

Perceptual Experience And Its Contents

Josefa Toribio

Indiana University

The contents of perceptual experience, it has been argued, often include a characteristic “non-conceptual” component (Evans, 1982). Rejecting such views, McDowell (1994) claims that such contents are conceptual in every respect. It will be shown that this debate is compromised by the failure of both sides to mark a further, and crucial, distinction in cognitive space. This is the distinction between what is doubted here as *mindful* and *mindless* modes of perceiving: a distinction which cross-classifies the conceptual / non-conceptual divide. The goal of the paper is to show that there can be both *mindful* personal level perceptual experiences whose content cannot be considered conceptual — pace McDowell (1994)— and that there are *mindless* personal level perceptual experiences whose content cannot be considered —pace Evans (1982)— non-conceptual. The resulting picture yields a richer *four dimensional* carving of the space of perceptual experience, and provides a better framework in which to accommodate the many subtleties involved in our sensory confrontations with the world.

The notion of internal representation plays a foundational role in our scientific understanding of cognition. Intelligent behavior is usually explained by positing some kind of inner representational entities: items that are semantically interpretable and causally efficacious. The idea underlying this plausible conviction is that cognition is not just a process of *taking in* information. We are only justified in talking about contentful mental states if our representations can also constitute a reason for action (Davidson, 1970). Psychological explanations, especially those pertaining to the domain of so-called folk psychology, invoke mental states with particular intentional contents in order to explain or justify a person’s course of action. The individuation of a mental state as a state with

such-and-such particular content is thus warranted by finding an explanatory link between the representational properties of such a mental state and the behavior exhibited by the system. Those representational properties, in turn, are usually characterized in terms of the system's possession of a certain conceptual repertoire.

Perceptual experiences, however, appear to present a special case. Do agents really need to possess the concepts involved in describing the contents of their perceptual representations? And, if not, what kind of content is characteristic of perceptual experience? One possibility, originally suggested by Evans (1982), is that the contents of perceptual experience include a distinctive "non-conceptual" component. We are still justified in individuating a particular mental state as a state with such-and-such a content in virtue of the explanatory connection between its representational properties and the system's behavior, but we don't need to posit the possession by the agent of the concepts involved in the description of such a mental representation. Perceptual experiences, on this view, provide the subject with a *primitive* version of reasons for action. Without needing to possess, for example, the concepts of "square" and "fitting," an agent may nonetheless see that the square peg will fit into the square hole. The non-conceptual content of perceptual experience is thus genuine content, since it is able to constitute a reason for action. Yet, it may be possessed by subjects lacking the concepts which external observers would use to describe the cases concerned.

McDowell (1994) famously rejects such views, and depicts the kinds of contents which might figure as genuine reasons for actions as conceptual through

and through. If the perceived squareness is indeed to count as a *reason* for the action, according to McDowell, the subject must be in command of at least a minimally articulate and conceptualized understanding of the situation. She might then appeal, for example, to the way the peg looks to her if asked why she expected it to fit through the hole. Such a minimally articulate response, however, demands the possession of conceptual capacities over and above those required for the putative primitive experience. For McDowell, it is only conceptualized perceptual experiences that can help us negotiate the “space of reasons.”¹ Non-conceptual contents, insofar as they remain non-conceptual, cannot constitute reasons for action. Such a view relegates the notion of non-conceptual content to philosophical limbo, leaving it respectable only when applied sub-personally and outside the space of reasons.

McDowell’s dismissal of the rich notion of non-conceptual content is, we shall argue, too swift. But it is invited by a failure on the part of *both* McDowell *and* the friends of non-conceptual content, to mark an additional distinction in cognitive space. This is the distinction between what it is called in this paper *mindful* and *mindless* modes of believing, judging, and perceiving. The mindful / mindless distinction, as we shall show, cross-cuts the conceptual / non-conceptual divide. Non-conceptual contents may be mindful or mindless, and conceptual contents likewise. Very roughly, a thought, experience, or belief is *mindful* insofar as we are disposed to fulfill certain additional cognitive duties in respect of its content, and *mindless* insofar as we are not. To be disposed to fulfill such duties is to be semantically responsible concerning the thought, belief, or experience in question. We shall argue that there can be personal level non-conceptual perceptual

experiences for which an agent is semantically responsible (i.e., *mindful* representational states), and, conversely, personal level conceptualized perceptual experiences for which an agent is not semantically responsible (i.e., *mindless*). The resulting picture yields a *four dimensional* carving of perceptual experience in place of the two dimensional carvings common to both the friends and foes of non-conceptual content.

Non-Conceptual Content: What is it and Why Do We Need it?

It is best to begin where the light is brightest. Under what conditions may the content of a representation be said to be conceptual? Conceptual contents, it is generally agreed, are subject to what Evans (1982) terms a “generality constraint.” According to Evans: “It is one of the fundamental differences between human thought and the information-processing that takes place in our brains that the Generality Constraint applies to the former but not the latter” (Evans, 1982, p. 104, fn. 22). The content of inferentially articulated contentful mental states such as beliefs is subject to this constraint. The content of perceptual experiences, Evans suggests, is not subject to the generality constraint. This means that the attribution to a system of contentful states of the form *a* is *F* and *b* is *G* does not here commit us to the idea that that system should also be able to represent *a* as *G* or *b* as *F* (Evans, 1982, p. 104). Sometimes, these ideas are formulated by saying that the content of perceptual experiences is non-conceptual because such content is not individuated according to Frege’s notion of a mode of presentation (Peacocke, 1986, p. 14).

