

The Nonconceptual World[†]

Brian Cantwell Smith*
University of Toronto

Over the past twenty years, the notion of non-conceptual content has played a prominent role in philosophical discussions of the relation between thought and language, on the one hand, and perception and action, on the other. In spite of its importance, however, the nature of non-conceptual content remains remarkably obscure. The fact that it is negatively defined is not helpful. Just about the only thing on which both proponents and detractors of nonconceptual content agree (beyond the fact that nonconceptual content is not conceptual content—though as we will see it is very likely that they disagree on what conceptual content is) is that the predicate ‘non-conceptual’ should be understood *epistemically*, as a predicate on content-bearing mental states.

What distinguishes nonconceptual content, it is thought, is the structure of the belief, or the attitude of the believer—not the world thereby believed in. I will argue that this purely epistemological focus is mistaken, in the following sense: that the phenomena that have driven at least many advocates of a notion of non-conceptual content to embrace the notion are not, in the end, best understood from an epistemological point of view. Instead, I claim, the character of and ultimate warrant for nonconceptual content is ontological.

There are two parts to the claim, to put it most baldly. First,

[†]...

*Coach House Institute, Faculty of Information, University of Toronto
90 Wellesley St W, Toronto, Ontario M5S 1C5 Canada

[†]© Brian Cantwell Smith 2009

Last edited: November 12, 2009

Please do not copy or cite

Comments welcome

Draft only (version 0.60)

brian.cantwell.smith@utoronto.ca

what is explanatorily fundamental about nonconceptual content, I will argue, is neither how it is used, nor the epistemic role it plays in the agent, but *how it takes the world to be*. And second, the world as truth-maker for such nonconceptual content bearing mental states is ...

Any distinctive epistemological characteristics of nonconceptual content is a consequence of that ontological commitment.

I will also argue that nonconceptual content is, in a specific sense, more accurate—more faithful to the world in detail—than conceptual content. This raises problems for the analysis of conceptual content, including for issues of realism and truth. Conceptual content, I argue, involves a form of abstraction, which in turn ties both semantic issues of truth and reference and ontological issues of objects and properties to dynamical human norms. At the same time, the story remains fundamentally realist, illuminating the metaphysical ground underlying the intimate relation among perception, thought, and action.

The overarching theme of the investigation is that epistemic issues of experience, representation, and thought, and semantic issues of truth, reference, and content cannot be solved without tackling fundamental metaphysical questions about the nature of objects, properties, relations, and the founding world. Only by understanding the mind against this metaphysical ground, and investigating both conceptual and nonconceptual representation's ontological commitments can we understand what these various kinds of content are like, what these various kinds of content are for.

1 The Nonconceptual Content Debate

The notion of non-conceptual content has played a major role in recent discussions about the relation between thought and language, on the one hand, and perception and action, on the other. Yet in spite of its importance, the underlying nature of nonconceptual content has resisted trenchant analysis. The fact that it is negatively defined does not help. It has also been recruited in diverse ways: some writers focus on nonconceptual content in judgment or thought, often perceptually based; others, on forms of intentional content that play a role in action. Some have asked

Nonconceptual World

whether nonconceptual content can be exhibited in a creature without conceptual skills at all, or whether instantiation of mental states with nonconceptual content requires a prior or concomitant mastery of concepts. Across this range of issues there is no doubt that the literature contains a number of intriguing insights and provocative proposals. Yet it is probably fair to say that nothing approaching a comprehensive theory of nonconceptual content has yet been presented.

The aim of this paper is to redress that situation. I will start by considering the case of non-conceptual content in judgment (including perceptual judgment): judging that the world is a certain way-but not a way that finds ready or even possible expression in conceptual form. Later I will extend the analysis to action, but even in the case of judgment substantial theoretical issues arise, in need of disentangling.

I will start by considering two major lines of argument: Gareth Evan's classic defense of nonconceptual content (*Varieties of Reference* [VOR]), and John McDowell's now equally classic rebuttal (*Mind and World* [M&W]). Evans and McDowell are talking past each other, I will argue—and therefore (to put it a bit anachronistically) are missing each other's point. Their views are held together only by an implicit assumption—an assumption it will be helpful to put on centre stage, in order to subject to challenge.

1a Evans vs. McDowell

Though Evans' avowed concern was with reference, truth, and objectivity, in fact he was almost equally concerned with the epistemic role that concepts and conceptions played in the mind of agents. As is especially evident in parts II and III of VOR,¹ for Evans a semantic account of a concept involved, among other things, *explicitating the role that the concept played in the mental life of a person who possessed*. His discussion of indexical concepts, for example, and his recruitment of notions of information delivered by the senses, are as much an analysis of their role in cognitive activity as they are of traditionally semantic issues such as truth. Indeed, one of the signal contributions of VOR is Evan's attempt to show how reference is achieved, not just what reference is.

¹Unless otherwise indicated, all Evansian references are to VOR.

Evan's epistemic or even cognitive bent is well kept in mind in understanding his analysis of conceptual content. For him, conceptuality has first and foremost to do with the semantic or propositional form of a judgment. As codified in his **Generality Condition**, Evans takes a judgment to be conceptual just in case it is of the form a is F , and the agent not only thinks that a is F , but is also capable of entertaining the thought that b and c are F , for any b and c of which it has a conception, and that a is G and H , for any G and H of which it has a conception (modulo various appropriateness conditions). Thus conceptuality, according to Evans, consists of content with something of an algebraic or compositional constituent structure.

It is content of this (at least potentially) recombinant form, according to Evans, that figures in linguistically articulate judgments and propositional attitudes—i.e., that is conveyed by such embedded English sentences as that there is a war in Serbia, that the sun is rising, that three people are standing at the door.

For McDowell, in contrast, conceptuality has first and foremost to do, not with the internal structure or form of a judgment, but with the fit of the judgment into a overall conception of the world: not a complete conception of the world (which, if even meaningful, would be impossible for a finite agent to achieve), but, as it were, a conception of a complete world—a cognitive grasp on the world's being whole—exhaustive, entire, complete in all details. What leaves McDowell unmoved by cases of perceptual judgment that lead others to embrace a notion of nonconceptual content is his (correct, in my view) recognition that we understand even spectacularly diverse arrays of colour as part of the world of our experience-as located, objective parts of the one comprehensive reality.

These at least superficially different concerns—between the internal form of a judgment for Evans, and its role in undergirding our grasp of reality in toto, for McDowell—are betrayed in a number of passages, in M&W, that to an Evansian, might otherwise seem perplexing: his unproblematic embrace of indexical concepts and indexical judgment, such as “that red”. This is conceptual, à la McDowell, because, in any situation in which it is uttered, it plays an unproblematic role in our overall cognitive grasp of the world. For Evans, indexical judgments are not so obviously con-

Nonconceptual World

ceptual, even if they have the structural form a is F , because two different thinkings or utterances of 'that red' could have such different contents. If, like Evans, one is concerned to explicate the role that thoughts play in a rational agent's cognitive economy, saying that a thinking of 'that red' refers to a particulate shade (I will get to the issue of individuating shades presently) doesn't do half of the work that is required. The problem is that one such thinking may connote blood, occult ceremonies, and Burgundy wine; another, lipstick, Orlon sweaters, and Muffy at a prep school dance. For an Evansian, they are different judgments—and the differences matter, as regards conceptuality, because how they differ is not obviously conceptually explicable. For McDowell, of course, they are different judgments too—but with respect to the defense of conceptuality, the differences don't matter—because both locate their referents in the world, in a way that can be comprehended objectively, and because of that fact, are, in McDowell's sense, conceptual.

There are others cracks suggesting that Evans and McDowell may be focusing on different aspects of judgment, if not on outright different phenomena. One is McDowell's comment, in passing, about the problematic nature of just-noticeable-differences (JNB) in colour perception.²

The only conceivable argument that McDowell's criterion entails Evans' criterion—i.e., that a grasp of the world must consist in judgments all of which have what Evans would call conceptual form—rests in part on what is ultimately an ontological assumption—an extremely common one, yes, but not something the student of nonconceptual content should blithely assume:

- A** • The world is exhaustively constituted of objects exemplifying properties, standing in relations, configured in situations or states of affairs, and gatherable together in sets, and so forth, with at least some of those objects (the concrete ones) spatio-

²The phenomenon of a just-noticeable difference, or JNB, well-studied in psychology, arises in cases where, for example, of three shades x , y , and z , human subjects, even in excellent viewing circumstances, cannot judge that x and y are different, nor that y and z are different, but *are* confident in being able to see a difference between x and z . In this situation (of having a non-differentiable neighbour in common) x and z are said to exemplify a "just noticeable difference."

temporally located or related.

