

interested in any of the concepts mentioned in its title will find something worthwhile in this collection.

MARIA ALVAREZ
King's College London, UK
maria.alvarez@kcl.ac.uk

The World in the Head

By ROBERT CUMMINS

Oxford University Press, 2010. x + 328 pp. £37.50

This volume collects Robert Cummins's papers on mental representation. With three exceptions, the guiding principle of these papers is that it is a mistake to read off the structure and content of mental representation from the syntactic and semantic features of language, respectively. The exceptions are the only previously unpublished paper in the collection, 'What is it like to be a computer?', which serves to explain Cummins's focus on cognition over consciousness (the 'hard problem' isn't yet sufficiently well posed); and two papers arguing against innate cognitive modules ('Biological preparedness', with D.D. Cummins, and 'Cognitive evolutionary psychology', by Cummins and Poirier). Cummins begins with a two-pronged attack on Jerry Fodor's 'Language of Thought' (LOT) hypothesis (Fodor 1987). According to LOT, thoughts are semantically composed of atomic symbols whose meanings are determined (roughly) by what would cause them to be tokened. Cognition is the process of manipulating such thoughts according to their syntactic structure. Cummins first argues ('LOT of the causal theory') that the causal and compositional elements of LOT are in tension. Primitive symbols derive their meanings from their roles in detection – I have the concept *cat* only if I can detect cats. But detecting cats requires a theory of cats – they are furry, have claws, etc. According to LOT, however, any theory mediating the detection of cats must have the concept *cat* as a semantic constituent. Having the concept *X* requires we can detect *X*s, detection of *X*s requires a theory of *X*s, but – given LOT – a theory of *X*s requires *X*, so LOT with a causal semantics for the atomic symbols is circular.

The next two papers – 'Systematicity' and 'Systematicity and the cognition of structured domains' (with Blackmon, Byrd, Poirier, Roth and Schwarz) – attack Fodor's central argument for LOT: that only semantically structured thoughts plus structure-sensitive cognitive processes could account for the systematicity of cognition. In LOT, the fact that anyone who can understand the sentence 'Mary loves John' can also understand 'John loves Mary' is explained as follows: (i) understanding a sentence is thinking a thought with the content it expresses; (ii) the thought that John loves Mary has the same atomic constituents as the thought that Mary loves John; (iii) there are structure-sensitive cognitive processes capable of rearranging the constituents, according to syntactic rules, to form new thoughts. Cummins argues that non-classical tensor product encodings of symbol systems, such as that

given in Smolensky (1990), can parse sentences systematically as well. Fodor's standard objection is that it's nomic that human cognition is systematic, but tensor product sentence parsing networks exhibit systematicity only if we train them to parse every member of a set S of sentences such that S is closed under systematic variation. Since it's nomically possible to construct connectionist networks that aren't systematic, the theory that we parse sentences using tensor product encodings doesn't explain why we do so systematically.

In 'Systematicity and the cognition of structured domains', Cummins et al. respond. First, they argue that various non-isomorphic cognitive domains exhibit systematicities. For instance, visual imagination exhibits a kind of systematicity, but doesn't have the same structure as propositional thought. LOT explains systematicity of linguistic understanding by positing a cognitive mechanism with a combinatorial syntax and semantics. We understand all systematic variants of a given sentence because cognition shares structure with language. Applying the same methodology across the board, the authors argue, LOT theorists ought to posit a distinct cognitive structure to account for systematicities in each non-isomorphic cognitive domain. Cummins et al. are clear that tensor product parsers don't employ representations that share structure with the sentences they parse, but argue that they do *structurally encode* sentences. The semantic structure of the encoded sentences is recoverable from tensor product encodings thereof by a recovery function, and it is '... only when cognitive systems employ such structure-preserving schemes that they can be causally sensitive to the structure of the domain cognized and thus exhibit systematicity effects' (63). But if this is right, then so is Fodor. If Smolensky's architecture is *causally* sensitive to the semantic structure encoded, in a way that explains systematicity effects, then it's an implementation of a classical parser. Tensor products, on this view, are little more than a way of storing classical structures that must be recovered by the architecture prior to cognition. By Smolensky's own admission, however, the semantic structure preserved in, and recoverable from, tensor product encodings, isn't causally available to the architecture at all, and isn't what causally explains cognitive transitions from one representation to another (Fodor 1997; Smolensky 1990). This is *why* Smolensky's architecture is non-classical. What's more, if we combine structural encoding with a recovery function that makes the recovered structure available to the architecture, there's no obvious reason why LOT architectures can't explain systematicities in non-isomorphic cognitive domains by structurally encoding them.

