

Intentionality Lite or Analog Content?

A Response to Hutto and Satne

Gerard O'Brien¹ · Jon Opie¹

Received: 6 February 2015 / Revised: 21 February 2015 / Accepted: 21 May 2015 /

Published online: 10 July 2015

© Springer Science+Business Media Dordrecht 2015

Abstract In their target article, Hutto and Satne eloquently articulate the failings of most current attempts to naturalize mental content. Furthermore, we think they are correct in their insistence that the only way forward is by drawing a distinction between two kinds of intentionality, one of which is considerably weaker than—and should be deployed to explain—the propositional variety most philosophers take for granted. The problem is that their own rendering of this weaker form of intentionality—contentless intentionality—is too weak. What's needed is a species of intentionality distinct from both the industrial-strength version beloved by philosophers and the intentionality lite recommended by Hutto and Satne. We briefly motivate and sketch this alternative, and say a few words about the account of cognition that it spawns.

Keywords Analog · Cognition · Intentionality · Mental content · Representation

Behaviourism Redux

A heresy is afoot in the philosophical community. Otherwise sensible folk are muttering that it's time to abandon the project of naturalizing mental content. Despite serious attention from a couple of generations of philosophers, we are very far from reaching a consensus on the nature and source of intentionality. One gets the feeling that everyone is ready for a break. Those of a heterodox disposition are even inclined to say that the whole project was a mistake from the get go.

Hutto and Satne are not heretics. But they do share the pessimistic analysis of the current state of play. And their response to this crisis is bracingly radical. They suggest we should consider the idea of *contentless* intentionality: that organisms engage in world-directed, action-guiding cognition *without* deploying internal structures that

✉ Gerard O'Brien
gerard.obrien@adelaide.edu.au

¹ Department of Philosophy, School of Humanities, University of Adelaide, Adelaide, SA 5005, Australia

represent, refer to, or stand in for features of the world. This way of conceiving intentionality provides the foundation for a project not of naturalizing mental content (since there really isn't any as such), but of explaining the socio-cultural origins of content-ascribing practices, especially those associated with natural language (Hutto and Satne, ms).

According to Hutto and Satne, providing a “target-based” account of intentionality is work best suited to the neo-Cartesians who occupy first-base in Haugeland’s “Intentionality All-Stars” (1990). To move things along, they recommend that first-basers try putting some familiar tools to fresh use. In particular, natural selection and learning histories might be employed not to naturalize mental content, but to explain why “certain organisms are responsive to a selective range of worldly items” (ms, p.20) and hence are capable of a form of “targeted directedness”.

We're not opposed to radical proposals, especially when a research program is genuinely stymied. But there is a fundamental problem with the idea of contentless intentionality: it's been tried before, and it doesn't work. Back then the scheme was known as “behaviourism”, rather than “targeted directedness”, but the two ideas are of a piece. Behaviourists sought to explain animal behaviour, including all the complexities of human problem solving and language, in terms of the history of stimulus–response events to which organisms (of each kind) are typically exposed. The bankruptcy of this approach consists in the fact that moment-by-moment stimuli are simply too impoverished to account for the richness, variety, and specificity of the behaviours that animals exhibit. It just isn't possible to explain the ability of evolved creatures to selectively engage with features of the environment—in other words, engage in targeted behaviour—without supposing they employ internal states that in some way *represent* those features.¹

Furthermore, many first-basers will point out that Hutto and Satne's proposal for contentless intentionality misunderstands the broader explanatory project that led them to occupy this position in the first place. The larger target was intelligence rather than just intentionality. And what took them to first base was a conviction that, so far as explaining intelligence is concerned, the computational theory of cognition is the only game in town. This is why Haugeland observes that “artificial intelligence and cognitive science belong at first base” (1990, p.394). Given the orthodox understanding of computation—“there is no computation without representation” (Pylyshyn 1984, p.62)—it is the computational theory of mind that motivates the project of naturalizing mental content. In this context, Hutto and Satne's advice to seek a contentless form of intentionality is really a call to abandon the computational approach altogether.