As is evident from the Generality Condition, it is just this sort of world that Evansian conceptual content represents. Content meeting the Generality Constraint is not only itself articulated, that is, and potentially recombinant; it also attributes an articulated and potentially recombinant structure to the world it represents. For issues of modal realism aside, most would agree that the objects, properties, and relations thereby represented are metaphysically distinct from each other in part because they, too, could have been differently combined. If a is F , then God—or another world, or even this same world, at another time and place—could have made it the case that b is F , or that a is G , for appropriate a , b , F , and G .

The picture is thus relatively clear. If A is true, and the world thus consists of objects exemplifying properties and constituting states of affairs (etc.), then it would be natural to conclude that McDowell's completeness requirement could be met by entertaining judgments that are conceptual in form, according to Evan's characterisation. This is not the realm of logical implication: nothing guarantees that even if A is true, all judgments of its being this way need satisfy Evan's condition; and conversely, nothing guarantees that meeting Evan's condition in and of itself need give the agent a comprehensive grasp of the world as whole. Rather, the point is that it is A that allows people to think that Evan's condition and McDowell's requirement characterise the same sort of conceptuality—and hence to conclude that McDowell and Evan disagree. For unless A is true, it is not clear that Evans and McDowell's construals of conceptuality are even compatible. If A is *not* true, in particular—as I am going to argue it is not—then an agent able to entertain only conceptual content in Evans' sense would not be able to meet McDowell's requirement: hence would not be able to have conceptual content in McDowell's sense at all.

My strategy, therefore, will be to take as conceptual any content that has this ontological character: that is, any content that represents the world in terms of objects, properties, and relations in the standard way. That is, I differ from Evans (and most other writers) in focusing: (i) neither on the articulate structure of the con-

Nonconceptual World

tent itself—i.e., qua sense, proposition, meaning, or other “intentional” entity; (ii) nor on the structure of the expression or representational vehicle; (iii) nor on the conceptual capacities of the agent—in the sense that the agent could rearrange the content, or the mental state that bears it, by substituting other pieces of appropriate type; but (iv) *on the structure of the world thereby represented*. This ontological focus does not automatically imply that conceptual content so defined is without distinctive epistemological character. Rather, what I want to argue is that conceptual content so defined is strong enough to entail Evans’ Generality Condition as a consequence.

For ease of discussion, I will extend standard usage and say that conceptual judgments take the world to be conceptual when they take it to consist of discrete, extended, concrete, reidentifiable objects, exemplifying properties, standing in relations, arranged in states of affairs, grouped together in sets, etc. That permits the following simple definition: conceptual content is content that takes the world to be conceptual.

I want to argue for this “ontologising” of the conceptual/non-conceptual distinction from considerations in cognitive science, my home field. I will have more to say about the specific character of cognitive science later; for now, it is enough to say that it is based on a broadly representational theory of mind, taking mental life to arise out of semantically-warranted causal transitions in material substrates. That is: the mind itself is taken to be physical instantiated (and hence, in a sense, to be a physical mechanism), but nevertheless to be distinctive, among physical mechanisms, in trafficking in representational, or semantically evaluable, states (and hence to be not “*merely*” a mechanism). What especially matters is that the transitions between and among these states are normatively governed. The most familiar—though, as we will see, not the only—norm is that the transitions should be semantically sensible (e.g., truth-preserving).

Within this context, the foregoing characterisation of conceptual content yields something like the following (familiar) image. An agent with a mind is taken to be a causally-realised, normatively-governed creature inhabiting and interacting with the world around it. Semantically-evaluable information derived

from the world—that x is F , that y is G , etc.—is encoded in causally-efficacious representational vehicles that lead the creature to act in (normatively) appropriate ways towards that same world. In perception, for example, an agent's encountering a situation of x 's being F would lead it into a representational state the (conceptual) content of which would be that x is F . If the creature had inferential powers, and believed (for example) that all F s are G , it might then conclude that x is G . Or something like that.

As I say, it is a familiar picture. And if that were all there were to it, then perhaps all content would be conceptual—at least all content in experience or judgment.

But in spite of its familiarity, what thirty years of cognitive science have shown is that it doesn't work.

It doesn't work because the world doesn't do its part.

There aren't any objects out there.

2 A world without objects

Some of you may fear that, in denying that the world contains objects, I have taken leave of my senses. On that I ask you to withhold judgment until the paper is done. What I can assure you now, however, is that I have not abandoned realism. Perversely, in fact, it is exactly in order to preserve realism that the story I am telling must be told. It is not me, but the person who clings to objects, that is, as it were, "unrealistic" (though I admit that the nature of the real is going to come in for something of a beating, in order to see why that is true.)

Now in order to defend this strong a claim—not that there aren't any objects, which would be false, but that there aren't any autonomous objects independent of subjects—we need an industrial-strength theory of what the world is like. Where do we find such a thing? There are at least three candidates: (i) commonsense, and the deliveries of introspection; (ii) science; and (iii) our experience constructing contentful systems. I will consider each, in turn.

2a Commonsense

Start with introspection and commonsense. It certainly seems as if the world contains objects. *Just look! Lo: a table! a chair! a person*

Nonconceptual World

sitting on a chair! a person sitting on a chair at a table! Nothing, most people think, could be more fundamental to unreflective, lay experience.³ Perhaps that is right. But to say that commonsense judgments take the world to consist of objects exemplifying properties and standing in relations is just to say that at least the sorts of commonsense judgment that philosophers bring forward have articulated conceptual content. We know that.

Problem is, the fact that the natural attitude takes the world to be conceptual begs the question: of whether the world is autonomously conceptual, independent of and prior to our so taking it. So we have not made any progress.⁴

2b Science

So look at our second source: science. How could science have possibly succeeded, to the extent that it has, unless the world really and truly consisted of objects exemplifying properties, independent of our so taking it?

Thing is, I am not convinced there are any objects in science—at least not objects of the right kind. At the very least science doesn't provide a theory of objects—a theory of concrete, extended, reidentifiable, particulars, in Strawson's sense. It doesn't provide a theory of identity or individuation, in particular—and an object is not an object without identity or individuation conditions.

What science does give us is theories of *properties* or *types*: such as trees. But ask whether that clump of redwoods is one tree, or seven, or thirteen—and you'll find that science is of no help. Or suppose an amoeba splits. Did the old amoeba die, and two new ones emerge? Or is the old amoeba still with us—just spatially distributed? Or is one of the new amoebas the same as the old one, with the other new one having just been born? It is not just that biology doesn't provide any answers; *biology doesn't care*. This is be-

³Actually I am not so sure of this. Explain ...

⁴This is too simple. Many people—from poets to painters to phenomenologists—deny that it is intrinsic to the phenomenological character of experience that it "objectifies" the world, as one might put it. I strongly agree. But that only strengthens the conclusion we are aiming at: that it is not the world, on its own, that presents in terms of objects, properties, etc., but that that is an epistemic way of taking it, with very particular merits and demerits.

cause no scientific regularities, I'll wager, hold in virtue of object identity. Scientific regularities, as I've said, care only about properties.⁵

More seriously, consider physics. I don't know much about quantum mechanics or relativity, except to know that they are surpassingly strange. It seems wild to suppose that they might provide the theory of individuals that classical physics does not. And it is by no means evident that there are any individuals in classical physics. For think of the world that physics depicts: a four-dimensional manifold of continuous spatio-temporally extended density, charge, force, mass, energy, etc. The best way to understand the ontological commitments of classical physics is field-theoretically: a stupefyingly complex superimposition of interpenetrating waves, vortices and fields and quiescence and turbulence, vibrations from glacially slow to blazingly fast, forces continuously impinging, forces falling continuously away. Imagine falling overboard in a storm at sea, surrounded by nothing but crashing waves, stinging spray, and undulating currents, as far as the eye can see—and then subtract you. That is approximately what the world is like, according to physics—except a zillion times worse.

