

Content Pragmatism Defended

Dimitri Coelho Mollo^{1,2} 

Published online: 23 November 2017
© The Author(s) 2017. This article is an open access publication

Abstract In the literature on the nature and role of cognitive representation, three positions are taken across the conceptual landscape: robust realism, primitivism, and eliminativism. Recently, a fourth alternative that tries to avoid the shortcomings of traditional views has been proposed: content pragmatism. My aim is to defend pragmatism about content against some recent objections moved against the view. According to these objections, content pragmatism (a) fails to capture the role played by representation in the cognitive sciences; and/or (b) is an unstable view that ends up collapsing into one of the traditional alternatives. I argue that those arguments fail. I show that content pragmatism has as much claim to descriptive adequacy as the traditional theories. Moreover, I defend the robustness of the view by arguing that it does not collapse into any of the traditional positions. Content pragmatism therefore offers a valid and coherent account of the nature of representational content.

Keywords Representation in cognitive science · Representational content · Content pragmatism · Philosophy of cognitive science

1 Introduction

For much of the centuries-spanning philosophical debate on the mind, the problem of the nature of representation and of representational content has had pride of place. This central role has not been jeopardised by the advent of the cognitive sciences, representation being one of the foundational notions behind that scientific endeavour. The philosophical problem of representation has become, therefore, scientifically interesting and pressing.

This paper is concerned with two central issues regarding the nature and role of representation in cognitive science, namely its ontological status—*are representations and representational contents real features of cognitive systems?*—and its explanatory value—*are representations explanatorily useful, and, if so, how?* I will focus on the notion of representation and content of subpersonal states—i.e. cognitive states short of beliefs, desires, and other propositional attitudes—which plays an important role in the cognitive sciences. To avoid confusion, I use the expression ‘representational content’ to refer to the contents of subpersonal states, which I normally refer to as ‘cognitive states’, while I reserve the expression ‘intentional content’ to refer to the contents of personal states, such as propositional attitudes, which I refer to as ‘intentional states.’¹

Traditionally, three mutually exclusive views about the ontological and explanatory status of representations were supposed to be the only conceptual options on the table: robust realism, eliminativism, and primitivism.

Roughly, robust realists hold that representations and representational content are real features of cognitive systems, which can be given reductive explanations in purely

✉ Dimitri Coelho Mollo
dimitri.coelho_mollo@kcl.ac.uk;
dimitri.coelho.mollo@hu-berlin.de

¹ Department of Philosophy, King’s College London, Strand, London WC2R 2LS, UK

² Institut für Philosophie, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany

¹ A similar distinction can be found in Cummins (1989).

naturalistic and non-representational terms (e.g. Fodor 1975; Dretske 1981; Millikan 1984). On this view, some inner states of cognitive systems have representational content due to their standing in some special natural relation to what they represent. Representations and contents are explanatorily useful, at least in part, because they consist of properties that cognitive systems actually possess.

Eliminativists, on the other hand, believe that our best cognitive sciences will do away with appeal to representation and content, producing instead explanations at a purely functional or neurophysiological level. Eliminativists hold that propositional attitudes and intentional content are probable candidates for elimination from future cognitive science (e.g. Churchland 1981; Stich 1983), while more radical eliminativists have claimed that even representational content will be eliminated (e.g. Van Gelder 1995; Hutto and Myin 2013). The issue of how to give a naturalistic reduction of representation thereby does not exercise the eliminativist, though they must show that equally satisfying explanations of cognitive phenomena can be provided without appeal to representations.

Primitivists, finally, hold that the notions of representation and content play a crucial role in the successful and fruitful cognitive sciences, and this suffices for taking them as real entities in cognitive systems (e.g. Burge 2010). They are scientific primitives, which need not be themselves further explained, or naturalised. Primitivists have been accused of holding a form of dualism, since they seem to consider representation and content to be notions that cannot be naturalised. I think this accusation is misguided. Primitivism in the lines of Burge (2010) is monist naturalism—it takes representation to be a scientific posit in no need of naturalisation, as much as fundamental entities posited in physics do not call for naturalisation, insofar as they are part of a successful and fruitful science. This view does not entail dualism, as it is compatible with materialism.

Recently, a fourth alternative, content pragmatism, has been forcefully put forward by Egan in a series of papers (1999, 2009, 2010, 2014b). She argues that representation and content should be seen as ineliminable parts of cognitive science, *contra* the eliminativist; but she rejects the view that those notions capture anything real in cognitive systems—*contra* realists and primitivists—content being rather part of an explanatory gloss dependent on pragmatic considerations.

In the last few years, Egan's attempt at building a fourth alternative has come under attack by Neander,² Ramsey,³

and Bechtel (2016). My aim in this paper is to defend content pragmatism from recent objections, showing that it is a valid alternative to the traditional positions. I do not want to argue that content pragmatism is true, but only that it is, those objections notwithstanding, coherent and motivated.

Here is how I will proceed in what follows. In Sect. 2, I present the main tenets of content pragmatism, without committing to the details of Egan's view. I examine four lines of objection to content pragmatism in Sect. 3, according to which the view is not descriptively accurate; and/or collapses either into primitivism; eliminativism; or realism. I argue that all four lines of objection fail, and thus the cogency of content pragmatism is preserved. In consequence, content pragmatism should be considered to be at least on par with traditional positions: a serious contestant in shedding light on the nature and role of representation and content in the cognitive sciences.

2 Content Pragmatism

The starting point for Egan's content pragmatism is the role the notion of representation is supposed to play in the cognitive sciences. Her project is that of describing and making sense of the commitments of our best sciences of the mind, with special attention to their more promising (and mainstream) branch: computational cognitive science.