What about Hutto and Satne's suggestion that the tools of teleosemantics—natural selection and learning history—might be redeployed to explain an organism's capacity to respond selectively to the environment? This idea falls foul of the crucial distinction between *ultimate* and *proximate* explanation. If one aims to explain why there exist organisms capable of intelligent behaviour, then an appeal to natural selection and

¹ The corollary in the field of AI is that it is impossible to construct physical mechanisms that exhibit targeted behaviour without equipping them with some knowledge of their targets. Some AI theorists did flirt with the idea of “intelligence without representation” in the latter years of the twentieth century, and even managed to construct devices that could circumnavigate a room by bumping off the walls (Brooks 1991). But after a brief diversion, most of these researchers recognised the limits of this approach. As everyone knows, Rodney Brooks has gone back to writing code.

individual learning is apposite. Intelligence, not to put too fine point on it, promotes survival and reproduction. This is the ultimate explanation for the presence of intelligent organisms on this planet. But if one's aim is to explain how such organisms get about so intelligently—judiciously avoiding predators, cleverly seeking out known sources of nourishment, shrewdly sidestepping ruses and traps—selective explanations draw a blank. Such explanations presuppose the adaptive traits in question, they don't explain them. For proximate explanations of intelligence, one must look elsewhere, hence the popularity of the computational theory of cognition.²

But having made these critical observations, we do think Hutto and Satne are correct in one crucial respect. This is their insistence that the only way out of the current impasse is by drawing a distinction between two kinds of intentionality, one of which is considerably weaker than—and should be deployed to explain—the propositional variety most philosophers take for granted (ms, p.19). It's just that Hutto and Satne's own rendering of this weaker form of intentionality—contentless intentionality—is too weak.³ What's needed is a species of intentionality distinct from both the industrial-strength version beloved by philosophers and the intentionality lite recommended by Hutto and Satne. Happily, just such a species of intentionality is already under discussion. It has emerged from a family of approaches that conceive mental representation in *structural* or *analog* terms. When intentionality is underwritten by physical analogy it is neither language-like nor contentless. In keeping with the spirit of Hutto and Satne's proposed reconception of the role of first-basers (ms, p.19), we will in what follows briefly motivate and sketch this alternative, and say a few words about the account of cognition that it spawns.

The Rise and Fall of Propositions

At least three significant theoretical developments in the second half of the twentieth century gave rise to the industrial-strength variety of intentionality which has proven so highly resistant to naturalization.

The first was the demise of behaviourism and in its wake the vindication of folk psychology as a credible basis for explaining human behaviour. It is folk psychology that licenses our attribution of mental states with the peculiar action-at-a-distance property of being “about” or “directed at” other things. Folk psychology populates our heads with all manner of intentional states—perceptions and thoughts, pains and memories, itches and inklings—whose interactions are the proximal cause of our behaviour. In the hands of philosophers, folk psychology was transformed into propositional attitude psychology. Under this regime, the content of a mental state, such as a belief, is the proposition expressed in the noun clause of a sentence ascribing that state to an agent,

² Incidentally, this analysis does vindicate Hutto and Satne in one regard: in their diagnosis of the failings of teleosemantics. As a number of philosophers have pointed out, teleosemantic theories get the explanatory dependencies backwards. Such theories ground mental content in the (success of) the very behaviours at issue, thereby rendering it immaterial to the actual mechanisms of intelligence.

³ Indeed, as they admit (fn.7), it's questionable whether the idea of contentless intentionality is conceptually coherent in the first place, given that both content and intentionality are unpacked by philosophers using the highly technical notion of “aboutness”.

e.g., “Jerry believes that Greycat is intelligent”. This seemingly innocuous prescription is the origin of the propositional turn in the philosophy of mind.

The second development was Davidson’s (1967) seminal work on linguistic meaning. Building on earlier work by Tarski (1956), Davidson sought to fashion a truth-conditional semantics for natural languages. His key idea was that meaning should be understood in terms of truth, and truth operationalized using the kind of formal scheme familiar in first-order logic. The meaning (the truth value) of a declarative sentence thus depends on its logico-grammatical structure and the reference (the extension) of its predicates. For those schooled in Davidson’s work, it was almost obligatory to think about mental representation in the same terms.

The third development was the proposal, from the new discipline of cognitive science, that cognition is the manipulation of in-the-head symbols. Fodor attributes the idea to Turing, whose work in mathematical logic led to the Turing machine: an abstract device designed to capture the essentials of computation by pencil and paper means. Turing is often heralded as having identified the limits of computability, but for Fodor his key insight was an idea about how to mechanize rationality. A suitably programmed Turing machine takes true input symbols and reliably convert them to further symbols that also true. This kind of semantic coherence is the essence of rationality and, according to Fodor, “the most important idea about how the mind works that anybody has ever had” (1992, p.6).