The investigative practice of physicists, of course, does deal in

⁵There are other objections. (1) Fields are nothing but space-time points, someone might argue; and a space-time point is a paradigmatic object. But I am not sure it is right that space-time points are paradigmatic objects. Space-time points are exactly not what ordinary objects are: extended. In fact it is not fully clear to me that we can even genuinely understand—can genuinely conceive of—space time points. We can understand representations of space-time points: temporally-durable points in a representation, such as a graph, in which time in the subject matter is represented by a spatial, not temporal, dimension in the representation). And we can arguably imagine space-points persisting in time. But no matter; perhaps I am alone in having trouble conceiving of a space-time point directly. The more serious comment is that, in virtue of not being extended, space-time points are intrinsically not subject to reidentification. There is no such thing as "encountering them again." Nor do they have paradigmatic properties of physical objects, such as being the common-cause of multiple effects, or the common effect of multiple causes. (2) What about chemistry? Surely chemistry deals in objects, such as molecules, or atoms? Or biology: what is a cell, if not an object? But is that really so? What are the individuation conditions on cells>

Nonconceptual World

concrete, discrete objects. Consider a high school physics problem: a mass of 3 kg traveling at 4 meters per second slides off a table 1 meter high; where does it land? There are at least two objects in this situation: a mass of 3 kg, and a table. Or think about calculating the gravitational attraction between two masses m_1 and m_2 . In such a calculation, those masses are likely to be taken as discrete objects. But of course they are not objects to which physics is ontologically committed. As objects, they are *idealizations*, are *approximations*. The true physical nature of the situation involves only a continuum of point-to-point forces and fields—a continuous manifold of physical disturbance. Yes, *qua* physicists, we traffic in discrete objects—but only for epistemic reasons, to make our calculations simpler—or even to make them possible.⁶ (Calculating the gravitational attraction between two extended objects, without this idealisation, would require solving a double triple integral.)

Some may agree that basic physics should be understood field-theoretically, but claim that ordinary individuals derive from (rise up on, emerge from, etc.) this field of physical forces by abstraction or idealization. That might be right. But—and in a way this is the point—‘abstraction’ and ‘idealization’ are terms from *epistemology*, not from physics. No natural science theorises abstraction as such (along with force, mass, and valence). No scientist writes “*abstraction(x)*” in their daily equations. On the contrary, like the discrete objects mentioned above, abstraction is part of the epistemic practice of scientists.

It follows that abstraction stands in need of explication. It especially needs explication if one is committed to anything like a naturalistic account of mind. In fact naturalising abstraction is one way to understand what we are doing here.

To make this clear, some terminology will help. By **physical** ontology I will refer to the world as theorised in physics; not the epistemically simplified version that permits calculation, but the ontological version to which the fundamental equations are committed. By **material** ontology, in contrast, I will refer to the furniture

⁶We may also treat higher-order objects as individuals. That is: it may be that physics cannot itself be formulated without objects (though I don't know for sure). But that's okay; my point is only that there are no first-order, discrete, concrete objects.

of everyday life: tables, chairs, dirty dishes, continents, détente. The claim I am exploring is that the objects of conceptual ontology are in the first instance *material*, not *physical*. That is: individuality, identity, being one as opposed to being none or being two, are not issues that any purely physical science can address.

2c Experience with constructing contentful systems

The third place to look for evidence as to the nature of the world is nonstandard, but I believe very significant: efforts in cognitive science and artificial intelligence to develop representational systems. That is: efforts to build semantic or representational systems that honour constitutive norms by representing the worlds they inhabit. This has proved to be an unbelievably sobering experience.

For several decades, Artificial Intelligence (AI) tried to build systems that took the worlds they inhabited to be (what I am calling) conceptual: to consist of well-behaved objects, properties, and relations, in the standard way. And AI failed. Virtually no one in the field any longer believes that the route to intelligence is via this kind of logical or conceptual representation.

Two developments are especially relevant. First, even in the early years, sobriety overcame students of perception. If you place a camera (or other sensor) on a robot, and show the signals to a first-time observer, they are stunned. The world does not “present” remotely as well-behaved as we imagine. Let me simply say, to any readers who have not worked with empirical data, that these results are unambiguous, repeatable—and unbelievably humbling. It is a bit tricky, of course: what you must not do is to display the incoming data in another medium that recruits the same sensor or perceptual apparatus that we would normally employ to deal with the world in that modality. For example, it does not work to display the signal coming in from a camera on, say, a TV monitor—because then our faculties of visual perceptual simply “parse” the image on the monitor in the same way that they would have parsed the scene that the TV camera is recording—giving us a false illusion that the world is well-behaved. But as long as you look at them in a different modality, or—better—construct algorithms to deal with them according to you think are the patterns holding them together, you are forced to conclude that *the messiness and partiality of the world outstrip anything untutored intuition*

Nonconceptual World

would ever have imagined. And this is not simply a result of poor instruments—noise in the signal, problems of transmission. The neatness is simply not there, in the world, for the instruments or signal to record. (Note that the problem is not just that the world recorded by instruments does not come pre-categorised. It is that the objects to be categorised cannot be *distinguished*, cannot be *segmented*.) The idealisation—and it is an idealisation—that the world is made of well-behaved, discrete objects has to be imposed on an unruly underlying reality.

The difficulties are not limited to perception—which leads to the second development bearing on these issues: the collapse of traditional or “symbolic” AI and cognitive science, quite apart from issues of perception. For its first few decades, the cognitive sciences operated on a model of mind in which the task of perception was to recognise essentially conceptual arrangements in the world, and to encode the results in explicit representations; the task of “mind” was to reason and solve problems with respect to these representations, and the task of “action” was to take representations of desired states and bring them about. That is the project that failed, and notoriously so. The reason normally proffered for this failure—a reason you still hear in the hallways of cognitive science centers—is that these systems failed because they were *representational*. At Indiana, for example, where I teach (you may not believe this, but I swear it is true) to claim that people represent the world is thought by many colleagues to be a recidivist, backwards view (about as popular as admitting that you listen to Mantovani and the 101 Strings). What is happening is that the traditional models are inexorably being replaced by a variety of non-traditional alternatives: connectionist systems, embedded agents (à la Brooks), systems that search high-dimensional state-spaces, dynamic systems, etc.

For several years, advocates of these new systems claimed that they were better because they did not represent (Rod Brooks, a champion of the new view, wrote a paper famously entitled “Intelligence without representation”). Indeed, antirepresentationalist tracts still regularly appear. On reflection, however, it is becoming clear that this characterisation—that the old school systems used representations, and the new school systems did not—is an inadequate way to characterise the sea-change. In spite of the press,

that is, the transformation has not been a shift from representational to non-representational systems. Rather, the situation is better described using the terminology laid out above. What was characteristic of the old school was that it used *representations with conceptual content*—that is, representations that represented the world of the agent to the agent in terms of what we are calling conceptual ontology: discrete well-behaved objects, properties, relations, states of affairs, and the like. What is characteristic of the new systems is not that they completely eschew representation. On the contrary: on the very general representational model adumbrated above, of normatively-governed, causal, contentful behaviour, the new systems are still fully representational. But these new representations represent the world in other-which is to say, in nonconceptual-terms.

2d Summary

In sum, neither science itself, nor our synthetic experience constructing systems to represent the world, supports the idea that the world *au fond*, consists of material objects. These two facts, in conjunction with the spectacular failures of early AI and cognitive science, suggest that we do well to be cautious in extrapolating from the content of our lay conceptual judgments to any view about the autonomous structure of the world.

3 Feature fields

But if not objects, then what?

My aim in this paper is to take a first, provisional step towards answering it. Not a step that gets us all the way. But a step that opens up a host of potent questions that any stronger answer will have to address. (Note: I am being conservative. I do not think that steps are discrete—and that we should take a lot of steps down this...plank. But bear with me; even this small step will have resounding consequences.)

In particular, with an eye on the fact that science was interested in types but not concrete individuals, I will follow Strawson, Cussins, and others, and characterise it in terms of what we might call a **feature-space**.