According to Egan, computational explanations in cognitive science are function-theoretic in nature: they unveil the mathematical functions computed by internal mechanisms of cognitive systems that help explain behavioural success. One of Egan's favourite examples is Marr's theory of how the early visual system detects edges by calculating the rate of change in the intensity of light reaching the retina.⁴ The ability to detect edges is enabled by the fact that the early visual system computes a mathematical function: the Laplacian operator convolved with the Gaussian operator. This mathematical function is general, inasmuch as it is not required that inputs be light intensities and outputs edges, and is applicable to several other domains; it belongs moreover to a class of well-understood mathematical tools. Egan claims that this is the primary characterisation of cognitive mechanisms in computational cognitive science: what explanations provide are purely computational, function-theoretic characterisations of the computations carried out by cognitive systems and parts thereof.

² 'Why I'm not a Content Pragmatist', unpublished paper presented at the 2015 Minds Online Conference. Available online at: <http://mindsonline.philosophyofbrains.com/wp-content/uploads/2015/09/Why-I%E2%80%99m-not-a-Content-Pragmatist.pdf>.

³ 'Against Representation Deflation', presentation at the conference 'Mental Representations, the foundation of Cognitive Science?'

Footnote 3 (continued)

Ruhr-Universität Bochum, 21–23 September 2015. Available online at: <http://www.youtube.com/watch?v=mgq6uo9cDdY>.

⁴ See Marr (1982), Egan (2010, 2014b).

Importantly, the function-theoretic characterisation does not involve distal content. The system is explained, and its parts individuated, in terms of the algorithms implemented in computing the relevant explanatory mathematical function. Distal representational contents are not mentioned in the theory itself, the latter being characteristically ‘environment-neutral’ (Egan 2014b, p. 122). The computational mechanism that computes the Laplacean convolved with the Gaussian does so regardless of the environment it finds itself in, and therefore regardless of the relations that input systems bear to things in the world. Were the inputs to the mechanism transduced from sound waves rather than light intensities, the mechanism would still compute the same mathematical function over those inputs. These features are not restricted to this particular example, but are supposed to be generalisable to most if not all endeavours in computational cognitive science.⁵

Crucially, for Egan, the relations between cognitive mechanisms and the environment, which typically are relied upon to determine representational content—according at least to mainstream theories of content, which are externalist—play no role in individuating cognitive mechanisms. Computational cognitive science individuates cognitive mechanisms not by their representational contents, but rather by their computational, function-theoretic, environment-neutral characterisation.⁶

A place for representational content in computational cognitive science is nonetheless preserved, for appeal to content is unavoidable in explanations of cognitive abilities. The *explananda* of cognitive science are typically tasks that are pre-theoretically characterised in distal terms, in terms of capacities sported by organisms in their interactions with their environments. Success in those tasks is presupposed, being what cognitive science tries to explain. In order to make clear how the computationally-individuated mechanisms unveiled by computational cognitive science can be the *explanantia* of those distally-individuated *explananda*, an interpretation of the computational goings-on in terms of distal features of the environment is called for. That interpretation involves representational content: states and processes of the computational mechanism are seen as representing entities and processes in the world. In this way, a connection

is established between the computational mechanism and the cognitive task it is meant to explain.

Importantly, such an interpretation is not part of the computational theory proper. Rather, the interpretation of cognitive mechanisms in representational terms is part of what Egan variously calls the ‘explanatory’ (2014a, b), ‘cognitive’ (2010), or ‘intentional gloss’ (2014b). The gloss, by including reference to distal content, allows us to see how computing certain mathematical functions in a certain environment constitutes the performance of a cognitive task.

Beside playing the role of ‘connective tissue’ between computational story and pre-theoretical *explanandum*, the gloss also has a heuristic value. It allows theorists more easily to grasp what the computational processes in the cognitive system are doing, and how the information flow inside the system proceeds. By seeing those processes as being carried out over states with distal representational content, the computational goings-on in the system become easier to understand and synthetically describe.

Other pragmatic considerations also contribute to the fixation of the content ascribed to cognitive states and processes (Egan 2014b, pp. 125ff.). Ascribed contents must be salient and tractable, and should not be too disjunctive or opaque. The representational interpretation may not perfectly match the computational description of the system—one-to-one mappings between computational states and contents may be given up, provided that this increases expository clarity. Moreover, depending on the level of generality that is supposed to be achieved by the cognitive gloss, more or less general contents may be ascribed.

The same computational system may contribute to different cognitive tasks, involving different types of input and output (e.g. visual vs. auditory), in different contexts. Depending on explanatory aims and interests, content ascription may try to capture the specific contributions made in the different contexts, thereby, for instance, ascribing visual contents in one type of case and auditory contents in another; or it may strive for generality by ascribing disjunctive contents, e.g. both visual and auditory; or it may strive for yet greater generality by ascribing proximal representational contents which are fully independent of context, e.g. changes in intensity of the input signal.

Individuation in computational cognitive science involves primarily the function-theoretic characterisation, which remains constant with change of explanatory context, embedding environment, different weight given to different pragmatic considerations, and so forth. Representational content, on the other hand, may vary with changes in each of these factors. Different content ascriptions do not affect the individuation of the computational states and processes that are the *explanantia* of computational cognitive science. What is essential to the *explanantia* is the computations

⁵ Egan (2010, 2014b) draws on several other examples from computational cognitive science, such as Shadmehr and Wise’s (2005) theory of motor control.

⁶ In her more recent work, Egan (2014b) has argued that computational states have one kind of content essentially, i.e. mathematical content. I think this move introduces considerable and unnecessary complications, and should therefore be rejected (cf. Piccinini 2015, pp. 137–8). I will thus stick to formulations of the view that do not posit this new kind of content.

they perform, not the representational contents that may be ascribed to them.

