more possibilities.²¹ (This distinction captures the important and nonmisleading sense in which perceptions are more like pictures than statements.)

We may put the point thus: the content of the perception that the table is brown *already contains* the perception of its shape. But the mere belief that the table has a shape can only be *inferred* from the belief that the table is brown, plus the *general* belief that whatever has a colour has a shape. The belief about the shape arises, plainly, from a process of reasoning. But the

²⁰Dretske 1981, chapter 6, p. 137. This is what he means by saying that the content is coded in *analogue* form. Note that Dretske means 'information' in his own sense, but in this context it can be easily read as 'content'. In this context, see also Peacocke, 'Scenarios, concepts and perception' (this volume) §2.

²¹I do not think, however, that the content of an experience should be represented in standard possible worlds terms – whatever their merits for the contents of beliefs – since it is possible to have contradictory experiences: experiences whose contents would be sets of impossible worlds. Again, see Crane 1988a.

perception of the shape doesn't, for the obvious reason that you *can't* have a perception of the general fact that whatever has a colour has a shape, since such general facts are not perceptible. But as we have seen, you don't *need* such a general 'perception' in order to perceive the shape of the table and its colour – which is just as well, since it can't be had. So although it is possible to infer ascriptions of perceptions from one another, this doesn't entail that the perceptions *themselves* enter into deductive relations.

This brings me to the question of whether those relations that I called 'semantic' – those relations between contents that are not just a matter of logic – apply to perception. The idea was that if you have the belief that *p*, then there are certain other beliefs that you *ought to have* if that belief is to have the content *p*. Once again, perception doesn't share this feature of belief: to perceive that *p*, there are no other *perceptions* that you *ought to have*. There is no 'ought' about it. You simply perceive what the world and your perceptual systems let you perceive. If these systems go wrong, then they can produce states with contents – e.g. contradictions – that the belief system would not tolerate. But unlike the case of beliefs, failing in this way does not stop the perceptions from *having* those contents. (A Davidsonian might put this point by saying that perception is not subject to the 'constitutive ideal of rationality': see Davidson 1970, p. 213.)

So I conclude that the structure in the contents of perception is not conceptual structure: that is, the inferential structure of the contents of beliefs. But how then should we understand its structure? My general account of nonconceptual content does not entail one particular answer to this question. But it is consistent with the most fully developed current theory of nonconceptual content – Peacocke's theory of 'scenario content'.²²

Peacocke's idea is that the content of a perception can be given by specifying a *scenario*: a set of ways of filling out the space around the perceiver with properties (colours, shapes, temperatures and so on) relative to an origin and a family of axes. The origin will depend on the position and posture of the perceiver. When specifying a given scenario, we will of course use concepts that pick out the various aspects of the scenario – the properties that constitute it. But, Peacocke claims, 'the fact that a concept is used in

²²See Peacocke 1989c, Peacocke 1989b and 'Scenarios, concepts and perception' (this volume). The last paper introduces a distinction between scenarios and another 'level' of nonconceptual content, the 'protoposition'. I do not need to discuss this refinement here.

fixing the scenario does *not* entail that that concept is somehow a component of the experience's representational content' (Peacocke 1989c).

This claim fits neatly into my account of nonconceptual content. According to the argument of Section 3, to say that concepts are not components of contents is to say that the subject does not have to possess the concepts used to characterise the content in order for his or her state to have such a content. Applied to scenarios, this means that the subject's being in a state with a scenario content does not entail that he or she possesses the concepts that characterise the scenario. And this in turn means that the subject does not have to be able to make inferences involving beliefs about the properties of the scenario, in order for his or her *perceptual* states to have such contents. And this seems right: surely if experiences have such scenario contents, we should not expect perceivers to have beliefs about all the myriad properties that characterise them, and make inferences involving those beliefs. So although a scenario content does have a *structure*, it does not have inferential structure – and so does not have conceptual constituents.

There is still an outstanding issue. How will this picture of perception explain the relation between perception and belief? Perceptions cause beliefs, and these perception-caused beliefs interact with other beliefs inferentially. It may then be asked: how can my account of the content of perception answer our question about the relation between perception and belief? Why doesn't this inferential relation impose structure?