By a feature, imagine something like a property, but logically simpler, in that it does not require an object for its exemplification.

Nonconceptual World

That is: features are logically simpler—more basic, less committing—than more complex objects and properties. As a paradigmatic case, consider how we describe the weather. “It’s raining,” we say; or “It’s foggy.” But as everyone knows, the ‘it’ in such sentences does not refer; there is no thing, x , such that x is raining, or that x is foggy. Rather, as various writers have put it, “It’s raining” means something like: “Rain, here, now!” or simply “Raineth!” Features do not characterise particulars; they do not serve as sortals. “Feature-placing,” as Strawson called it—this is the crucial point—does not commit one to any discrete, concrete, extended thing that can be reidentified.⁷

So that is the image I want to examine: an extraordinarily rich, four-dimensional world (of time and space), that instantiates a bewildering array of features—colours, smells, textures, fogginess, whatever—without any concomitant commitment to individuals or particular identity.⁸ And remember that this is the structure of the world I am talking about; not the structure of the incident sensory array, the press of local causes at the agent’s periphery—not a manifold of sensation, received at the sensory organs of the creature, or a manifold of energy, pressing in on the creature’s skin. Rather, it is an arrangement, laid-out in space-time, of differentiation—not yet “grouped” into the rather large-scale, coarse, “synthesized” or “abstracted” individual objects of conception.

In particular, the suggestion is that the world presents to embodied intentional creatures as a (literally) unutterably-rich spatially and temporally continuous array of spatio-temporally instantiated features, sans identity. The question is how those creatures manage it—how, given finite physical resources, they orient towards it, get around in it, in ways that satisfy the governing semantical norms.

Part of the answer (as is clear from modern neuroscience and biology and cognitive science and AI) is that they have different strategies, for different purposes—even if in our own case we do a stunning job of putting them together in a seamless whole.

The alternative ways create something of a tension. At the local

⁷For a discussion of feature-placing see Cussins “Content, Embodiment, and Objectivity: the Theory of Cognitive Trails.”

⁸Say something about how this is only a first step ...

level, the richness of the real-time feature array is a boon: necessary in order to control the fine-grained detail of action and perception. If you want to place your finger *just there*, if you want to track that pattern of animal motion against a background of grasses waving in the wind, if you want to slip just so through this crack in the wall, then the rich detail provides critical resolution for fine-grained action. The superfluity of detail is not so good, however—exactly because it is *so* detailed—for long distance inference, reasoning, planning.

Some of the difficulties are epistemological (as we saw in the case of science): managing that complexity would swamp any finite computational mechanism. Some of the difficulties are ontological: long-distance correlations are often not sufficiently strong to warrant being framed in such detail. Fortunately, these two limitations conspire together: it is a good thing that we can refer to far-away objects without needing all the fine-grained detail, because by and large we do not know what that detail is. If we couldn't refer to Pompeii except by laying out the spatial configuration of every inhabitant, we could not refer to Pompeii at all, since we don't have any idea of exactly how many people lived there, let alone where there all were distributed. By the same token, if I couldn't remember you unless I was able to represent the exact position of your arms, I similarly could not remember you at all, since most of the time (unless I happen to be facing you) I do not know how your arms are arrayed. This is all banal, of course—no one doubts the ubiquity or utility of abstraction. My point is only that this abstraction is not something that happens *to* objects; rather, abstraction *enables* objects; it goes on *underneath* them. That is: objects are related to ur-reality by synthesis and abstraction. And synthesis and abstraction are *intentional*. To say of a patch or region of the world that it is an object already involves an intentional—and, to up the ante a bit, even a purpose-relative, which is to say, normatively governed—abstraction or synthesis away from the underlying sea of features.

This then is the picture I want to explore. Representing or otherwise dealing with the local, ever-so-rich feature array (the feature array in the world) is good for controlling action and responding to fine detail; it is bad for long-distance inference and generality. The non-local, long-distance abstractions underwrit-

Nonconceptual World

ing conceptual ontology are good for long-distance inference and generality, but not very good as a way of capturing the unutterable fine complexity of local circumstances. Achieving a realistic picture of the world requires mutual support of both. Intelligence involves being able to move back and forth between the two, flexibly, plastically-correctly-as appropriate.

I believe this is a sensible story, making sense of a variety of truisms: why you have to “be there,” in order to fully understand a situation; why there are such limits on “book knowledge,” and the like. I will also want to claim, ultimately, that it has various overarching theoretical benefits—giving us for example the wherewithal to begin to steer an appropriate path between realism and social construction. Nevertheless, the picture needs defense, if for no other reason than because it flies in the face of the “well-entrenched” (to put it mildly) intuition that there really are autonomous objects out there, independent of us.

I will start on that defense by considering, respectively, what nonconceptual content, and what conceptual content, come to, on such a view.

4 Featural content

... I am not sure whether this section has been adequately written. I have the following set of notes as to an outline of what may go into it, but then there is the text that follows—which may or may not instantiate the outline. All needs to be checked ...

- A. Summary
 - 1. Don't have time
 - 2. Investigate one specific kind: featural content
- B. Three major questions
 - 1. Featural content: really content?
 - a. Yes, because eminently revisable
 - b. I.e., can serve as reasons
 - c. Cf. McDowell's argument against it
 - d. He doesn't understand *semantically-warranted* (normatively governed) *causal transitions*
 - 2. Featural content: subpersonal?
 - a. No; perfectly accessible to consciousness

- b. Cf. Lowe slide
- 3. Featural content: really conceptual, after all?
 - a. No; not the same content
 - b. Doesn't figure in right generalisations
 - c. Would swamp finite mechanism
 - i. Cf. trillion lines of commercial software
 - ii. Illustrates all points (normative governance, etc.)
- C. Connection to action
 - 1. Not just a question of richness of details
 - 2. Also: indexical, differential, closer to physical coupling
 - 3. Therefore: shade into this kind, in order to control activity
 - 4. Tie to field-theory, differentials (deixis, etc.
 - 5. That in turn generates why more modality specific
 - a. Cf. Cussins: motorcycle, Evans "behind you", etc.
 - b. And non-generality: drive as fast as hit the tennis ball

Nonconceptual content is negatively defined, hence weak. That is not its only problem; I am not convinced there is just one kind of nonconceptual content—that it names a single, unified phenomenon. For the conceptual/nonconceptual distinction to be taken seriously, we need positive accounts: of rich, delineated, self-standing kinds of content—to compare and contrast with conceptual.

In this section I want to consider perhaps the simplest such possibility, generated by the picture of finite embodied agents operating in a feature space. I will call this species **featural content**—content that takes the world to be feature-instantiating. But even it is a bit of a grab bag. I will not here (though a proper account should) examine the space of features—*egocentric*, *allocentric*, more or less physical, etc.—nor say much about how different kinds of featural content can figure in different kinds of agent activity. A genuine theory awaits such detailed cartography. Here I must limit myself to some broad remarks.

If, as suggested above, the world in itself is more featural than conceptual, then the semantics of featural representations are in a sense simpler than the semantics of conceptual representations (which we will examine later), because of the more direct "fit." The basic idea is similar to Peacocke's scenario content: featural representations take the world to consist of a three-dimensional array

Nonconceptual World

of feature instantiations.

Because of their potential for richness, featural representation is paradigmatically indicated when an agent is in close contact with the world's details, presumably including perception and action. As I will explain in a moment, featural content is also the easiest kind of content for a physical mechanism to attain—it is of a sort that the laws of physics can underwrite rather directly. Intuitively, this makes sense: think of simple instruments, such as thermometers, light meters, etc. Setting aside for the moment the (critical) question of whether such devices really have content, or are only interpreted as such, it is nevertheless intuitive that, when we interpret them, a featural reading is most natural. Thus a thermometer indicates that “it is 90° here, now”—where, as in the case of weather statements, no ontological commitment is made to any object's “being 90°.”