To token a specific non-conceptual content, subjects need not themselves possess the concepts that an external observer would use to characterize the content of the experience. Instead, non-conceptual content is *unmediatedly* and *intimately* connected to the subject's abilities to act upon an object or to perform a particular task (cf. Evans, 1982, p. 146). To adjust one's grasping actions to suit different kinds of object, or to recognize similarities and/or dissimilarities in shape and structure, are examples of such abilities, which "... are not available to the subject as the content's referent, but they are available to the subject as the subject's experience-based knowledge of how to act on the object, and respond to it" (Cussins 1992, pp. 655-656). To take one of Evans' examples, consider the experience of hearing a sound as coming from "over there." What makes it the case that a subject hears the sound in that way is, according to Evans, the subject's particular ability to negotiate the domain in which she is embedded (Evans 1982, p. 154). It is to possess a set of skills, and other know-hows that will enable her to carry out a particular task, where such skills need not include any explicit theoretical knowledge. She need not have, for example, the concept "north" to be able to have an experience of a sound as coming from that direction. In Evans' account, it is, however, required that the subject be able to feed these non-conceptualized contents into the conceptual apparatus of reasoning for them to be considered contents at all. Otherwise, such perceptual states are just the informational states of sub-personal mechanisms. This extra requirement is not endorsed by every non-conceptualist (see e.g., Bermúdez, 1994, 1995, 1998; Bermúdez and Macpherson, 1998; Davies, 1992, 1997; and McGinn, 1989, for

positions according to which non-conceptual content can exist even for organisms with no concepts whatsoever).

But why invoke a notion of non-conceptual content at all? The rationale and the arguments are different for different authors. The desire to provide a naturalistic account of semantic behavior in general, and of concept possession in particular, is prominent in Christopher Peacocke's work (1992a) who believes that the explanation of what it is to possess a concept would be circular if we always used the concept whose possession conditions we are trying to explain in the specification of the content. The normativity of concept-involving behavior needs, Peacocke claims, to be *explicable* in terms of the properties of more basic non-conceptual, yet still representational, cognitive states. Non-conceptual content thus plays a crucial role in Peacocke's attempts to provide a non-circular, naturalistically acceptable, account of what it is to possess a concept since *non-conceptuality* both helps avoid circularity in the determination of a concept's possession conditions and allows for a naturalistic explanation of concept-involving behavior.²

Another reason to posit non-conceptual content turns on the idea that the action-controlling aspects of perceptual experiences require a more behaviorally oriented understanding than that provided by the Fregean notion of sense. For those with this motivation (e.g., Bermúdez, 1995, 1998; Bermúdez and Macpherson, 1998; Cussins, 1990, 1992; Davies 1992, 1997), it is important that non-conceptual content be considered more than just a primitive notion, that is, that it not be invoked merely to account for the contentful states involved in the behavior of pre-linguistic children and animals.

Among Evans' original reasons for introducing the notion of non-conceptual content was also the idea that the richness and grain of perceptual *experience* cannot be unpacked using the standard notion of belief. In fact, "belief," according to Evans, really names a " ... far more sophisticated cognitive state: one that is connected with ... the notion of *judgement*, and so, also, connected with the notion of *reasons*" (Evans, 1982, p. 124). The non-conceptual content of perceptual experiences need not be so directly connected with judgments. Yet, this does not imply that such contents are just informational states of sub-personal cognitive mechanisms. On the contrary, such contents are meant to be personal level experiential contents nevertheless, and, at least according to Peacocke, having an experience involving such-and-such non-conceptual content can constitute a reason, in fact, a *good* reason for an agent's believing something (cf. Peacocke, 1992a, p. 80).

How are we to resolve this appearance of contradiction? First, Peacocke defends the idea of non-conceptual content having associated correctness conditions. As such, he argues, non-conceptual content is genuine content, and not just a type of sub-personal information state. However, while the correctness conditions associated with any given concept are answerable to the norms that govern the use of the concept, the correctness conditions associated with non-conceptual contents need not be taken as answerable to anything else. They are self-evident; they are not inferred from other principles. They are, as Peacocke uses the term, *primitively compelling* (cf. Peacocke, 1992a, p. 6).³ The content of perceptual experiences is thus not revisable on the basis of any other inferential or evidential relations (cf. Crane, 1992). Using standard cognitive psychology terminology, we

can say that the content of perceptual experiences is *cognitively impenetrable*. Unlike, for example, the content of a belief, the content of a perceptual experience is independent of what we might *think* about it, in the same sense in which perceptual illusions retain their character even after we've learnt that they are indeed an illusion (cf. Fodor, 1984).

The second, and more important step toward resolving the alleged contradiction consists in moving from the simple idea that non-conceptual content has associated primitively compelling correctness conditions to considering them *a subject's reasons for believing* something. Thus Peacocke, while talking about positioned-scenario content —a kind of non-conceptual content (see below)—claims:

By perceiving the world, we frequently learn whether a judgement with a given conceptual content is true or not. This is possible only because a perceptual experience has a correctness condition whose holding may itself exclude, or require, the truth of a conceptual content. Some conceptual contents are actually individuated in part by their relations to those perceptual experiences that give good reasons for judging those contents. (Peacocke, 1992a, p. 66)

However, as McDowell points out, to claim that there are correctness conditions associated with non-conceptual contents does not yet imply that such correctness conditions constitute a *subject's reasons for believing* something (cf. McDowell, 1994, p. 163). The key move here seems to be the distinction between a reason's being the reason *for which* a subject believes or does something and a reason's being the

reason *why* a subject believes or does something. This is the example McDowell uses to illustrate such a distinction.