Since truth-conditional semantics requires representing vehicles with a combinatorial syntax—something, arguably, only symbol systems possess—these last two developments are a match made in heaven. They come together in the classical computational theory of mind, the acme of the propositional turn, perhaps best articulated in Fodor’s defence of the language of thought (1975).

All of this is old news. But this theoretical context has set the parameters for most of the efforts towards naturalizing mental content that have followed. The name of the game is to find plausible naturalistic reference conditions for in-the-head symbols, subject to the proviso that their rule-governed combinations—our thoughts—have propositional contents, and hence are truth-evaluable. This turns out to be far more difficult than it first appeared. Most of the current proposals share the fault of presupposing a degree of intelligence in the creatures whose mental contents they seek to naturalize. This is a fault indeed, because philosophers typically use content to explain intelligence, not the converse. We’ve already seen, very sketchily, how this problem plays out in the case of teleosemantics—it is patently circular to ground mental content in biological success if one is seeking to use content to explain that success. But a similar problem besets causal theories of content too, since (as Fodor has discovered to his chagrin) it turns out that a cognitive system must already possess a great deal of perceptual and cognitive smarts in order for bits of its nervous system to causally covary with particular distal features of its environment. Again, if one requires intelligence to explain content, one can’t deploy the latter to explain the former.

The propositional turn creates other problems, beyond those associated with the naturalization of mental content. It promotes extreme intellectualism about cognition, such that even minimally intelligent creatures—for example, the desert ant (Gallistel and King 2009)—are fancied to possess a combinatorial symbol system and cognitive processes that are sensitive to the structure of those symbols. Furthermore, given that learning a language of thought would itself be a sophisticated cognitive achievement,

one faces the absurd prospect that the language of thought is not learned, but innate—that every creature comes equipped with a hardwired set of primitive symbols from which its thoughts are assembled.

This, then, is the bitter harvest of the propositional turn in the philosophy of mind. We find ourselves with no credible way of naturalizing mental content, and a computational account of cognition that is at best problematic. No wonder that many philosophers, and even a few cognitive scientists, have begun to avow some form of anti-representational or anti-computational doctrine.

Analog Representation

Human languages undeniably provide a powerful symbolic representational medium. But it would be a mistake to underestimate the significance of analog forms of representation, such as pictures, maps, graphs, and diagrams. Analog representing vehicles differ from symbols in two principal respects. First, their contents are determined locally, by intrinsic structural properties of the vehicles themselves. For this reason they can be used to represent anything which shares those properties and hence resembles them in the relevant respect(s). For example, variations in brightness on an X-ray image can be used to represent analogous variations in bone or tissue density. Second, while their contents are often complex, and highly structured, analog vehicles are not inherently suited to representing propositional contents.

Although sometimes overlooked, there is a long and venerable history of attempts to understand the human mind in analog terms. Aristotle was the first to articulate the idea, but it was later tweaked by the Scholastics, toyed with by Descartes, and then resuscitated by the British Empiricists. In more recent times it has been developed, amongst others, by Craik (1943), Johnson-Laird (1983), Cummins (1996), Churchland (2012), and ourselves (O'Brien and Opie 2004). All of these theorists dispense with the view that cognition relies on symbolic *descriptions* of worldly stuff in favour of the idea that mental representations are *analog*s of their objects, with which they share physical structure of some kind.

Given the recent fortunes of descriptivist approaches, it's not surprising that Hutto and Satne have argued for doing without content altogether. What is a little surprising is that philosophers haven't paid more attention to the possibility that biological intelligence is founded on analog forms of representation. There are several reasons for this. One is the deep-seated influence of propositional attitude psychology. Although many philosophers allow that analog representations may play a role in the operation of the brain's various input and output systems, they argue that such representations must ultimately "interface" with central systems that are in the business of fixing our beliefs and regulating our responses in light of our desires. In other words, the real work of rational agency relies essentially on the activity of representing vehicles with propositional contents.

A second reason is the charge that the naturalizing project is likely to be no more successful when its target is non-propositional mental content than it has been with propositional mental content. Hutto and Satne argue, for example, that the problems associated with teleosemantics—which they observe is widely regarded as the most promising naturalizing strategy—apply just as much to a

story that replaces truth conditions with something more deflationary, such as accuracy or veridicality conditions (ms, p.9).

A third, related, reason is that any move away from propositional content will bring with it a destructive form of content indeterminacy. References and truth-conditions anchor content in something determinate (particulars and states of affairs, respectively). Accuracy and veridicality, on the other hand, are much fuzzier notions, allowing a degree of uncertainty in content determination that makes it effectively impossible to specify what contents are being conveyed by representing vehicles.