For essentially engineering reasons, it is likely that fine-grained motor control will want access to the rich, detailed, pre-conceptualised (non-objectified) structure of its environment that featural representations are ideal at conveying, and that sensory mechanisms will yield these at the onset of perception. But no logical commitment is being made, in this account, as to whether any of perception, motor control, and conscious perceptual experience use the same representational vehicles, or even represent the world in compatible ways. It is certainly no part of the view being expounded that either perception or sensation must go with featural representation; conception with conceptual.⁹ By the same token, there is no logical requirement that featural (or other sorts of non-conceptual) states can only be engendered by direct encounter with the world. It is a benefit of an ontological approach that, instead of building such claims into the theory itself, it provides the wherewithal for giving such claims genuine empirical content, if they turn out to be true.

Rather than go architectural issues, however, what I want to do here is to consider three questions that will naturally be asked: (i) whether such featural states warrant the label content at all; (ii) whether they are not intrinsically sub-personal; and (iii) whether,

⁹Unless of course one were to stipulate the difference between perception and cognition in such terms.

as described, they are not really conceptual, after all.

4a Is featural content really content?

What warrants the claim that nonconceptual states—in particular, any states that arise in a creature with featural content—can legitimately be said to have content, at all? That is: how do I know that my describing states as nonconceptual isn't merely word-play?

There are two facets of essentially a single answer. Overall, nonconceptual states are governed by the same normative/semantical considerations that apply to conceptual content. As such, they can serve as rational (or at least normatively appropriate) reasons for an agent's action. There is no implication, from that critical fact, that they must thereby be able to be given "conceptual voice." It is not an *a priori* truth, after all, that if a content-bearing state is a reason for an agent's doing something (bending its arm in a such a way, feeling that danger impinges, whatever) that, if asked, the agent must be able (even potentially) to articulate its reason. In fact commonsense, lay experience, artistic sensibilities, psychoanalysis, cultural anthropology, and just about every other form of human study suggests the opposite. That is not to deny that it is philosophically common to assume that reasons must be articulable; the point is only that, in the current context, to presume that would be empty. All that that assumption comes to is a claim that the only genuine form of content is conceptual—exactly what is being denied.

Moreover, featural content is eminently revisable. There is nothing about nonconceptual content that means that it just is what it is, independent of semantic constraints from the world it represents. Nor—it must be emphasised—is nonconceptual content in any sense "given." By analogy, consider photographs. The content of photographs is surely derivative, not original—but qua derivative content, it is also surely nonconceptual. The photograph itself is a two-dimensional spectral density array, representing something like a three-dimensional reflectance array that laid out in front of the camera. Suppose, while looking over some photographs, you come upon one which does not look right, and say "Oh dear; this one is a double exposure!" Why do you say that? Because the world represented by the photograph is not possible,

Nonconceptual World

or anyway so unlikely as to warrant extreme doubt. By the same token, imagine a depiction (painting or photograph) of dense shrubbery surrounding a path across a rocky slope, with wisps of fog swirling up from a valley below. Now that description is necessarily in language (this is a philosophy paper), and will therefore have conceptual content;¹⁰ but what I am asking you to imagine is one of the infinite variety of visual scenes from which that conceptual abstraction is possible. Now ask what that scene would look like if one were to take a few steps forward. There are constraints on what would be a legitimate answer. Some depictions are incompatible, some depictions are wrong—depictions that, if they were delivered to you in consciousness by your perceptual system, would make you rub your eyes and look again (or perhaps rub your eyes and try to wake up).

Conceptual representation, in sum, has no patent on revisability. On the contrary, featural and other nonconceptual states are just as amenable to semantically-warranted transition—and hence are just as capable of semantically-unwarranted transition. So they can serve as reasons.

4b Is featural content subpersonal?

Some may argue that if there are states with featural content, in the way I am proposing, then they must be “subpersonal.” But I believe that is manifestly false. For nonconceptual states as I have described them are available to consciousness. Indeed, it is (perfectly realist) consciousness of the non- or pre-conceptualised world, I believe, that has led legions of philosophers to believe in a spate of such ontologically unfortunate entities as sense data, visual fields, “experiences,” opaque mental states that can be reflected upon independent of their content, etc.

Another example: suppose, visiting in California, you say “The fog has come in again.” Suppose I, in turn, say to you: “You have a PhD. In saying ‘the fog has come in,’ do you mean the same instance of fog? or a new instance of fog of the same type?” There is no appropriate reply. There is no reason to suppose that, in making your original statement, you are epistemically committed, or have made any ontological commitments to, an object that is “the

¹⁰I don't quite believe this, of course ...

fog,” Rather, you have merely judged that the feature “fog” is re-instantiating itself again, around here, around now.

But that is a linguistic example—which is distracting (since we are largely associating conceptual content with articulated language). A better example is given in figure 1. This image (by Adam Lowe) can be “parsed.” It depicts the painter’s studio, with a door on the right, a waste can on the floor to the left of the door, a cheap phone attached to the door frame, also on the left, and a person’s body moving towards the door, from the right.

Pictures such as this—and even more so, pictures that look roughly like this, but that cannot be so readily parsed, or that cannot be conceptually parsed at all—are paradigmatically labelled abstract. I think that labeling is exactly backwards. It is so-called “representational” pictures that are abstract, by my lights. They discard the rough and tumble of the world—“over-neaten” it, pull out and present to conceptual judgment, in imagistic form, what our conceptualising faculties do to the world, in normal perception. It is Lowe’s painting, in contrast, that is concrete-concrete in the sense that it depicts the concrete world as it is.

Admittedly, this image may be slightly exaggerated—but only a tiny bit (remember those robot cameras). That is because what the painter is trying to do, with this image, is to bring to our conscious awareness an inkling of how the world presents to our unconscious or artistic sensibilities.¹¹ Forget the image, therefore, and reflect on the world. The image is merely intended instrumentally: to bring us to reflect, consciously—perhaps even with conceptual supervision and commentary—on the pre-conceptualised reality that is what we regularly look out upon, whenever we open our eyes.

4c Isn’t featural content conceptual after all?

One final objection must be dealt with, which is likely to be raised to the picture developed so far: that there is nothing nonconcept-

¹¹If one were mathematical, one might say that if the conceptualising faculties of judgment that conceptualise the world transform it according to some function f , then Lowe has painted this image to be something like f^{-1} of what the world is really like, so that what arrives into conscious conceptual judgment is, as closely as possible, the nature of the pre-conceptualised world.

Nonconceptual World

tual about the sorts of featural representations being discussed. “Surely,” such a person might say, “a nonconceptual featural representation of the world can be expressed, conceptually. Simply take each point in space and time as an object, and predicate, of that space-time point, a property that corresponds (isomorphically) to the there-placed feature. For each point in the depicted region, one can simply predicate the appropriate colour value of that point.”

There are number of things to say, by way of reply. First, there are problems of continuity. Conceptual representation seems to involve a certain degree of digitization, and featural representation, at least as I have presented it, can be, and perhaps most often is, continuous. That suggests that one would need to digitise the image first. And imperceptible digitization is no simple affair—as decades of work on graphic displays and audio in the computer and entertainment industries attests. In particular, there is perception’s notorious non-transitivity of indiscernibles: the fact that there can be three colour patches, x , y , and z , such that a subject cannot tell the difference between x and y , or between y and z , but *can* distinguish x and z .¹²

¹²For example, someone might suggest that one could digitize an image, at some level of resolution finer than the optical resolution of the eye, and store colour values for each pixel, making just as many distinctions as are required to meet “JND” properties of the visual recognition system. This is not such a simple thing to do. The point of JNDs is that there can be three colour patches, x , y , and z , such that a subject cannot tell the difference between x and y , or between y and z , but *can* distinguish x and z . If one were to record colour at the level of resolution of the eye—i.e., one bit per JND—then it is clear that x and z should receive different codes. But what about y ? It cannot be given the code for x , or the code for z , and reproduce the same phenomenology.

This means that the digitization of recording must be finer than the resolution of the sensory system. How much finer is a matter of intense debate. The case of digitised audio is instructive. When compact discs were first produced, it was widely believed that the then-standard digitisation standard—16 bits of information 44,100 times a second—would suffice, since the upper cutoff Nyquist frequency was above the range of human hearing, and the resolution of a 16-bit amplitude (1 part in approximately 32,000) was well below the audible JND of volume discrimination. It is now widely believed that those standards were insufficient: CDs are audibly inferior to the resolving power of the human ear. The professional audio industry is now moving to a more informationally-dense standard of

But in a way that misses the point. For while in some abstract sense it may be possible to construct a conceptualisation that has the same information content as the original image, it does not follow that it would have the same content. It will not figure in the same semantically-warranted generalisations. For imagine: take Lowe's image, given above, to be an image of his studio, taken from a perspective a few feet away from the door. How would the image change, if one were to step a few feet to the left? Nothing in the brute-force conceptual recording of the image gives one any help with that.