In sum, representational content, though not part of the purely function-theoretic characterisations employed by computational cognitive science, plays a pragmatic role in allowing us to understand how cognitive computational systems work, and how their workings amount to the performance of pre-theoretically characterised cognitive tasks. Content is thereby fixed not by means of natural relations between certain cognitive states and entities in the world, as per the robust realist view, but rather by pragmatic considerations motivated by the explanatory aims of theorists, as well as the heuristic value of ascribing content to computational states and processes.

3 Content Pragmatism Assessed

In recent years, partly due to the frustration felt in some philosophical quarters with the project of naturalising content, and partly thanks to Egan's compelling advocacy, content pragmatism has attracted considerable attention, as well as criticism. The most vocal direct opposition has been put up by William Ramsey, Karen Neander, and Bechtel (2016).⁷

Ramsey and Neander make use of the same strategy in arguing against content pragmatism. As we have seen, the view is supposed to provide an alternative path in the philosophical attempt of understanding the nature and role of representation and content, eschewing robust realism, eliminativism, and primitivism—the traditional positions, often taken to be exhaustive of the conceptual landscape. Neander and Ramsey argue that this is an illusion: content pragmatism is not a stable position, and thus ends up collapsing into one of the traditional views.

Bechtel, on the other hand, focuses on the putative lack of descriptive adequacy of content pragmatism in capturing the way cognitive scientists use the notion of representation in their theories. I will start with Bechtel's critique, and then move on to Neander's and Ramsey's.

3.1 Descriptive Accuracy

Bechtel (2016) argues that content pragmatism is at odds with the practices of cognitive science. It seems, indeed, that cognitive scientists make use of representations not as explanatory glosses to be provided once their theories are complete, but rather as research-guiding hypotheses about the nature of states and processes in the cognitive system. Bechtel examines one case study, the discovery

and development of theories of spatial representation in the rodent brain—which came to be known as 'cognitive maps' (Tolman, 1948; O'Keefe and Nadel, 1978). In the decades that followed the original discovery of place cells, grid cells, head-direction cells and other processes that play a role in spatial representation and navigation have been found (cf. Moser et al., 2008). Bechtel's aim is to show, *contra* Egan, that the positing of representations, rather than working as a mere gloss, actually involves strong ontological commitments, drives new research endeavours, and deeply informs theorising in cognitive science.

... treating brain processes as representations is foundational to these research projects as they have been pursued. The research pursuits are focused on identifying what neural processes serve as representational vehicles and especially determining what they represent. The research efforts are designed to answer these questions and it would be difficult to understand why the researchers pursue these projects if their goal was not to identify representations and determine what they represent. (Bechtel 2016, p. 1289)

Bechtel then describes in great detail the process of discovery and theoretical development that have marked scientific work on place, grid, and head-direction cells in the past four decades. He uses this careful reconstruction of progress in this area, what motivated it and what it was directed at explaining, to argue that a commitment to the reality of representations was essential for that scientific project. The whole process would have seemed absurd and meaningless were talk of representation to be taken as merely providing an explanatory gloss, rather than as capturing real entities in the rodent cognitive system. Though Bechtel's careful treatment of the discovery of place, grid, and head-direction cells is historically and philosophically interesting, I think that he fails in his attempt to undermine content pragmatism by means of this case study.

Bechtel is right that *prima facie* content pragmatism is not descriptively accurate. As he points out, most of cognitive science traffics in representational talk. However, more must be shown if his line of objection is to be successful: he must establish that talk of representation is justified and substantive, and not a mere matter of scientific heritage. Ramsey (2007), examining this issue, concludes that the latter is more often the truth: most of cognitive science works with an excessively weak notion of representation based on mere causal correlation, what he dubs the 'receptor notion'.⁸ The receptor notion, he then argues, fails to

⁷ Sprevak (2013) presents and criticises a related view—fictionalism, treatment of which is beyond the scope of this paper.

⁸ This is indeed how Sullivan (2010) defines the notion of representation at play in cognitive neuroscience and neurobiology. She then goes on to argue that in many cases talk of representation only plays a minimal, heuristic role, in research and explanation.

distinguish representations from mere causal relays. Most talk of representation in cognitive science could be replaced with talk of causal relays without loss of explanatory value. Appeal to receptor-representations, for Ramsey, does not buy us anything explanatorily useful in comparison to appeal to causally mediating states, and should thereby be excluded from theorising in cognitive science.

Ramsey's sceptical results about the presence of a substantial notion of representation in contemporary cognitive science have been compellingly countered by Sprevak (2011) and Shagrir (2012), at least for what regards some areas of enquiry. At any rate, the cautious note that Ramsey urges on us remains in force: talk of representation in the cognitive sciences, though widespread, may be largely empty, and we must examine in each case whether it is playing its proper explanatory role. The fact that cognitive scientists routinely deploy the term 'representation' should not thereby lure us into believing that they are actually positing the existence of robust representations in cognitive systems—they might be simply talking about causal relays, and using representational talk for its heuristic value.

Furthermore, Bechtel's case study is fully compatible with the tenets of content pragmatism. As the content pragmatist points out, the notion of representation plays a crucial role in cognitive science inasmuch as it allows us to connect the computational *explanantia* with the distally-individuated *explananda* of cognitive science. It is natural for cognitive scientists to start with the cognitive task they want to explain, and which is individuated in terms of the relationships between organism and environment—in the case at hand, the ability of rodents to navigate space. That choice of explanatory target is essential in directing scientific research toward the discovery of the computational mechanisms behind the feat.