Consider, for instance, the bodily adjustments that a skilled cyclist makes in rounding curves. A satisfying explanation might show how it is that the movements are as they should be from the standpoint of rationality: suited to the end of staying balanced while making progress on the desired trajectory. But this is not to give the cyclist's reasons for making those movements. The connection between a movement and the goal is the sort of thing that *could* be a reason for making the movement, but a skilled cyclist makes such movements without needing reasons for doing so. (1994, p. 163)

What McDowell suggests is that, in granting that non-conceptual content has associated correctness conditions, we only get to treat such a notion as being a reason *why* a subject does something, that is, in McDowell's example, the kind of reason that sustains an explanation of the subject's behavior from the point of view of a rational external observer. However, that in itself does not guarantee that the subject's action is done *for that reason*, or even that the subject has any reasons at all!

Peacocke's move from the idea that protopositional (non-conceptual) content has associated primitively compelling correctness conditions to the idea that such a content constitutes reasons for action thus stands in need of further justification. Our final verdict will be that the invocation of protopositional content is a mistaken solution to a very real problem, the problem of *entry into the space of reasons*. A better candidate, and one which accommodates McDowell's intuition about the possibility of acting without reasons, is developed later in this paper.

Entry into the Space of Reasons

In developing an account of the representational content of experience, Peacocke introduces two different kinds of non-conceptual content. The first kind, considered the most fundamental of all, is *positioned-scenario content*. This kind of non-conceptual content is introduced to explain the individuation of first-person concepts. It is determined by “specifying which ways of filling out the space around the perceiver are consistent with the representational content’s being correct” (Peacocke, 1992a, p. 61). This space, considered a *type*, is in turn determined by the labeled fixing of an origin (usually one of the perceiver’s bodily parts, e.g., the center of the chest) and axes (e.g., directions with respect to the center of the chest). Once the origin and axes have been labeled and the time of the experience has been assigned, we can talk of a *positioned-scenario*. Positioned-scenario—the content of the perceptual experience itself—is to be distinguished from mental representations of such a content, since it is clearly the case that the same positioned-scenario can be the content of different mental representations (cf. Peacocke, 1992a, p. 65).

Essential to the notion of positioned-scenario content is the idea that the perceiver need not be aware of her bodily parts nor have the concepts that an observer may use in specifying the labeled ways of “filling out” the space around her. Having a perceptual experience with such-and-such positioned-scenario content does not yet present the world to the subject as being thus-and-so. In order

to describe the positioned-scenario content of a perceptual experience, the external observer may use concepts, “[b]ut it is crucial to observe that the fact that a concept is used in fixing the scenario does *not* entail that the concept itself is somehow a component of the representational content of the experience, nor that the concept must be possessed by the experiencer” (Peacocke, 1992a, p. 68). A perceiver may have a perceptual experience whose positioned-scenario content is, for example, the shape of a table being a certain way —let’s say, square— at a given time and from a particular angle without seeing the table *as* square, without the concept “square” being a component of her experience, and without having the concept “square” herself.

Positioned-scenario content thus fails to exhaust all the possible layers of complexity of human perceptual experience. To account for the representational content of a perceptual experience in which the world is indeed presented to the perceiver *as* being thus-and-so (e.g., square), we now seem to require a second layer of non-conceptual representational content. This is what Peacocke calls “protopositional content.” Unlike positioned-scenario content, the protopositional content of perceptual experience does not depend on fixed origins and axes. The protopositional content of, for example, my perceptual experience of seeing a table in front of me *as* square remains the same even when I see the table from a different angle, or in a new orientation, or at a distance, while the positioned-scenario content would vary accordingly. Perceptual experiences with different protopositional contents may thus have the same positioned-scenario content. Think, for instance, about the famous duck/rabbit image. When we see the figure *as* a rabbit, the protopositional content of that perceptual

experience is different from the (protopositional) content of the perceptual experience in which we see the figure *as* a duck, but the positioned-scenario content—as Peacocke understands it—is the same in both cases.

As in the case of positioned-scenario content, the fact that we mention certain concepts (e.g., “square”) in fixing the protopositional content of a perceptual experience does *not* yet require that the experiencer already possess those concepts. Protopositions contain objects and properties, not the concepts thereof. However, protopositional (non-conceptual) content attributions do require that the perceptual experience feed into a concept-possessing system. The line between non-conceptual and conceptual content is thus very fine, since the claim seems to be that the content of perceptual experience is protopositional only when the conditions for concept-possession are already in place. As a result, protopositional (non-conceptual) content occurs, on this account, only in organisms which can already display conceptual thought in other areas of cognition. As already stated, a perceptual experience with such-and such protopositional content has associated correctness conditions which, according to Peacocke, provide *primitively compelling reasons* for the formation of a belief based on such a perceptual experience (cf. 1992a, p. 7). Perceptual experiences have protopositional content because, although themselves non-modifiable by way of reasons, the experience itself is judged to be of the kind which *could* be the object of some reasoned articulation.

It thus looks as though the notion of protopositional content is Peacocke’s attempt to find what we may call an *entry point into the space of reasons*. Protopositional (non-conceptual) content is indeed content, because it has

associated correctness conditions and such conditions are taken to constitute reasons for actions. But it is non-conceptual because the subject need not possess the concepts that an external observer may use to describe the experience.