Moreover, there is no reason to suppose—and every reason not to—that we represent images in such a way as to have ready access, meshed with our more general conceptual powers, to such a detailed recording of our visual experiences. The data and processing load this would demand would be overwhelming. And if cognitive science has taught us anything about the architecture of the mind, it is that considerations of computational complexity, even in a mechanism as dauntingly impressive as the brain, are of the utmost importance. Moreover, as mentioned above, a conceptualisation that relies on space-time points (i.e., on non-extended, non-durable objects) is conceptualisation in letter only; it does not deal with any of the issues—of synthesis, abstraction, reidentification, life-time variability, etc.—that are constitutive of objects in real life. Space-time points are not really objects.

In passing, it is perhaps worth pointing out, in this vein—for people tempted by such conceptualised recording—that essentially all modern computer software design can be viewed as implementing semantically-warranted causal transitions on nonconceptual representations.¹³ I estimate that something like a trillion lines of computer programs have been written, to date; virtually none of it uses conceptualised representation. Because of the egocentricity, purpose-specificity, and contingency of the tasks that computers are up to, it is vastly more efficient and reliable—indeed, it may be

24 bits of information sampled 96,000 or even 192,000 times a second. But the data implications are enormous: uncompressed, that implies something approximately 1 gigabyte for 40 minutes of music, which even today is substantial storage.

¹³Except for the vanishingly small number of AI expert systems.

Nonconceptual World

the only possible way to get anything done—to use purpose-specific representations. So no one should think that purpose-specific algorithms, such as “abstracting” algorithms, in order to get to the long-range conceptualised skeletons of the world’s nonconceptual detail, are in any sense precluded.

4d Connection with causation

To conclude, set objections aside, and consider one final argument in favour of featural representation.

Two things that I have not emphasised here, but that need to be dealt, include:

1. Nonconceptual representational vehicles, including featural vehicles: i.e., the structure of the representation that bears nonconceptual content, viewed as a causally-efficacious entity (the problem of mental content, after all, stands as need of solution in the nonconceptual case as in the case of conceptual content); and
2. The semantic (interpretation) relation between the vehicle and the (nonconceptual) world thereby represented.

But some suggestions along these lines have been implicit in various examples I have used, such as photographs, TV cameras, and recording instruments. This ties into the statement made earlier, about the field-theoretic nature of physics, and the suggestion that science is interested in types (or features).

According to the broadly representational theory of mind we are working under, an agent works, causally, but subject to governing norms in such a way that it is not a “mere” mechanism. That means that an embodied, embedded agent—as all agents must be—will be causally plugged into their environments (though not, again, merely causally plugged in). As I explore in detail elsewhere,¹⁴ it turns out, given the way causation works—i.e., as a consequence of the nature of physical law—that the form of representation that is easiest to have is nonconceptual: one whose features correspond, moment-to-moment, with the features of that with which it is causally coupled. That is why instruments—thermometers, microphones, photographs, cameras, etc.—

¹⁴«Reference O3; also “Who’s on Third?”»

tend to be such good examples of featural (nonconceptual) signifiers.

It also turns out, for similar reasons, that the content of the simplest form of representational mechanism will be egocentric or indexical—“deictic,” as I have put it. This fact has profound implications for the nature of first-person reference and consciousness awareness; it also establishes the nature of the task that an agent faces in order to have objective content (be it conceptual or nonconceptual). For now, though, the point is that it is a consequence of the nature of underlying physical laws that the sorts of structure that must guide action must be vehicles with egocentric, nonconceptual content that structurally, qua vehicles, as they get closer and closer to the world, grow increasingly isomorphic or iconic to the content they carry.

5 Conceptual content

*... Again, unsure about the relation between this outline and the following text
... whether it is complete, etc. ...*

- A. Conceptual content
 - 1. Turn then to conceptual content
 - 2. In a way, even more interesting, because of lack of fit
 - 3. Once again, deal with three major topics
- B. Skip the first two
 - 1. Structural correspondence
 - a. If underlying reality is fields of features ...
 - b. What is object reference (objectification) like?
 - c. Two cross-cutting algebraic kinds of correspondence
 - i. Property / object (as in Generality condition)
 - ii. Singular object reference
 - Has to do with name \approx type, object \approx instance
 - Not point to point
 - iii. Field theory of object reference
 - d. Complicated by stabilisation, deixis, first-, second-, and third-person objectivity, etc.
 - 2. Relation to nonconceptual
 - a. Basic: non-conceptual anchors conceptual

Nonconceptual World

- b. Not logically identified with perceptual / cognitive (())
 - i. Allows cross-fertilisation (interpenetration)
- c. Make sense of many truisms
 - i. Limits of book knowledge
 - ii. “Had to be there” (walk in another one’s shoes, etc.)
- C. Norms
 1. Third one has to do with norms.
 2. Want to wrap up with this.

Turn then to conceptual content: what it is, how it could arise in a featural world, how it is tied to (or grounded on) nonconceptual content. In a sense, this is the interesting case, because of the lack of obvious fit. How, if the world is not (aboriginally) conceptual, can claims with conceptual content be true?

Conceptual content takes the world to consist of objects, exemplifying properties, standing in relations, grouped in sets. Rather than being metaphysically basic, conceptual ontology is “constructed” by intentional creatures, using processes of abstraction, out of an explanatorily and ontologically prior world (a world we are for now taking to consist of a vast array of instantiated features). The question is how those abstraction processes go.

I want to consider three issues, as a way of getting at the answer. The first has to do with the form of correspondence that conceptual representation bears to the world. The second concerns the relation between conceptual and nonconceptual content. The third, which is also the most consequential, has to do with the norms on which the abstraction processes are based.

5a Structural correspondence

We have identified two features of conceptual content: abstraction and recombination. A modal claim was made, about the (potentially) recombinant structure of conceptual ontology: that if a is F , and b G , then a might have been G , and b F . This potential for recombination must be reflected in the conceptual vehicles. It is a constitutive condition on conceptual abstraction, that is, that it eventuate in a kind of Evansian generality—guaranteed by appropriate patterns of rational inference.

The potentially recombinant world, that is, is reflected in potentially recombinant representation. It is not, of course, ulti-

mately reflected isomorphically: negation, disjunction, numerals, and quantifiers all famously break any one-to-one correspondence between language and world. Historically, however, it is possible that it started out isomorphically:¹⁵ “This is blue,” “Pat is eating dinner.” But isomorphism is broken the minute ‘two’ is introduced: ‘two’ represents duality with unity.

That claim, about isomorphic or partially-isomorphic mappings between representation and represented, is couched in terms of conceptual registrations of both realms. Things get much more interesting when we look at the correspondence (interpretation) of conceptual representations, but understands the represented world featurally. For what emerges is that the recombination of parts is only one form of vaguely algebraic coupling; another one, underneath the objects, is explanatorily more basic. One of the characteristics of featural representations, mentioned above, is that they typically (at least in the simplest cases) involves a point-to-point correspondence between vehicle and content: at time t , the sunflower points at the direction of the incident sunlight at time t ; at time $t+1$, it points at the direction of the light at time $t+1$.

But think about reference to an object—say, with a proper name. Suppose the name ‘Pat’ refers to a person, Pat. We think of this as a one-to-one correspondence: one name, one person. But of course that is a distracting way to put it, since the name is a *type*, the person, an *instance*. There are instances of the name-uses or utterances, that typically occur at a specific moment in time. As our experience with indexicals has taught us, it is these temporally-specific uses that refer. Since objects do not exist in the world independent of being objectified, on this story, there is also a sense in which there are “instances” of objects: namely, those temporally-specific “manifestations” or “time-slices” of objects that also “occur,” at different moments in time. Crucially, however—and this is the important point—*individual temporal utterances do not refer to individual time-slices of their referents*. When I thought of you last night, I did not just think of you-last-night, or (even less) of the then-occurrent instantaneous time-slice of you. Rather, I

¹⁵Something of the sort is suggested by Terrence Deacon—though since his semiotics is not very developed, it is hard to know whether he thinks that this was only an evolutionarily transitional stage.