The fact that an 'early and integral step' in cognitive science is that of "using investigations of content to help identify the vehicles"⁹ is thereby in harmony with content pragmatism. The cognitive task, which is externally-individuated, is typically the starting point for research in cognitive science: it is the choice of *explanandum* that constrains the scientific investigation that follows. This does not mean that talk of representation must be taken as ontologically committing; it may well be, as the content pragmatist insists, a useful way of connecting goings-on in the cognitive system with the behaviour to be explained—an explanatory gloss.

After determining the explanatory target, cognitive scientists try to figure out what internal computational processes can explain behavioural success—what the computational vehicles are, *e.g.* in spatial navigation. This involves discovering which stimuli neurons respond to, as well as their

causal relationships with other neurons and networks. This allows cognitive scientists to individuate the neurons and networks that are good candidates for realising the computational mechanisms behind the successful performance of the cognitive task. Once plausible candidates are found, it is only by taking them to be representing (or, more weakly, carrying information about) entities and processes in the environment that the connexion between computational mechanisms and cognitive task target of the explanation can be established.

The fact, repeatedly appealed to by Bechtel (2016), that cognitive scientists are interested in finding out what neural states represent is perfectly in line with what the content pragmatist holds. It is indeed a central step in trying to individuate the computational mechanisms that explain successful cognitive behaviour. The content pragmatist does not deny the importance of the notion of representation in the 'context of discovery.'¹⁰ The gloss that representation provides not only allows us to keep the explanatory target in view, but it also makes it simpler to understand the role that computational goings-on in the cognitive system play in successful completion of the task. However, this does not entail that, in the 'context of justification,' an ontological commitment to cognitive representations is necessary. It does not follow from anything that Bechtel claims about how research in rodent navigation has proceeded that a robust realist position about representation needs to be accepted in order to do justice to scientific practice. That practice is compatible with the tenets of content pragmatism as well. Therefore, I submit, the content pragmatist escapes Bechtel's objection from descriptive accuracy.

3.2 Pragmatism and Primitivism

The second line of attack against content pragmatism involves claiming that the view is too unstable, its instability making it collapse into one of the traditional views of representational content. Neander worries that content pragmatism might fall into some form of dualism, or rather primitivism. For the content pragmatist, the argument runs, content ascription depends on explanatory aims and interests, which are ultimately based on the intentions of humans. However, intentions are states with content. For the account to work in a non-circular way, the contents of the intentional states grounding the aims and interests of humans seem to have to be taken as primitive.

Egan insists that her account is targeted only at representational content (short of beliefs and desires), and not at the intentional content of propositional attitudes. The content of

⁹ Bechtel (2016, p. 1291).

¹⁰ I am using the term, and its usual companion, in a somewhat loose sense.

cognitive states would be parasitic on intentional content, for it is the latter that helps determine the content of the former. In this way, she can avoid committing to primitivism, leaving open the question of the nature of intentional content. But then, Neander presses on, content pragmatism accounts for only a part of the problem of content, being completely silent on intentional content, on which representational content depends. The content pragmatist thus fails to offer a complete theory of how content is possible in our world.

Moreover, they would have to subscribe to a heterodox view of how the project of explaining representation should go. The strategy typically taken to be more promising has been to try and give an account of simpler cognitive states, so as to later build a more encompassing theory on this basis, perhaps by adding extra factors in order to include intentional states.

These are fair worries, I believe, but they fall short of undermining content pragmatism, let alone lead it to collapse into dualism or primitivism, or even robust realism about intentional content. The view does not collapse into dualism or primitivism because the kind of content it appeals to in helping to determine representational content is intentional content. Being a different kind of content, intentional content is plausibly to be explained by a different theory, targeted to that domain of phenomena.¹¹ Moreover, the content pragmatist need not endorse robust realism about intentional content. Alternative views of intentional content are available. For instance, Dennett's (1981, 1987) Intentional Stance view is supposed to account for the intentional content of propositional attitudes without thereby endorsing robust realism about intentional content.¹²

This latter line of reply makes the prospects of content pragmatism partially depend on the fate of non-robustly realist theories of intentional content, such as Dennett's. This is, though, no reason to reject content pragmatism: the game is still on as to which kind of approach to intentional content will eventually prevail.¹³

3.3 Pragmatism and Eliminativism

Content pragmatism and content eliminativism share one crucial element: both views deny robust realism about representational content. However, content pragmatism and content eliminativism are at odds in an equally crucial point: while the latter advocates elimination of the notion

of representation from cognitive science as lacking explanatory value, the former claims that cognitive science cannot do without it. According to the content pragmatist, the explanatory gloss that accompanies cognitive computational explanation cannot be eliminated. For the gloss in terms of representational content plays an essential role in allowing cognitive science to explain what it sets out to explain: behavioural success in cognitive tasks.

As Neander points out, content pragmatism would collapse into eliminativism were the pragmatist to hold that future advancements in the cognitive sciences would discharge the need for the gloss. In that case, the explanatory gloss would be a non-essential epistemic device that we now need because of the relatively early stage of scientific development we are in. In the future, with a fuller picture and a better understanding of how it all works computationally and neurally, the explanatory gloss, with all its representational baggage, would be done away with. If content pragmatism corresponded to that view, then it would indeed be a species of eliminativism.

On the contrary, content pragmatism sees the explanatory gloss as essential to cognitive science, something that further scientific developments will never eliminate. Though representational content is not part of the computational theory proper, it is an essential part of the explanatory gloss, without which it would be impossible to connect the *explanantia* and *explananda* of cognitive science. Therefore, content pragmatism does not give way to eliminativism.