For McDowell, by contrast, if the content of our perceptual experiences is genuine at all, it has to be subject to the operations of the creature's spontaneity, and has therefore to be already conceptual. There cannot be a distinction between "reasons for which a subject thinks as she does and reasons she can give for thinking that way" (1994, p. 165). According to McDowell, the richness of our perceptual experiences is best accounted for, not in terms of their possessing non-conceptual contents, but by appeal to the notion of *demonstratives*. Non-conceptual content plays a role *only* at the level of our *cognitive machinery*:

I am not saying that there is something wrong with just any notion of non-conceptual content. It would be dangerous to deny, from a philosophical armchair, that cognitive psychology is an intellectually respectable discipline, at least so long as it stays within its proper bounds. And it is hard to see how cognitive psychology could get along without attributing content to internal states and occurrences in a way that is not constrained by the conceptual capacities, if any, of the creatures whose lives it tries to make intelligible. But it is a recipe for trouble if we blur the distinction between the respectable theoretical role that non-conceptual content has in cognitive psychology, on the one hand, and, on the other, the notion of content that belongs with the capacities exercised in active self-conscious thinking —as if the contentfulness of our thoughts and conscious experiences could be understood as a welling-up to the surface of some of the content that a good psychological theory would attribute to goings-on in our cognitive machinery. (1994, p. 55)

McDowell's position is thus that *either* non-conceptual content must be restricted to the sub-personal level *or* there can be two kinds of reasons for a subject's behavior. Since reasons are all of the same kind, McDowell claims, there cannot be an intelligible notion of non-conceptual content at the personal level. The reasons for a subject's acting the way she does, even when prompted by a particular perceptual experience, can still be conceptual albeit in a quite minimalist way: a demonstrative way (cf. 1994, pp. 162-174). Given a perceptual experience, the subject can simply articulate its content by saying that something looks *that* way. The reasoning-weighted role of demonstratives in a subject's explanation of her perceptually-based representational states seems thus to be crucial for the justification of the conceptual character of the content of perceptual experience. McDowell's account thus offers no *entry point* into the space of reasons. It is reasons all the way down.

Peacocke's (1994) argument against this view is based on the idea that demonstrative-perceptual concepts are just *too* fine-grained to be considered appropriately applicable to perceptual experiences, since there are many kinds of properties and objects which can be the referents of a demonstrative, yet which cannot themselves be discriminated by perceptual means. We shall not pursue this analysis of the role of demonstratives any further. But the picture we want to draw emerges more clearly if we keep these remarks in mind.

The point we wish to highlight is that McDowell's view of thought as conceptual through and through leaves us with no understanding of how it is that rational animals know *when* to reflect on their reasons for belief and action. Yet Peacocke's notion of primitively compelling correctness conditions (associated with protopositions), although designed to establish a bridge between perceptual-

based belief formation and the conceptual realm, does not seem to locate the reasons for action within the subject. They belong to the external observer's explanation of the subject's behavior.

It will be argued that the space of possible links between perceptions, reasons, and actions is larger and more complex than either Peacocke or McDowell here allows. In the case of McDowell, being able to conceptually articulate the reasons for a particular behavior and/or belief is a necessary, but not sufficient condition for another, equally important, aspect of our cognitive life, namely, being able to modify such a behavior/belief when things go astray. McDowell acknowledges that "being at home in the space of reasons involves the standing potential for a reflective stance at which the question arises whether one ought to find this or that persuasive" (1994, p. 125). For McDowell, this ability can only be displayed when other —mostly non-conceptual— skills are already in place. McDowell's reflective stance only makes sense if the thinker is able to recognize a situation as reflection-demanding. But the ability to detect such a property cannot be, on pain of circularity, reflective itself.

Peacocke's position is slightly more complicated. Here the content of a perceptual experience is considered protopositional only if it feeds into a system which already possesses concepts. Interestingly, the addition of such an extra cognitive ingredient seems to weaken Peacocke's attempt to use protopositional (non-conceptual) content as the desired point of entry into the space of reasons. First, although less central to the current discussion, there looms a certain *explanatory gap* between the semantic properties of protopositional content and the normative status of the type of correctness conditions associated

with it. For what is required in order to allow for genuine protopropositional content attributions —according to Peacocke— is that the creature be *disposed to act* (*ceteris paribus*) in a self-reflective way regarding its perceptual experiences. The notion of primitively compelling correctness conditions thus reveals itself as a dispositional notion. As such, the normative status of Peacocke’s notion of primitively compelling correctness conditions has a *causal* —rather than *constitutive*— character.⁴ However, the claim that the individuation of the correctness conditions associated with the content of a perceptual experience is based on the creature’s finding the content of the experience *primitively compelling* —taken as a dispositional claim— cannot have direct *semantic* implications regarding the *nature* of the content itself.⁵ Secondly, the appeal to the existence of a putatively concept-possessing system waiting to use the protopropositional contents does not yet guarantee that the system be able to spot situations in which to invoke the full apparatus of conceptualization and reasoning. This argument will be developed further in the remainder of this paper.

To better accommodate the complexity of the possible links between perception, reason, and action, we shall next propose a finer —and hopefully more fruitful— carving of the normative space in which to locate perceptual representational states. The idea is to show that some (but not all) of our perceptual experiences are full-blooded, semantic representational states only *because* other —mostly non-conceptual skills— are in place and automatically act so as to prompt us to enter the space of reasons. An added layer of complexity in perceptual experience —one not captured by the simple conceptual / non-conceptual

dichotomy— is thus revealed. In order to begin to justify these claims, let's next introduce the concept of semantic responsibility.