Nonconceptual World

thought of how you were doing at school, whether you were prepared for today's exam, etc. Today, when I thought of you, I again did not think of today's time-slice, but again of you as a temporally extended entity. In other words—as depicted in figure 2—each instance of a proper name type refers to the full extended space-time worm (or whatever region of the infinite flux) you constitute. So the featural (or field-theoretic) structure of even a simple name-object relation involves various forms of *cross-cutting coupling*. All of this is required—is an achievement of subjects—in order to refer to an object as an object (and referring to objects is surely one of the most basic capabilities of conceptual representation). The complexity of this cross-correlational mapping, as compared with the simple form of point-to-point correspondence characteristic of featural representation, underscores the significance of the accomplishment that is intrinsic to conceptual abstraction.

5b Relation to nonconceptual content

Even on a classical account, few would deny that thoughts are anchored in perception and action—that human reference is grounded in our engagement with the world. But on the story being told here, a much stronger moral emerges, having to do with the relation between conceptual and nonconceptual representation.

We can get at this moral by noting two facts about conceptual content.

First, conceptual content involves loss. When we take the world to consist of objects exemplifying properties and standing in relations—when, that is, we “objectify” the world—we discard staggering amounts of information (the vast majority we are presented with, in fact). Remember those robot cameras; once again, computational experience is a sober reminder of the prowess—this time, the “forgetting” prowess—of the brain. And as I have said, it is fortunate that we shed this much detail. Given finite computational resources, it is only with the radically pared-down result that we have even a prayer of doing passable inference. (That's one reason conceptual representation is valuable; if one tried to compute with

full featural maps, the computational load would be intractable.¹⁶⁾

Second, conceptual representations are *disconnected* from their referents. That representation be able to be disconnected from its reference is well-recognised; it is that ability that allows us to hypothesize, to refer to things beyond the reach of our senses, to have a sense that there is a world out there, beyond the reach of our senses. Once one recognises that objects are an abstraction over the world, not part of the world's aboriginal structure, though, one is forced to realise that essentially all representation is disconnected- if for no other reason than that objects by and large are perduring- exist through time-whereas all that physics allows us to couple with (because of its locality) is the present moment.

In sum, conceptual content sees the world "through a glass, darkly." It "lets go" of the world, discarding vast amounts of detail, so as to support generalisation and long-distance inference. There is a worry, however, given the depth and human centeredness of this abstraction (more on that in a moment), that the conceptual content will take leave of the messy details of the world altogether, and float entirely free.

But of course that is exactly one of the roles of featural content: that it anchors the "abstracted" conceptualised objects built on top of it. Nonconceptual content is the "glue" that binds abstracted objects and properties to the pre-objectified world. It is what keeps the fact that even concrete objects are abstracted from implying that they take leave of reality. This is why I said above that nonconceptual content is necessary in order to retain what is right about realism.

What this consideration shows, however, is that abstraction is just half of the story. For as normally conceived, the term 'abstraction' refers to the processes of "letting go" of the world: to the discarding of the mass of featural (and other nonconceptual) information, so as to achieve a finite, compact, gloss on what is the case.

¹⁶Why the world should be such that conceptual abstraction works is a non-trivial metaphysical question, which I will not address here except to note that what it seems to work best for are artifacts, which we build- perhaps with malice aforethought. Just as language has evolved subject to the constraint that the human brain can speak it, so too artifacts may have developed subject to the constraint that the human brain can understand them.

Nonconceptual World

But that describes perception; what about action? What happens when, on the basis of a conceptualisation of the world, we reach out to do something?

What happens, I believe (in part for the reasons cited above, about the closer fit between featural representations and causally coupled mechanisms) is that *our nonconceptual faculties enter into the equation so as to fill back in the requisite detail*. Thus imagine deciding to reach for a cup. You have that thought: “I will pick up this cup.” And it may even be that in the conceptual thought is some indication of the size of the cup: perhaps it is a small latté, or large mug. But then, as your hand approaches the cup, your fingers adjust, through representational mechanisms (prior to contact), so as to be prepared much more exactly than they could have been, in virtue of conceptual content alone. The same for serving tennis, for leaning into a corner on a motorcycle, and so forth.

In previous work I used the term ‘reconciliation’ for this process that is the opposite of abstraction—this “filling back in” of the world’s detail that is lost when one conceptualises.

... use: ‘concretisation’ ...

In sum, conceptual abilities are required to conceive of an object as an object, to conceive of an object as conceptualised. Nonconceptual abilities are required in order to understand that which is conceptualised as an object. Only if you understand that an object is a conceptualisation of reality do you really understand what an object is. So nonconceptual content is thus not “optional”; it is a necessary ingredient to objectivity.

This conclusion contains strong lessons for AI. It implies that purely conceptual creatures have no chance of achieving objectivity, because they exactly lack the critical (nonconceptual) glue that binds their conceptual conceptions to the gritty stuff and substance of the world. No wonder “book learning” is limited-and eerily detached.¹⁷

¹⁷Note an irony to the story I’ve been telling. I started out saying that I wanted to “ontologise” conceptual and nonconceptual content-to push it out from heads into the world. Now, however, it seems that I am allowing the nature of what is “in the world” (particularly in the case of conceptual

5c Norms

Third and finally, consider norms. In a way, the point is simple: Those constitutive processes of abstraction, qua epistemic practices of rational, norm-governed agents, happen for a reason. It is the ontological character of that statement that makes it strong. It is not just that there is a reason people represent objects, in other words. That much is obvious—and anyway guaranteed by the normative character of the semantical story within which we are working. The point is stronger. Which abstractions a creature makes—and as a result, what objects there are in the world, for that creature—arises out of the constitutively norm-governed life that that creature leads.

Objects themselves, that is, not just their representations, have their existence in worlds governed by significance, interpretation, and importance. To put it in a two short words: **objects matter**.

This claim has a happy terminological consequence. In §1, I distinguished *physical ontology* (the strange world described in modern physics) from *material ontology* (the everyday world of human experience). Now it is common to suppose that the word ‘material,’ in English, has two distinct senses: (i) a more common one, meaning something like physical or bodily, as in “living in the material world,” materiality, materialism, and so forth; and (ii) a less common, vaguely legalistic one, meaning something like important, as in a “material argument,” or “material consideration.”

In calling everyday ontology material, it may have seemed as if I was recruiting the first, roughly physicalist sense. But my intent was more devious. For what I am suggesting—which we can now see—is that there are not really two senses, after all. If, as I claim, the processes of synthetic abstraction constitutive of conceptual content are anchored in the norms governing the lives of conceptualising creatures, then material ontology (as I have defined it) is in part normatively derived. *To be an object is to be important* (to

content) to slip back partway into the agents that inhabit it. There is a grain of truth in this blurring of the subject/object boundary. Still, this is absolutely not a story that devolves into pure idealism or vacuous relativism; that would only be true if the reality (the reality that for the moment we are characterising as featural) had no grip on the thereby-conceptualised objects’ nature. But that is no implication of what is being claimed.

Nonconceptual World

someone), in the world. It is because of this fact—that objects *do* matter, not just that they *are* matter—that I called them material.

Part of what we are doing, that is, in this project of naturalising abstraction, is healing the 300-year gap between matter and mat-
tering.

To see what this comes to, let's bring it to bear on an issue that has been lurking in the background for some time: what it is (on this picture) for a conceptual judgment to be *true*.

I trust it is clear why this is a non-trivial issue, on the picture I am painting. I have claimed, after all, that objects, *qua* objects, are not wholly independent of people (of us). I have depicted featural content as closer to the “mind-independent” structure of the world than is conceptual content (though, to repeat, it is only *closer* to reality; I am not saying that it is reality—remember, we are just taking a first step down that plank). So it looks as if featural content has a better claim than does conceptual content on being true, perhaps even on being objective. But that cannot be. It would be perverse—even nihilistic—to deny to conceptual judgments the possibility of truth. Rather, the question we must ask (in a spirit of reclamation) is this: *what can or does truth mean, for conceptual judgments, on an abstracting, human-implicating picture?*

I want to get at the answer by going back to where we started: with the representational theory of mind. In broad brush strokes, I characterised that view as committed to a picture of mental life as involving semantically-warranted, normatively-governed, causal processes. At that very general level, I remain sympathetic to the view. It is (among other things) in the details of how the norms are treated that the view I am proposing radically parts company with standard accounts.