3.4 Pragmatism and Robust Realism

The most insidious challenge to content pragmatism is avoiding its collapse into some form of robust realism about representational content. In order to assess this challenge in detail, it is worthwhile to recall the essential differences between realists and pragmatists about content.

Realists hold that representational content captures some observer- or mind-independent feature of cognitive systems.¹⁴ Even though contents are normally posited in the course of attempts to explain cognitive phenomena, their explanatory value is supposed to hinge on their corresponding, to some degree of accuracy, to features that cognitive systems possess, and would possess even in the absence of any sentient being having cognitive states about those features, as well as the explanatory purposes and practices in which those notions are employed.

¹¹ See also Cummins (1989, pp. 12–3, 88).

¹² In a reply to Neander during the MindsOnline 2015 Conference, Egan seemed attracted to a view on these lines. A related account, also compatible with this line of argument, is the measure-theoretic view of propositional attitudes put forward by Matthews (2011).

¹³ See Haugeland (1990) for an only partially out-dated overview of the options.

¹⁴ Of course, the notion is not independent of sentient beings and their cognitive states insofar as it captures properties of the cognitive systems of those beings. But this is not the kind of mind-independence that is at issue here.

In contrast, for a content pragmatist, ascription of contents to cognitive states depends on conceptual apparatus that we as cognitive scientists impose on the world, without commitment to their being faithful to its nature and organisation. Contents may be ascribed to systems as part of our explanatory practices in order to simplify explanations, to allow us more easily to grasp interesting connections, or merely because regarding some systems as representing proves to be heuristically useful in helping us make sense of the goings-on in cognitive systems. The positing of content is done with an eye to the fruitfulness and heuristic value of using this notion, rather than to its corresponding to some observer-independent feature of the world. According to pragmatists, in other words, it may well be that the notion of content does not capture anything that cognitive systems possess independently of the explanatory interests and practices that we bring to bear in our scientific understanding of the world.

I will examine four arguments that purport to show that content pragmatism, in some way or another, fails to provide a principled alternative to realism about content. I will try and defuse each of them, securing the coherence and interest of content pragmatism as a genuine competitor in the search for a satisfying theory of the nature and role of representational content. Once again, please note that I am not arguing that content pragmatism is the correct view, only that extant objections to the view fail to undermine its coherence.

3.4.1 Neander's Compatibility Argument

Neander argues that the main claim of the content pragmatist, namely that content *ascription* is relative to explanatory aims and interests, is perfectly compatible with robust realism about representational content. What the realist must hold on to is that there is, as a matter of fact, an ultimate answer about what content those cognitive states actually have. Though different content *ascriptions* are possible, given pragmatic considerations, the realist about content must insist that there is one (or at most a few) content ascriptions that actually get it right, that actually say something true about that cognitive state. The fact that content *ascriptions*—the emphasis is present throughout Neander's argument—are variegated is no reason to embrace pragmatism about content. Competing ascriptions are compatible with content realism. I think Neander is right.

However, the presence of competing ascriptions is not the crucial tenet of content pragmatism. The content pragmatist claims, much more strongly, that there is nothing over and above content ascription. There is no fact of the matter about what a cognitive state represents that is separate and independent from practices of ascription informed by pragmatic considerations. It is not the fact that different content ascriptions are possible given different explanatory interests

that should motivate the content pragmatist. As Neander correctly points out, this is compatible with the realist picture. What is incompatible with that picture is the further claim that there is no fact of the matter about what a cognitive state represents—different content ascriptions are possible, all of which may play crucial roles in different explanatory glosses (or even in one and the same), without any involving ontological claims about cognitive states having representational contents.

This is, of course, just conceptual clarification about the tenets of content pragmatism. It does not tackle the motivations for upholding the view, nor defends its cogency.

3.4.2 Ramsey Against the Argument from Environmental Neutrality

An argument that plays an important role for Egan in motivating content pragmatism is the argument from environmental neutrality of computational individuation. Egan (2009, 2014b) shows that computational mechanisms that contribute to a cognitive task in one type of organism in one type of environment (say, computing shape from shading) can be transferred to a different environment or a different organism and play a different role (say, in audition), despite the fact that they still compute the same mathematical function. Thus different representational contents are ascribed to one and the same computational mechanism when embedded in different organisms and environments. This fact, Egan believes, argues against robust realism about content—it would show that representational contents are not essential to computational mechanisms, since ascription of the former varies wildly when the latter are kept fixed, and only the embedding context is changed.

Ramsey has recently argued that this argument will not do. I believe that he is right—as it stands, Egan's argument does not provide reasons to prefer content pragmatism over content realism. For the fact that the same computational mechanism, by computing the same mathematical function, can have different contents when embedded in different organisms and different environments, does not jeopardise its having fairly determinate contents in the actual organisms it is embedded in. Computational mechanisms have been plausibly selected for the functions they compute because they contributed to successful behaviour in a specific species of organism and type of environment—alternative contents they might have had given different evolutionary stories or different embedding contexts are thereby irrelevant.

To avoid the objection, the content pragmatist must be more radical in their claims. They must claim that representational content is ascribed to different computational states and processes, and different combinations thereof, given different contexts and cognitive tasks in the same type of organism and the same type of environment; and this due

to heuristic and pragmatic reasons tied to our interests and practices. If this is so, then content is not essential to computational states even in the same organism in the same environment—an outcome difficult to square with robust realism about content, given the considerable indeterminacy of content and of vehicle individuation across contexts that follows.¹⁵ This is largely an empirical hypothesis, which only empirical work can vindicate or disprove.¹⁶ It is though to some extent also a warning: unproblematically endorsing the robust realist view of representation and vehicle individuation can skew theorists' and scientists' interpretation of the available data, generating puzzles and questions that would be misplaced given a more flexible (pragmatist) view of the role of representational content in cognitive science.