Semantic Responsibility

The notion of *semantic responsibility* has its anchor point in a deontological conception of *epistemic* justification. The defining characteristic of “deontological” in this context is the idea that epistemological terms such as justification, evidence, warrant, etc. are best understood in some strongly normative fashion, that is, as relating to notions of requirement, duty, blame, obligation, and the like: being justified in believing that *p* is a matter of fulfilling certain epistemic duties so that we tend to accept only those beliefs that are most likely to be true. To reach what Chisholm (1982) calls “positive epistemic status” is for an agent to fulfill a certain epistemic responsibility, that of trying achieve a certain state of intellectual excellence by bringing about a situation in which one's beliefs are mostly true.⁶

The aim in this paper is not, however, to contribute to the debate on epistemological justification. Instead, our suggestion is that certain responsibilist ingredients may play an important role not (or perhaps not only) in the epistemic justification of some mental states, but in their individuation. The notion of semantic responsibility involves an agent's duties regarding particular cognitive states. A subject fulfilling these duties can engage *both* in inferential and non-inferential practices in which she is prompted to corrections by the provision —by

herself or others— of reasons and arguments. A characterization of semantic responsibility can thus be provided along the following lines. A subject *S* is semantically responsible for a cognitive state *p* iff the following three conditions are met: (i) *S* is a well functioning cognitive system, (ii) *S* could reasonably have been expected to fulfill her critical duties relative to *p*, and (iii) *S* is skilled at detecting the kind of situations in which those critical abilities should be deployed.⁷ To fulfill one's critical duties is, broadly speaking, to appreciate the various commitments that accrue as a result of taking the world to be thus-and-so, and to be ready to engage in evidence and reason-weighting activity as a result. When these three conditions (for semantic responsibility) are met, we can talk about the representing as *mindful*. Failure to meet any of the three conditions results in what we shall call *mindless* representations.

Condition (iii) needs to be spelled out a bit further. The idea is that a crucial part of what makes any contentful state count as *mindful* is a subject's background ability to spot the kinds of situation in which she needs to step back and critically appraise her first "impressions" rather than simply trusting her automatic responses. The process of *spotting* these situations must, however (on pain of regress) itself be automatic. The "care-needed" situations must simply "pop-out" to the expert, that is, they must directly present themselves as requiring care or caution without any further reflection or intellectual analysis. Think, for example, of the air traffic controller who must learn when *not* to trust her otherwise automatic reactions to a radar screen pattern. This crucial cognitive skill, a kind of "critical pop-out," has been unjustly neglected in both philosophical and cognitive scientific discussion.⁸ Notice, also, that the notion of "critical pop-out" is not to be

assimilated to the much simpler notion of (if you like) “novelty pop-out.” Novel situations do, to be sure, tend to engage our critical and reflective faculties. But the kind of expertise envisioned here is often best displayed in situations with which we are highly familiar.

What carving of normative space results from applying this notion of semantic responsibility to perceptual experiences? At the bottom level there will be the representational products of sub-personal cognitive functioning. These can be considered a special kind of *mindless* representational state. Think, for instance, of peripheral sensory neurons like the rods in the retina. The amount of neurotransmitter these cells release varies as a result of changes in light intensity. It would not be unreasonable to claim that these neurons represent the property of there being a certain temporal and spatial density of photons at a certain retinal location. Certain perceptual contents concerning the intensity of ambient light could thus be considered sub-personal representational states. The agent is not semantically responsible since she cannot alter these voltage discharges by way of reason, nor is she able to spot the kind of situation in which such corrections would be advisable.

One way to display the specificity of this category is to think in terms of how to characterize errors. When the state we are considering is of this sub-personal (mindless) type, to make a mistake seems to be tantamount to a malfunction in the system. Sub-personal (mindless) content is thus the version of non-conceptual content that both Peacocke and McDowell acknowledge as scientifically respectable. Notice, however, that this kind of mindless content cannot play the role that Peacocke assigns to the notion of protopositional (non-conceptual)

content at the personal level. The reason for this is that the normative status of sub-personal (mindless) content does not leave room for an important distinction: the distinction between a representation being in error as a result of some malfunction in the system and as a result of normal functioning in a system that nonetheless actually makes mistakes (see, e.g., Haugeland, 1998, p. 310).⁹

At the intermediate level there are representational states that are available to the agent (i.e., there are personal level cognitive states) for which she bears no semantic responsibility. The specific nature of such mindless representational states depends, however, on whether the condition that is not met is (ii) or (iii). Examples involving perceptual illusions are the obvious paradigm cases when the missing condition for semantic responsibility is the second one, that is, when a properly functioning agent lacks the ability to learn by exposure to reasons to modify the content of her perceptual experiences.¹⁰ Perceptual illusions are the kind of situations in which, despite normal functioning, the person accrues no semantic responsibility for the content of her perceptual experience.