The appeal to structural encodings is unfortunate, because as the book progresses, it becomes clear that Cummins really wants to account for cognition without *any* appeal to semantically structured contents. Although we can use sentences to express thoughts, it's a mistake, Cummins thinks, to suppose that thoughts themselves have propositional contents. Psychological externalists fall foul of the intuitive tendency to equate the contents of cognitive states with the truth conditions of the sentences we use to express them, but what's really going on in twin-cases is that we use sentences with different truth-conditions to express states with the same cognitive content ('Methodological restrictions on belief'). As Cummins puts it, the fact that p and q are distinct propositions doesn't entail that the belief that p and the belief that q are distinct cognitive states. Representation, for Cummins, is *communicative* in a

way that reference isn't, and this, for Cummins, undermines all causal theories of content, including teleosemantic variants ('Representation and unexploited content'). Conceptual representation involves the kind of isomorphism we see in maps, and it's representing in this informative way that enables us to detect what our concepts apply to ('Truth and meaning', esp. 156–61). Representation is ubiquitous, for Cummins, because isomorphism is ubiquitous, but this will seem problematic only to those who confuse the contents of representations with their *targets*. For Cummins, an 'intender' mechanism M has the (non-teleological) function to represent a target T iff M's capacity to represent T features in a functional explanation of some capacity of a system S containing M (See 'Haugeland on representation and intentionality' for the content/target distinction; and "'How does it work?" vs. "What are the laws?"' for Cummins's classic account of functional explanation). Map-like representations aren't true or false, and need a graded notion of accuracy, depending on how closely the structure of a representation matches what it is applied to by an intender ('Representation and indication').

The question remains, however: if mental representations don't have semantic structure, how are we to explain systematicity? Cummins thinks of linguistic meanings as 'recipes' for constructing non-propositional contents ('Meaning and content in cognitive science'). However, without the illicit appeal to structural encodings and architectures sensitive to the structure encoded, there's no explanation here as to why we can construct a representation for 'Mary loves John' iff we can construct one for 'John loves Mary'. Cummins argues that meaning is an explanandum, rather than an explanans, of cognitive science, but offers little by way of explanation. Cognitive scientists don't need the compositional semantics of LOT, he claims; what they need is '... an understanding of how information can be acquired, stored and manipulated in a way that gives rise to ... the ability to emulate [belief-desire intentionality] machines and use a propositional language', ('Truth and meaning', 183). But that, as far as I can tell, is exactly why Fodor thinks cognitive scientists *do* need LOT.

Despite my misgivings about Cummins's project, this book is stimulating, engaging and – while difficult – hugely rewarding. Proponents of LOT will find much here that is challenging, even if they don't find an alternative explanation of systematicity. What's more, Cummins offers important arguments against causal and externalist theories of psychological content, based on distinguishing content from linguistic meaning, which may be endorsed even by those who agree with Fodor that cognition must have a combinatorial syntax and semantics. In a brief preface, Cummins tells us that the papers are ordered '... by topic and, where possible, by an attempt to put papers that presuppose *x* after papers that explain *x*'. This is useful, but the book is not always an easy read, and would benefit from a more detailed introduction making explicit the various interconnections between the papers. Nonetheless, this is a welcome addition to the literature, and will be of great interest to anyone who wishes to get in the head of this interesting and important philosopher.

DAVID YATES
 King's College London
 London, UK
 david.yates@kcl.ac.uk

References

- Fodor, J. 1987. *Psychosemantics*. MIT Press.
- Fodor, J. 1997. Connectionism and the problem of systematicity. *Cognition* 62: 109–19.
- Smolensky, P. 1990. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence* 46: 159–216.

The Really Hard Problem: Meaning in a Material World

By OWEN FLANAGAN

MIT Press, 2009. xvi + 288 pp. £22.95 cloth, £11.95 paper

We are all familiar with the alleged ‘hard problem’ of consciousness – explaining how the subjective world of conscious awareness could arise out of neuronal activity. Owen Flanagan, an out-and-out naturalist, is committed to dissolving this much discussed problem without any appeal to ‘spooky’ immaterial entities or properties. Instead, he proposes that the solution depends on accepting what he calls ‘subjective realism’: ‘it is simply a unique but nonmysterious fact about conscious mental states that they essentially possess a phenomenal side’ (29). Having, to his satisfaction, disposed of the puzzle of consciousness, Flanagan now turns to the arguably much harder problem that is the focus of the book’s title. The *really* hard problem is to explain ‘why and how, in the greater scheme of things, any human life matters’ (xii).

There seems to be no immediate connection between the consciousness problem and the meaning problem, except that both provide test cases for the scientific–naturalistic worldview that Flanagan espouses. He wants an account of human life and its meaning that is wholly compatible with ‘the picture of persons that emerges from neo-Darwinian theory and from the best current mind science’, according to which ‘we are fully embodied thinking-feeling animals who live and achieve meaning – if we do – in a world that is fully natural’ (61).

The answer that emerges is that human beings (or at least 21st century Westerners) live within a set of six ‘spaces of meaning’ – art, science, technology, ethics, politics and spirituality (12). These, Flanagan argues, are the arenas within which we pursue the goals of truth, beauty and goodness. This initially sounds rather Platonic (and hence, to Flanagan’s ear, suspiciously ‘spooky’); but to allay this worry, he offers us a ‘naturalised’ interpretation of Plato: “‘The good,” “the true,” and “the beautiful” are ways of gesturing at... the three fundamental and universal ways humans orient themselves in and toward the world in order to live well and meaningfully’ (40).

Consistent with his professed scientific and empirical outlook, Flanagan needs to be able to show that the human tendency to ‘orient’ ourselves in these ways is, in the end, simply a matter of the dispositions that we have naturally evolved to have. Many people suppose that the Darwinistic picture of our biological and social origins leads to a depressing and disenchanting picture of nature as ‘red in tooth and claw’; but Flanagan aptly resists this: ‘we were designed to be fit as *social* animals... to care about more than individual fitness’ (43).