3.4.3 Arguing from Despair

A further objection to content pragmatism has it that it is *prima facie* radically revisionist of scientific practice. Robust realism about content, it is argued, is the default position, and thereby the burden of proof lies with the pragmatist. The pragmatist, that is, must show that content realism is false, or at least questionable, calling hence for exploration of alternative paths.

As already hinted in Sect. 3.1, this line of argument is problematic. Scientific practice, though apparently committed to the ontological reality of representational content, is compatible with the move recommended by the content pragmatist, namely relegating content to the explanatory gloss. Taking scientists at their word can be misleading—talk of representation might not really refer to robust notions of representation; and, argues the content pragmatist, talk of representation need not be ontologically committing. So the claim that content pragmatism is revisionist is too quick, and too rash. Indeed, Egan's motivation for developing her version of content pragmatism stems from a close examination of work in computational psychology and neuroscience done by the likes of David Marr, Reza Shadmehr and Steven Wise. Content pragmatism has as much claim to being the 'default scientific position' as content realism, appearances to the contrary notwithstanding. Scientists are typically satisfied with making use of conceptual apparatus that lead

to successful theorising and prediction, without worrying enough about the ontological status of that apparatus to allow us to determine whether they would rather endorse content realism or content pragmatism.

Most importantly, extant robust realist theories of content face enormous difficulties in delivering what they set out to deliver. Robust realists hope to provide naturalistic conditions that bestow fairly determinate content on cognitive states and processes (e.g., Dretske 1981, 1988; Millikan 1984; Fodor 1987). Representations are seen as identifiable, repeatable, and composable cognitive structures that have specific contents due to their standing in special naturalistic relations to what they represent in the world (Fodor, 1975). However, such theories have so far failed to provide accounts of content that satisfactorily deal with liberality of representational status, indeterminacy problems, and misrepresentation. This is, I take, the most powerful motivation for developing an alternative to robust realism about content on pragmatist lines—what Neander calls the 'argument from despair'.

I will not repeat here the rich debate that, especially in the 80s and 90s, explored robust realist theories of content and their shortcomings. I will rather just point to the well-known problems that plague these theories.

First, there are issues regarding the liberality of the notion of representation that follows from robust realist accounts. The proposed naturalistic reductions of representation make it so that states and processes of non-biological systems, as well as very simple biological systems, count as representational. There is nothing specifically cognitive or mental about those states that realist theories identify as representational (Burge, 2010; Morgan, 2014). It follows that representations are widespread and have no special cognitive value, 'debasement' the notion of representation and hurting its explanatory value in the cognitive sciences.

Second, and more damningly, robust realist theories have failed to provide satisfactory conditions for bestowing determinate representational content.¹⁷ Several indeterminacy problems have been discussed in the literature, many of which have not been satisfactorily solved, including: the problem of error (Fodor 1984), the problem of distality (Dretske 1986), Quinean indeterminacy (Gates 1996), and functional indeterminacy (Fodor 1990).

Indeterminacy of content is widely seen as a serious shortcoming for theories of content. Indeterminacy jeopardises the explanatory role that representation is supposed to play in the cognitive sciences. If we cannot determine at least to some level of precision the content of a representational state, appeal to representation in order to explain

¹⁵ This is not incompatible with content realism *per se*, but it is incompatible with mainstream robust realism about representation, which sees representations as stable and repeatable cognitive structures. Non-robust realist views are therefore in the cards, but outside the scope of this paper. At any rate, the pragmatist's insistence that such variety of content ascriptions stems at least partly from heuristic and pragmatic considerations is incompatible with realism *tout court*—provided that the points in Sect. 3.4.1 be kept in mind.

¹⁶ Such a hypothesis is in line with neural reuse theories. For a review of the theories and the evidence for them, see Anderson (2010).

¹⁷ For a review of the most influential theories, and their main problems, see Ryder (2009).

specific behaviour loses much of its value. Moreover, indeterminacy of content imperils the possibility of giving an account of misrepresentation, generally taken to be an essential requirement for a satisfactory theory of content (Cummins 1996). Given indeterminacy, some content can in most cases be found such that a representation correctly represents it.

The second step of the argument from despair has it that content pragmatism can succeed where robust realist theories fail.

Content pragmatism has indeed the tools to avoid the indeterminacy problems that plague robust realist theories. First, indeterminacy of content is a problem only if there is a commitment to there being cognitive states that essentially bear content. The content pragmatist can deny that representational content is a real property of cognitive states, and reject the view that cognitive states are essentially individuated by their contents. Second, representational content, being part of the explanatory gloss, is ascribed with pragmatic considerations in view. Depending on explanatory aims and interests, different specific contents may be ascribed to the same cognitive state. Determinacy is achieved if and to the extent that it is explanatorily useful relative to the relevant pragmatic considerations. Misrepresentation is similarly handled. It is only given a certain pragmatic background in which a choice has been made on what content to ascribe to a certain computational state that questions about representational correctness or misrepresentation can be evaluated.

In sum, the argument from despair gives us good reason to take content pragmatism as a serious candidate; the failures of current robust realist accounts, which content pragmatism does not share, puts the latter, if not in advantage in the overall debate, at least on an equal footing. The argument from despair cannot be simply dismissed.

3.4.4 *Glossing Reality*

The final argument against content pragmatism that I will examine is, I believe, also the most threatening. It tackles head-on a crucial difficulty that content pragmatism has to face, namely providing grounds for seeing representational content as explanatory, whilst withholding ontological commitment to the notion. That is, the content pragmatist must answer the pressing question: if talk of representational content cannot be eliminated from cognitive science, even after future progress, why then reject the idea that content ascription is capturing something real about cognitive systems?