If due to some optical illusion, I fail to discriminate the twelve steps down from my house (perhaps seeing only ten) and I fall, my *mistake* need not be the result of anything going wrong at the level of my physical or physiological functions. My visual system and neural mechanisms may all be working just as they are supposed to, yet (under these specific ecological conditions) they fail to deliver a veridical representation of the number of steps, much as a perfectly functioning heart may (under extreme pressure) fail to pump sufficient blood. In such a case it is surely correct to say that I *misrepresented* the number of steps, but there is nothing that I should have done to avoid the mistake. I can in future try to

pay more attention, be more careful —perhaps marking this situation as one where greater care is called for— but even with all the care in the world, I cannot rule out the possibility of yet another such optical illusion. This is a *mindless* way of representing. It is mindless because the subject cannot be expected to learn to perceive better.

The normative status of perceptual experience, in general, is usually taken to involve *no* personal level cognitive duties of the kind just mentioned. As Crane (1992, p. 154) says: “to perceive that *p*, there are no other *perceptions* that you *ought to have* ... (A Davidsonian might put this point by saying that perception is not subject to the ‘constitutive ideal of rationality).” However, as we shall argue, at least some perceptual experiences are indeed subject to this constitutive ideal. The resemblance between the kind of mindless cognitive states just described —one in which the second condition for semantic responsibility is not met— and the standard notion of non-conceptual content does not imply that a mindless cognitive state is just a state whose content is non-conceptual (but personal). The mindful / mindless distinction cross-classifies cases relative to the traditional conceptual / non-conceptual content scheme, turning the overall classification into a four-dimensional one.

At the top, there are representational states (both conceptual and non-conceptual) which are consciously available and toward which the agent may be expected to fulfill her cognitive duties —these are the *mindful* representational states. When in error, if the representation is mindful, there is something like *culpable error*: the mistake is in a deep sense a *cognitive* mistake. By that we mean not only that the error belongs, or could belong, to the stream of consciousness but

also, and much more importantly, that we could have expected the agent to do better. It is this crucial, but difficult feature —the feature of *culpable* cognitive failure— that characterizes the cases of representation that we call *mindful*. Some of our perceptual experiences fall into this category without thereby being conceptual representational states.

Perceiving Well

We turn now to the discussion of two scenarios which will help show how our taxonomy cross-cuts the conceptual / non-conceptual divide. The first attempts to show that there are *mindful* perceptual experiences whose content is not conceptual. The second aims to display a case of *mindless* perceptual experience whose content is not non-conceptual.

It may be helpful to rehearse the three conditions for semantic responsibility. A subject S is semantically responsible for a mental representation *p* if and only if S is a properly functioning cognitive system, S could reasonably be expected to critically appraise *p*, and S is able to detect the kinds of situation in which such critical appraisal is called for. The first scenario is one in which all three conditions for semantic responsibility are met. Representational states meeting all three conditions of semantic responsibility are *full-blooded*, genuinely semantic or, in my terminology, *mindful* representational states. Think, for instance, of a teacher perceiving boredom on the faces of her students. The experience might be quiet subtle, something difficult to “put in words,” but any teacher who has been

involved in pedagogical tasks for some period of time will have come to recognize that particular look. And at that moment, one knows to engage one's critical and reflective faculties so as to (hopefully) *correct* that situation. This type of perceptual experience, we contend, involves a fully *mindful* representational state: one for which the subject is indeed semantically responsible. Notice that the teacher is skilled at recognizing the situation, thus perceptually presented, as one which requires some correction of her own activity, and is capable, in considering how to put the situation right, of listening to reasoned argument or suggestions.

Such a description, however, makes it sound as if what makes the perceptual experience *mindful* is simply (à la McDowell) this engagement with the space of reasons. But this is too quick, for it is not just that the subject —the teacher— *takes in* this situation *as* being thus-and-so, and exploits it “in active thinking, thinking that is open to reflection about its own rational credentials” (1994, p. 47). It is also, and more importantly, that the teacher should be the kind of agent for whom a perceptual experience of this kind *pops-out* as requiring the engagement of certain faculties. As explained earlier, if such an account is not going to become circular, this “pop-out” must be viewed as an automatic process, something entirely supported by sub-personal goings-on. The perceived need for critical engagement is thus as *automatic* as the experience itself. What *makes* a representational state *mindful*, on this account, is thus in part the relations it bears to bodies of (sub-personal) expertise that enable the agent to know *when to worry*.¹¹ But those enabling relations are indeed present even in some of our non-conceptualized perceptual experiences, which thus belong to the same (mindful) category as other

mental representations available to the subject *via* more conceptualized modes of presentation.

Now, it could be argued that the teacher's perceptual experience counts as an experience of *boredom* only because the experience is poised to act as input to a concept-exercising system. Such a characterization recalls Peacocke's depiction of protopositional content, as well as Evans' original story above. One may be tempted to say that the experience is presented to the teacher as being a (non-conceptual) experience of her students' boredom.

The important point of difference here, as before, lies in the automatic, sub-personal nature of the process of "critical pop-out" itself. This notion has no analogue in the accounts of Peacocke or Evans (or, indeed, McDowell). It is, however, the natural response to the requirement for an entry point into the space of reasons. For without some such mechanism in place, the realms of experience and reason remain uncoordinated, and the potential for rational engagement cannot be efficiently realized. Peacocke, Evans, and McDowell, all concentrate, in their different ways, on the conceptual abilities that render perceptual content genuinely contentful. But in so doing they miss the deep and important sense in which this crucial consummation itself depends on an unremarked substrate of sub-personal activity, the presence or absence of which affects the functional poise of the experience itself. Our claim is that this functional poise enters into the individuation of the perceptual experience itself. As characterized here, the idea that *some* (but by no means all) perceptual experiences involve *mindful* representational states thus emerges as different from both McDowell's *conceptual*

characterization of all perceptual experiences and Peacocke's depiction of perceptual experience as the input to a concept-exercising system.