To see why, consider the classical (logicist) picture. It works as follows: one starts by distinguishing (static) states from (dynamic) processes defined over states. Given this distinction, the norms then attach in stages. In the first stage, semantic evaluation is defined for the states—in a way that is assumed to be explanatorily prior to, and independent of, their use in inference or reasoning. Then, with truth and reference in place, a second set of norms is defined for the processes, in terms of that presumptively prior semantic valuation. Once we realise that semantic evaluation

is a species of normative governance (truth being better than falsehood, information better than misinformation, etc.), this classic view can be summarised as follows. What I will call the **dynamical norms** (norms on activity or use, such as on inference) are assumed to be explanatorily derivative on the **statical norms** (norms on states—such as reference, truth, etc.). This form of asymmetrical dependence underwrites all standard accounts of soundness, completeness, truth-preservation, etc.

Experience with real-world systems, however, shows us that this strategy does not work. It turns out to be impossible to assign semantic evaluation prior to and independent of activity. Rather, over the last few decades, in a perhaps unwitting endorsement of a vaguely Wittgensteinian doctrine of “meaning is use,” computer and cognitive scientists have all come to lean in the opposite direction. They have shifted to the opposite form of explanatory dependence, with semantic evaluation, content, interpretation, etc., taken to derive from large-scale dynamic activity. To put it in terms of the terminology just introduced, it may not be recognized as such, but it is nevertheless virtually universally assumed that *statical norms derive from dynamical norms*.

This shift is *unimaginably consequential* (in spite of the fact that no one seems to be noticing it).¹⁸ What makes it so important is that it requires, on pain of circularity, that something else (other than “preservation of the static norms”) ground, or serve as the origin of, the dynamic norms.

What are some plausible dynamic norms? Several alternatives have been pressed into service: meeting a specification, maximizing an externally-supplied value, etc. These days, however, especially in cognitive science, philosophy of mind, and evolutionary epistemology (to say nothing of Artificial Life and the theory of complex adaptive systems), the dynamic norm most in favour is that of *adaptability* or *evolutionary survival*. (I have in mind for example Ruth Milikan, teleofunctional semantics, notions of proper function, etc.)

Whether evolution will prove strong enough to anchor the

¹⁸There seems to be a kind of “Road-Runner” effect: everyone has rushed off the cliff, but no one has yet looked down.

Nonconceptual World

range of dynamical norms needed in order to explain the human condition is of course a matter of intense debate. I am not going to broach that topic here, except to say that I do not believe it. I do not think evolution is remotely strong enough; I am not even sure it has the right categorical form. For note how much is being asked. The full dynamic norms on the condition of a person will among other things include *ethics* and *significance*—how to live, what to do, what constitutes a meaningful life.

However those questions work out—whatever dynamical norms govern human (or humane) lives—what is relevant here, and this I do want to insist on, in the present context, is that they do not just apply to reference and truth; *they also apply to ontology*. If the story I am telling about conceptual content is right, that is, dynamical norms underwrite not only the semantical content of an agent's representational and intentional states, but also the constitutive normative standards in terms of which the agent forms its conceptual abstractions.

And with that we are finally ready to answer the question of what it is for a conceptual judgment to be true. A conceptual judgment is true when the thereby-abstracted situation satisfies the dynamic norms governing the lives of the creatures who perform it (i.e., the creatures who objectify its constitutive objects, delineate its constitutive properties, and so forth). Roughly, that is, a conceptual judgment is true just in case conceptualising the world in that way—including not only the patch of the world thereby conceptualised, but also the act of so conceptualising it—is a “success maker” for the objectifier with respect to the overarching norms that govern that objectifier's projects.

It is a theorem of this view, that is—a consequence of the recognition that objects matter—that (conceptual) truth is to a degree pragmatic. To put it in a slogan, conceptual truth (that is, truth applied to states or judgments with conceptual content) ultimately depends on living truly.

Some may still object that even if the choice of how to abstract or conceptualise is purpose-relative—perhaps norm or project relative—the “space of possible abstractions over real-world features” is nevertheless already out-there. And so, they might say, the story is realist after all. Conceptual ontology is not so much *constructed*,

on this story, they might say, as *selected*. If one chooses to call this realism, that is okay with me. But I believe it is a misleading way to view things. Consider something as simple as a 16×16 array (the number of pixels in used for the cursor on your computer). The number of ways to colour them (i.e., the number of possible distinct cursors) is hundreds of thousands times greater than the number of electrons in the universe. The intrinsic choices for grouping are so vast, that is, that most of the nature of the choice derives from the constraints one obeys in making it—constraints (to return to the case at hand) that derive from the creature's normatively-governed life. The nature of the abstraction, that is—the abstraction that the object, *qua* object, must normatively honour—derives in part from the intentional practices of the representing agent. And that agent's practices are grounded, ultimately, on anchoring in the sustaining field of features—or more generally, in the nonconceptual world.

Moreover—to shift up one level—there is no guarantee that the norms that ground this conceptual abstraction will themselves be conceptualisable without loss. Moral realism, to take one striking example, does not imply moral effability.

6 Conclusion

More can be said—but time has run out. What have we learned?

The metaphysical morals are the strongest. It is not just semantics that needs naturalising; ontology needs naturalising too. That is the first lesson. At least material ontology needs to be naturalised: the familiar everyday world of objects, properties, and relations. The processes of abstraction that underwrite the grouping and individuation of nature into material individuals, essential to our understanding of cognition, are no less mysterious, no more secured by a mechanistic or causal scientific world view, no more automatically integrable with results in contemporary science, than any of the other challenging features of intentionality. To assume that the world of objects, properties, etc., exists independent of us, in fact, as naive realism and any commitment to the “natural ontological attitude” would suggest—to think that the world is autonomously conceptual—is an especially pernicious way of succumbing to the Myth of the Given.

In sum, this is a picture of *metaphysical monism*, but *ontological*

Nonconceptual World

pluralism. And to my sensibilities, that seems just right. It allows us to do justice to the humility that underwrites constructivism: the recognition that we are here, and have an impact on the worlds we inhabit. And it allows us to do justice to the humility that underwrites realism: the recognition that we are not all that is here, that ultimately we have to defer to that world of which we are a part, on which we have an impact, but that we do not control—and certainly cannot ultimately grasp.

What we do is to make our way as best we can—neither impotent nor omnipotent, neither ignorant nor omniscient. We live, that is, in something of a middle ground—in a continuous creative tension between the incredible richness and unabstractable detail of local coupling, on the one hand, and the long-distance utility of language, inference, and abstraction, on the other. When we engage directly with the world, we want to do the opposite of “abstract”: we want to *concretize*, to reconcile our ideas with reality, to let more of the world’s ultimately ineffable details fill our representations, in order to be appropriately responsive-in action, in perception, in local, contingent reasoning—to the world’s fine-grained, particular, structure. When we want to travel long distances—in order to conceive of the world as a whole, in order to create complex institutions, in order to do science—it pays to let go of that overwhelming profusion of local detail, and employ sparer, more efficient methods—methods purpose-designed for inferential travel. If we are clever (and surely we are clever) we can—in fact must—do both, in such a way that each props up the other, thereby allowing the nonconceptual representations to approach a kind of objectivity, and the conceptual representations to, in their own way, be true.

Perhaps the best way to summarise this is by an analogy. I sometimes think of objects, properties, and relations (i.e., conceptual, material ontology) as the long-distance trucks and interstate highway systems of intentional, normative life. They are undeniably essential to the overall integration of life’s practices—critical, given finite resources, for us to integrate the vast and open-ended terrain of experience into a single, cohesive, objective world. But the cost of packaging up objects for portability and long-distance travel is that they are thereby insulated from the extraor-

NC · 38

Indiscrete Affairs

dinarily fine-grained richness of particular, indigenous life-insulated from the ineffable richness of the very lives they sustain.