Critics of content pragmatism argue that it cannot justify the fracture between explanation and ontology that lies at

the heart of the view.¹⁸ Importantly, this fracture is not a general one; it does not apply to every, or even most, scientific posits—a position characteristic of full-blown scientific pragmatism. By content pragmatism's lights, the fracture comes in for representational content, but not for other scientific entities, such as computational vehicles and processes. The guiding idea is, indeed, to cling to realism about the latter, while shifting to pragmatism about the former. Content pragmatists must provide principled reasons to believe that a notion that is explanatorily ineliminable from cognitive science—representational content—should not be seen as ontologically committing; whilst other putative explanatorily ineliminable notions—such as computational mechanisms—should be so seen. Unless such principled reasons are provided, the critic presses on, content pragmatism is unjustified. Why set the pragmatism/realism border there, rather than somewhere else, or not at all?

I think there are good, albeit defeasible, motivations for placing the border right there. Cognitive science aims at furnishing a fully naturalistic account of what cognition is and how it works. Notions whose naturalistic credentials are dubious should not compose the theory proper, on pain of endangering the scientific status of the field. Representation and content are clearly problematic in this regard. It is indeed both due to their apparent resistance to naturalisation, and to their explanatory role in cognitive science, that philosophers have been so keen on trying and giving robust naturalistic treatments of these notions.

Forty years or so of focused efforts, however, have yielded little success, as the 'argument from despair' suggests. It is thus an open question whether representation and content will ever be naturalised, at least in the way robust realism would want. This fundamental doubt is enough, I submit, to justify scepticism toward a realist take on those notions. It is enough, crucially, to justify quarantining representational content in a pragmatically-motivated, non-naturalistic, explanatory gloss. Perhaps future work will give us the much sought naturalisation of representational content. If this should be so, content pragmatism will arguably lose most of its motivation. However, until that happy day arrives, if it does, content pragmatism is justified in its project of pursuing a different path—an alternative way of seeing the role played by representational content in cognitive science, i.e. not as part of the naturalistic theory proper, but as part of a supplementary non-naturalistic explanatory edifice, built by us, for us, in light of our interests, capacities, and aims.

This is just the first step of the rejoinder, however. I have so far showed only that, given our current knowledge, representational content possibly lies outside the border of what

¹⁸ I thank an anonymous reviewer to this journal for framing the criticism at hand in this clear and pointed way.

we should ontologically commit to. Similar considerations to the above suggest that, as content pragmatism would want, computational mechanisms are to be found on the inner side of that border. For concrete computation, i.e. computation in physical systems, seems a much more promising candidate for naturalisation than representational content. Despite some early scepticism about the prospects of naturalising computation (e.g. Putnam 1988; Searle 1992), recent proposals are much more robust, in particular the mechanistic view (Piccinini 2015; Milkowski 2013; Fresco 2014). If concrete computation is a natural phenomenon, there is no impediment to its appearing in scientific theories. The fate of content pragmatism thus hinges on whether concrete computation, in contrast to representational content, will be satisfyingly naturalised.

In sum, there are plausible, though certainly defeasible reasons to believe that computational states and processes are acceptable components of a naturalistic story about cognition. And there are equally plausible, and equally defeasible, reasons to believe that representational contents are not. This is justification enough to hold, in the current state of play, that the localised fracture between explanation and ontology may fall where content pragmatism claims it does; though this may prove to be wrong given future developments, it is not an incoherent or ungrounded position.

Even conceding that there are cogent reasons to uphold the localised fracture between explanation and ontology that content pragmatism advocates, the critic may still not be satisfied, and rightly so. For now content pragmatism seems to flirt dangerously with eliminativism. The burning question thus becomes: if representational content is a bad candidate for figuring in a naturalistic theory of cognition, why not get rid of the notion entirely, as the eliminativist recommends?

To assuage that worry, content pragmatism reminds us that representational content can never be eliminated from the explanatory gloss because of the nature of the *explananda* in cognitive science. Since those *explananda* are characterised as cognitive abilities having to do with robust successful interaction between organism and environment, the only way to make the computational *explanantia* cogent as explanations of those abilities is to see their components and processes as representations of the body and environment. A commitment to robust realism about content does not follow, but neither does an elimination of the notion from explanations in cognitive science.

Content becomes explanatorily ineliminable once we see physical systems as cognitive systems, bringing thus to bear our explanatory interests in making sense of their behaviour, understood in its turn as performance of cognitive abilities. It is only when we take a specific perspective—a ‘stance’—toward physical systems that representational content becomes something we cannot do without in making perspicuous how the computational goings-on explain

behaviours characterised as cognitive. Taking such a stance brings with it a host of pragmatic factors that inform content ascription in the pragmatist picture. Therefore, the fact that, as the content pragmatist claims, the explanatory gloss in representational terms will never be eliminated from cognitive science (though it may undergo changes as the science progresses), does not entail that content corresponds to some property that cognitive systems and their internal states actually possess.

Thus the collapse into content realism is avoided. At the same time, collapse into eliminativism is also averted: the *explananda* of cognitive science require that content be ascribed to (some) of their *explanantia* in order to make intelligible to us how the latter enable the former. We get what the critic feared was not to be had: explanatory ineliminability of content, without ontological commitment—i.e. the localised fracture between explanation and ontology that content pragmatism champions. This retort has an obvious shortcoming: it makes representational content parasitic on the intentional contents involved in the relevant perspective-taking, an account of which the content pragmatist does not offer. However, as we have seen above, there are possible non-robustly realist accounts of intentional content that are compatible with the spirit of content pragmatism about representational content.