Notice finally that there will be perceptual experiences (and representational states in general) for which the subject is not semantically responsible (i.e., *mindless*), which count as conceptual according to the original distinction. Thus consider a *novice* air traffic controller. She sees the radar screen, knows the rules, guidelines and procedures, and knows that certain situations will require prompt action (telling one plane to climb, another to bank, and so on). But unlike the *expert* air traffic controller, she has yet to tune her cognitive-perceptual capacities so as to allow such care-needed situations to *pop-out* from the ongoing display. She is fully capable of the required kinds of critical reflection and judgment, and has (intuitively) the right perceptual experiences. But they remain *mindless*, because they do not automatically activate her critical faculties. She is, *qua novice*, not culpable for failing, in real time, to pursue a critical agenda when certain perceptual experiences obtain.

To say that the novice air traffic controller's experiences are mindless is not equivalent to considering the content of such experiences as non-conceptual. The subject in this example is able to engage in critical reflection and even judgments about the on-going experiences and the actions that the experiences trigger. What makes this case an example of a mindless representational state is that the subject's experiences are not marked as calling for further critical appraisal of her own actions. The *novice* air controller's perceptual experiences, at one level (perhaps at the level of positioned-scenario content?), may be the same as those of an *expert* air

traffic controller. But in another, more cognitively interesting sense, they are surely quite different.

The essential point is that, in order to enter the space of reasons, a creature should not only be able to engage in some sort of critical evaluation of her perceptual experiences, but also, and importantly, be able to *spot* those experiences which stand in need of such critical reflection. This latter requirement, on pain of circularity, cannot be interpreted as a demand for additional conceptual cognitive skills. Such critical pop-out must be as automatic a phenomenon as the perceptual experience itself.

The discussion between the friends and foes of non-conceptual content is marked by a failure to appreciate the varieties of perceptual experience. Using functional poise as a tool for individuating experiences, and incorporating the additional dimension of critical pop-out, we can now divide perceptual experiences into (at least) four types: those conceptualized and poised-to-enter-the-space-of-reasons; those unconceptualized yet poised-to-enter-the-space-of-reasons; those conceptualized and not thus poised; and those unconceptualized and not thus poised. Marking those possibilities, and drawing attention to the additional problems raised by the need to *know when to reflect* (and when not to reflect), is a small first step, we hope, toward a better understanding of the complex relations between perception, experience, and reasoned action.

References

Armstrong, D. (1969). *A materialist theory of the mind*. London: Routledge and Kegan Paul.

- Axtell, G. (1997). Recent work on virtue epistemology. *American Philosophical Quarterly*, 34(1), 1-26.
- Bermúdez, J. L. (1994). Peacocke's argument against the autonomy of nonconceptual content. *Mind and Language*, 9, 402-418.
- Bermúdez (1995). Nonconceptual content: From perceptual experience to subpersonal computational states'. *Mind and Language*, 10, 333-369.
- Bermúdez, J. L. (1998). *The paradox of self-consciousness*. Cambridge, Massachusetts: The MIT Press.
- Bermúdez, J. L., and Macpherson, F. (1998). Nonconceptual content and the nature of perceptual experience, *Electronic Journal of Analytic Philosophy*, 6, 2.
- Bermúdez, J. L. (1999). Naturalism and conceptual norms, *The Philosophical Quarterly*, 49(194), 77-85.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge: Harvard University Press.
- Chisholm, R. (1977). *Theory of knowledge*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.
- Chisholm, R. (1982). *The foundations of knowing*. Minneapolis: University of Minnesota Press.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46, 279-295.
- Crane, T. (1988). The waterfall illusion. *Analysis*, 48, 142-147.
- Crane T. (1992). The non-conceptual content of experience. In T. Crane (Ed.), *The contents of experience: Essays on perception* (pp. 136-157). Cambridge: Cambridge University Press.
- Cussins, A. (1990). The connectionist construction of concepts. In M. Boden (Ed.), *The philosophy of artificial intelligence* (pp. 368-440). Oxford: Oxford University Press.
- Cussins, A. (1992). Content, embodiment and objectivity: The theory of cognitive trails. *Mind*, 101(404), 651-688.
- Davis, M. (1991). Individualism and perceptual content. *Mind*, 100, 461-484.
- Davis, M. (1992). Perceptual content and local supervenience. *Proceedings of the Aristotelian Society*, 92, 21-45.
- Davis, M. (1997). Externalism and experience. In N. Block, O. Flanagan, and G. Güzeldere (Eds.), *The nature of consciousness: Philosophical debates* (pp. 309-327). Cambridge: MIT Press.
- Davidson, D. (1970). Mental events. In L. Foster and J. W. Swanson (Eds.), *Experience and theory* (pp. 79-102). London: Duckworth.