My answer is this: when a perception that *p* causes a belief that *p*, the *whole* contents of these two states are of the same type – *p*. The (causal) relation between perception and belief takes place at the level of whole contents. But on the perception side of this transaction, the contents are not composed of concepts: concepts come later when thinkers employ the beliefs they thus formed, and the desires they have, in reasoning – belief *conceptualises* the content of perception. So treating the transition from perception to belief in terms of whole contents allows us to explain how perceptions have contents that can be the contents of beliefs. It is crucial here to remember the argument of Section 3: the idea of a content was not explained as a combination of concepts. We extracted the idea of a concept from the idea of a system of beliefs; we did not extract the idea of the content of a belief from an antecedently given idea of a concept.

That completes my argument for the nonconceptual content of perception. But what about the other two examples of nonconceptual content that I mentioned at the beginning of Section 2? Will my account apply to them too?

I have nothing to say on the subject of the rings of the tree representing its age. Since for me, possessing concepts entails at least having beliefs, then since the tree has no beliefs, *a fortiori* it possesses no concepts. And I need not take a stand here on the controversial issue of whether the tree is in a state with *content*.

The case of the visual system is more complex. As I said in Section 2, many people think that the contents of states attributed by computational theories of vision are nonconceptual. But these theories often describe the stages in visual information-processing in terms of making inferences.²³ So if these theories are right, how can my argument possibly apply?

I shall end this paper by sketching an answer to this question. There is an important difference between the ways in which intentional states are inferentially related in thought, and the way computational states are related. The inferences a thinker is disposed to make, on which the conceptual structure of his or her thought depends, are constrained only by rationality, which allows the mind to range over its whole territory for its material. The resources of the visual system are, by contrast, severely restricted. Although we can treat it as deducing consequences from premises, the contents of these premises are not holistically related in the way the contents of beliefs are.

Consider, for example, how Irvin Rock explains one striking effect of the Julesz random dot stereograms:

it would seem that the perceptual system ‘knows’ certain laws of optics that normally obtain, and then ‘interprets’ seeming departures from these laws in such a way as to be compatible with them.

(Rock 1983, p. 10)

Although it may well be correct to attribute to the visual system representations of certain laws of optics, this is not the same as what goes on when we know or believe laws of optics. To echo a remark of Gareth Evans: to make sense of this case, we do not need to suppose that the visual system is capable of entertaining contents about any other sort of law.²⁴ But this is

²³See Segal 1989, at p. 192. Cf. Marr: ‘the true heart of visual perception is the inference from the structure of an image about the structure of the real world outside’ (Marr 1982, p. 68).

²⁴I have in mind a well-known passage from Evans 1982, p. 104: ‘When we attribute to the brain computations whereby it localises the sounds we hear, we ipso facto attribute to it representations of the speed of sound and the distance between the ears, without any commitment to the idea that it should be able to represent the speed of light or the distance between anything else.’ See also Davies 1989.

precisely what we *do* suppose when attributing to someone a belief about a law – we suppose that the subject can make inferences about laws, and thus entertain other contents about laws. If subjects couldn't do this, then they wouldn't have the concept of a law.

But this isn't so with the visual system – for its states to have content, they do not need to be so holistic. And I have argued here that the holistic constraints of rationality are the only motivation for postulating concepts. So the fact that the visual system does not meet these constraints helps explain why it possesses no concepts. And this would be why its states are not *composed* of concepts, in the sense I have defended in this paper.²⁵

²⁵Earlier versions of this paper were presented at University College London, the Department of History and Philosophy of Science, King's College London, and at the Universities of Sheffield, Stirling and York. I am grateful to the audiences on these occasions, and especially to John Campbell, Bill Hart, Stephen Makin, Mike Martin, Hugh Mellor, Alan Millar, Philip Percival, Tom Pink, Mark Sainsbury, Barry Smith and Michael Tye for discussions and written comments that have helped me greatly with the paper.