These are quite legitimate concerns for philosophers who operate within the “classical” milieu. But what has not been widely appreciated to date is how profoundly an analog framework changes our thinking about mental content and its role in cognition.

To begin with, from an analog perspective, linguiform styles of representation are among the explananda of cognitive science, rather than part of the ground floor furniture of intelligence. This renders the capacity to represent propositions an emergent feature of certain cognitive systems, perhaps only brought to full fruition in language users. The working assumption here is that cognition in general, not just perception or motor processes, involves representations with non-propositional contents. In this regard, there is no question of an interface between fully-fleshed conceptual thought and non-conceptual aspects of cognition. The only interface that needs explaining is the one that bridges our thoughts and our words. Of course, it remains an open question whether this take on cognition is viable, but the classical computational paradigm has a rival, in the form of connectionism, which has already begun to demonstrate the power of neural systems to process information encoded by analog means.

As for the problems facing teleosemantics, we agree with Hutto and Satne that these are significant. Significant, but beside the point. The content of an analog representing vehicle isn't determined by learning or selection history, but by a structural resemblance between that vehicle and its object. Notice that this kind of resemblance depends only on intrinsic properties of the vehicle (and object) and not on any causal connection between them. For this reason, if we conceptualize mental representation in analog terms, a promising strategy for naturalizing mental content is already at hand. This strategy avoids the mistaken, if largely inadvertent, tendency to explain mental content in terms of intelligence, and also side-steps the epiphenomenalism which also plagues the project of naturalizing symbolic content.

Finally, it is true that resemblance relations don't generate highly determinate content. However, it begs the question against analog accounts of the mind to suppose that the degree of determinacy achievable by analog means will be insufficient to explain human intelligence. Publically available analog representations are determinate enough to support some pretty sophisticated abilities, such as the ability to diagnose a cancer via quick inspection of an X-ray, and this level of specificity may suffice for many cognitive processes. Where greater specificity is required, we suspect natural language plays a role, and hence may be accounted the result of inherently non-propositional processes, not the outcome of an already determinate propositional encoding.

Conclusion

We congratulate Hutto and Satne for so eloquently articulating the failings of most current attempts to naturalize mental content. Were it not for the alternative, we would be inclined to agree that intentionality is a social phenomenon pure and simple, which therefore requires nothing deeper than a socio-cultural explanation. However, we've argued that the propositional turn in the philosophy of mind has made it difficult to recognise (or easy to overlook) another significant approach to intentionality. That approach, like its principal rival, takes its inspiration from folk wisdom about the mind. But instead of finding there a story writ in language-like strings of symbols, it discovers a canvas painted with mental images and other more abstract mental picturings. It remains to be seen whether the analog approach to mental representation and cognition will win the day, but for fans of robust, causally efficacious mental contents it may be the only real hope. Instead of opting for intentionality lite, as Hutto and Satne advise, we therefore recommend a generous helping of full-strength analog content.

References

- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–59.
- Churchland, P. (2012). *Plato's camera: How the physical brain captures a landscape of abstract universals*. Cambridge: MIT Press.
- Craik, K. (1943). *The nature of explanation*. Cambridge: Cambridge University Press.
- Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge: MIT Press.
- Davidson, D. (1967). Truth and meaning. *Synthese*, 17, 304–23.
- Fodor, J. (1975). *The language of thought*. New York: Harvester Press.
- Fodor, J. (1992). The big idea: can there be a science of the mind? *Times Literary Supplement*, 5–7.
- Gallistel, R., & King, A. (2009). *Memory and the computational brain*. Oxford: Willey-Blackwell.
- Haugeland, J. (1990). The Intentionality All-Stars. *Philosophical Perspectives*, 4, 383–427.
- Johnson-Laird, P. (1983). *Mental models: Towards a cognitive science of language, inference and consciousness*. Cambridge: Harvard University Press.
- O'Brien, G., & Opie, J. (2004). Notes towards a structuralist theory of mental representation. In H. Clapin, P. Staines, & P. Slezak (Eds.), *Representation in mind: New approaches to mental representation*. Amsterdam: Elsevier.
- Pylyshyn, Z. (1984). *Computation and cognition*. Cambridge: MIT Press.
- Tarski, A. (1956). The concept of truth in formalized languages. In *Logic, Semantics, and Metamathematics*. Oxford: Oxford University Press.