- DesAutels, P. (1996). Gestalt shifts in moral perception. In L. May, M. Friedman, and A. Clark (Eds.), *Minds and morals: Essays on cognitive science and ethics* (pp. 129-143). Cambridge: MIT Press.
- Evans, G. (1982). *The varieties of reference*. J. McDowell (Ed.). Oxford: Clarendon Press.
- Fodor, J. (1984). *The modularity of mind*. Cambridge: MIT Press.
- Fodor, J. (1988). *In critical condition: Polemical essays on cognitive science and the philosophy of mind*. Cambridge, Massachusetts: MIT Press.
- Frege, G. (1956). The thought: A logical inquiry. Translated by A. M. Quinton and M. Quinton [Trans.]. *Mind*, 65, 289-311.
- Haugeland, J. (1998). Truth and rule-following. In J. Haugeland (Ed.), *Having thought: Essays in the metaphysics of mind* (pp. 305-361). Cambridge: Harvard University Press.
- Karmiloff-Smith, A. (1998). Development itself is the key to understanding developmental disorders. *Trends in Cognitive Science*, 2(10), 389-398.
- Jackson, F. (1977). *Perception*. Cambridge, England: Cambridge University Press.
- McDowell, J. (1994). *Mind and world*. Cambridge, Massachusetts: Harvard University Press.
- Millikan, R. (1984). *Language, thought and other biological categories. New foundations for realism*. Cambridge, Massachusetts: MIT Press.
- Millikan, R. (1989). In defense of proper functions, *Philosophy of Science*, 56, 288-302.
- Neander, K. (1991). The teleological notion of "function". *Australasian Journal of Philosophy*, 69(4), 454-468.
- Papineau, D. (1990). Truth and teleology. In D. Knowles (Ed.), *Explanation and its limits* (pp. 21-43). Cambridge, England: Cambridge University Press.
- Peacocke, C. (1986). Analogue content. *Proceedings of the Aristotelian Society. Supplementary Vol. 60*, 1-17.
- Peacocke, C. (1992a). *A study of concepts*. Cambridge, Massachusetts: MIT Press.
- Peacocke, C. (1992b). Scenarios, concepts and perception. In T. Crane (Ed), *The contents of experience: Essays on perception* (pp. 105-135). Cambridge, England: Cambridge University Press.
- Peacocke, C. (1998). Nonconceptual content defended. *Philosophy and Phenomenological Research*, 58(2), 381-388.
- Sellars, W. (1956). Empiricism and the philosophy of mind. In H. Feigl and M. Scriven (Eds.), *The foundations of science and the concepts of psychoanalysis: Minnesota studies*

in the philosophy of science, Vol. I (pp. 253-329). Minneapolis, Minnesota: University of Minnesota Press.

¹Wilfrid Sellars introduced the notion of a “space of reasons” in his seminal paper “Empiricism and the Philosophy of Mind” as follows: “In characterizing an episode or a state as that of *knowing*, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says” (Sellars, 1956, pp. 298-299). We make extensive use of Sellars’ notion in this paper to distinguish the domain of rationality and normativity from the domain of empirical description.

²Or so it seems. The project only works if the correctness conditions associated with non-conceptual content are themselves *naturalistically respectable*. Some doubts about whether they are indeed so can be found in Bermúdez (1999).

³The idea of non-conceptual content having *primitively* compelling correctness conditions is not devoid of complications (see e.g. Bermúdez, 1999). See also below.

⁴For an insightful and extremely elegant argument in which this option is presented as one horn of a dilemma for Peacocke’s position, see Bermúdez (1999).

⁵Fodor uses a similar argument in a completely different context: that of (surprisingly) defending functional role semantics against some (according to him) unfair criticisms. Fodor’s idea is that the functional individuation of e.g. hearts as devices that pump blood does not involve any semantic thesis about the word “heart” or the concept HEART, or the possession conditions for that concept (cf. Fodor, 1998, Chapter 6, esp., pp. 71-72).

⁶For deontological approaches to the notion of epistemic justification see, for instance, L. Bonjour (1985), R. Chisholm (1977, 1982) and S. Cohen (1984). For a very detailed map of current positions, especially of what it is known in the literature as “virtue epistemology” see G. Axtell (1997).

⁷This notion of proper functioning is not meant to be contentious. Subtle differences regarding this notion, such as the ones between, e.g., Millikan (1984, 1989), Neander (1991), or Papineau’s (1990) accounts, although playing an important role in discussions of the possibility of a teleobiological account of semantic properties, do not matter for present purposes.

⁸An exception includes DesAutels (1996) account of learning to perceive moral situations.

⁹Of course, one can imagine a case in which a system *malfunctions* in a way that *causes* it to misrepresent. For example, if a sudden burst blood vessel caused double vision resulting in a fall. Thus certain malfunctions may be sufficient for *mindless* (but personal-level) misrepresentation to occur.

¹⁰Evans (1982, p. 123) explicitly mentions perceptual illusions as supporting evidence for the non-conceptual character of perceptual experiences. Thus a subject who already knows about the Müller-Lyer illusion, believes (mindfully believes) the lines to be of equal length. The fact that the subject can’t help but still perceive one line as longer than the other is distinctive, Evans suggests, of a more primitive informational state with a different *non-conceptual* content. In Evans’ treatment of this kind of case, as in my example, the subject can’t help but believe that the lines look the way they look. No amount of learning would stop her from having the experience of seeing one line as longer than the other. If prompted, she might give us a detailed description of her experience and even admit that she just can’t stop herself from seeing the lines as being different, regardless of her knowledge of the illusion. There is thus no room for reason-sensitive tuneability of the *experience*, even though her *actions* could be tuned by e.g.

learning to ignore the appearance. This is one of the cases where the notions of mindless cognitive state and non-conceptual content coincide. There will, however, be cases where this is not so (see next Section).

¹¹Such expertise comes in degrees, and so, in consequence, do the notions of mindfulness and mindlessness as applied to mental states.