4 Concluding Remarks

Much can be said about the conceptual and empirical advantages and shortcomings of content pragmatism vis-à-vis its realist, eliminativist, and primitivist competitors. Such treatment is though beyond the scope of the paper. My quite limited aim was not to show that content pragmatism is the right, or even the best theory of representation currently on offer. My rather more modest objective has been to show that it is a coherent and plausible position which deserves careful consideration in the debate. I have argued that extant objections to content pragmatism are less convincing than they might seem at first glance. Content pragmatism is not at odds with scientific practice, and it is not a disguised version of the traditional options normally taken to exhaust the conceptual landscape. The pragmatist can therefore maintain their allegiance to the ineliminability of the notion of representation and content from cognitive science, while at the same time avoiding the view’s collapse into primitivism, eliminativism, or realism.

Acknowledgements Author indebted to Nir Fresco, Guido Löhr, an anonymous referee, and the audience at the 2016 NCH Mind and Brain conference for valuable feedback on earlier versions of this material.

Compliance with Ethical Standards

Conflict of interest The author declared that he has no conflict of interest to declare.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Anderson ML (2010) Neural reuse: a fundamental organizational principle of the brain. *Behav Brain Sci* 33:245–313
- Bechtel W (2016) Investigating neural representations: the tale of place cells. *Synthese* 193:1287–1321
- Burge T (2010) *Origins of objectivity*. Oxford University Press, Oxford
- Churchland PM (1981) Eliminative materialism and the propositional attitudes. *J Philos* 78:67–90
- Cummins RC (1989) *Meaning and mental representation*. MIT Press, Cambridge
- Cummins RC (1996) *Representations, targets, and attitudes*. MIT Press, Cambridge
- Dennett DC (1981) True believers: the intentional strategy and why it works. *Scientific explanation: papers based on Herbert Spencer lectures given in the University of Oxford*. Clarendon, New York
- Dennett DC (1987) *The intentional stance*. The MIT Press, Cambridge
- Dretske F (1981) *Knowledge and the flow of information*. Basil Blackwell, Oxford
- Dretske F (1986) *Misrepresentation. Belief: form, content, and function*. Oxford University Press, Oxford
- Dretske F (1988) *Explaining behavior*. MIT Press, Cambridge, MA
- Egan F (1999) In defence of narrow mindedness. *Mind Lang* 14(2):177–194
- Egan F (2009) Is there a role for representational content in scientific psychology? In: Murphy D, Bishop MA (eds) *Stich and his critics*. Wiley, Malden
- Egan F (2010) Computational models: a modest role for content. *Stud Hist Philos Sci Part A* 41(3):253–259
- Egan F (2014a) Explaining representation: a reply to Matthen. *Philos Stud* 170:137–142
- Egan F (2014b) How to think about mental content. *Philos Stud* 170:115–135
- Fodor JA (1975) *The language of thought*. Harvard University Press, Cambridge, MA
- Fodor JA (1984) *Semantics, Wisconsin style*. *Synthese* 59(3):231–250
- Fodor JA (1987) *Psychosemantics*. MIT Press, Cambridge, MA
- Fodor JA (1990) *A theory of content and other essays*. MIT Press, Cambridge
- Fresco N (2014) *Physical computation and cognitive science*. Springer, Heidelberg
- Gates G (1996) The price of information. *Synthese* 107(3):325–347
- Haugeland J (1990) The intentionality all-stars. *Philos Perspect* 4:383–427
- Hutto DD, Myin E (2013) *Radicalizing enactivism: basic minds without content*. The MIT Press, Cambridge
- Marr D (1982) *Vision*. Freeman, San Francisco
- Matthews RJ (2011) Measurement- theoretic accounts of propositional attitudes. *Philos Compass* 6(11):828–841
- Milkowski M (2013) *Explaining the computational mind*. The MIT Press, Cambridge, MA
- Millikan RG (1984) *Language, thought, and other biological categories*. MIT Press, Cambridge
- Morgan A (2014) Representations gone mental. *Synthese* 191:213–244
- Moser EI, Kropff E, Moser M-B (2008) Place cells, grid cells, and the brain's spatial representation system. *Annu Rev Neurosci* 31:69–89
- O'Keefe JA, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford University Press, New York
- Piccinini G (2015) *Physical computation: a mechanistic account*. Oxford University Press, Oxford
- Putnam H (1988) *Representation and reality*. MIT Press, Cambridge, MA
- Ramsey WM (2007) *Representation reconsidered*. Cambridge University Press, Cambridge
- Ryder D (2009a) Problems of representation II: naturalising content. In: Garzon F, Symons J (eds) *The Routledge companion to the philosophy of psychology*. Routledge, London
- Searle JR (1992) *The rediscovery of the mind*. MIT Press, Cambridge, MA
- Shadmehr R, Wise SP (2005) *Computational neurobiology of reaching and pointing: a foundation for motor learning*. MIT Press, Cambridge, MA
- Shagrir O (2012) Structural representations and the brain. *Br J Philos Sci* 63(3):519–545
- Sprevak M (2011) Review: William M. Ramsey, *Representation Reconsidered*. *Br J Philos Sci* 63(3):669–675
- Sprevak M (2013) Fictionalism about neural representations. *Monist* 96:539–560
- Stich SP (1983) From folk psychology to cognitive science: the case against belief. The MIT Press, Cambridge, MA
- Sullivan JA (2010) A role for representation in cognitive neurobiology. *Philos Sci* 77(5):875–887
- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55(4):189–208
- Van Gelder T (1995) What might cognition be, if not computation? *J of Philos* 92(7):345–381