Body Representations and Cognitive Ontology: Drawing the boundaries of the body image

Stephen Gadsby
Cognition and Philosophy Lab, Monash University, Australia

Abstract

The distinction between body image and body schema has been incredibly influential in cognitive neuroscience. Recently, researchers have begun to speculate about the relationship between these representations (Gadsby, 2017; 2018; Pitron & de Vignemont, 2017; Pitron et al., 2018). Within this emerging literature, Pitron and colleagues (2018) proposed that the long-term body image and long-term body schema co-construct one another, through a process of reciprocal interaction. In proposing this model, they make two assumptions: that the long-term body image incorporates the spatial characteristics of tools, and that it is distorted in the case of Alice in wonderland syndrome. Here, I challenge these assumptions, with a closer examination of what the term "long-term body image" refers to. In doing so, I draw out some important taxonomic principles for research into body representation.

Keywords

body image; body schema; long-term; anorexia nervosa; Alice in wonderland syndrome; body representation; cognitive ontology

1. Introduction: Body Representation and Cognitive Ontology

The distinction between a body image—underpinning bodily perception—and a body schema—underpinning bodily action—has played an incredibly influential role in cognitive neuroscience (Gadsby & Williams, 2018; Gallagher, 1998; de Vignemont, 2010). Nevertheless, much debate, speculation, and confusion still surround these terms, and the kinds of cognitive entities to which they refer (Ritchie & Carruthers, 2015; de Vignemont, 2018). One recent point of interest has been the functional relationship between the body image and schema. Specifically, a debate has emerged over the relationship between the *long-term* body image and long-term body schema, where the term 'long-term' signifies that the representations under discussion (body image/body schema) hold content related to the spatial dimensions of the body. Briefly recapping this debate: I proposed that these representations both derive content from a shared prior (Gadsby, 2017); Pitron & de Vignemont (2017) suggested they have a more direct connection; I pointed out some issues with their proposed model (Gadsby, 2018); and they made some amendments (Pitron et al., 2018).

In the background of this debate—and indeed any debate over body representation—is a question of *cognitive ontology* (Price & Friston, 2005). In particular, we can draw out two issues related to the cognitive ontology of body representations (Janssen et al., 2017). Firstly, there is the issue of how to carve out the set of entities to which our theories of cognition refer. How many body representations are there? What function do they serve? How should we draw the boundaries between them? Tied up with this project are some thorny empirical issues, such as the question of how we measure and intervene on these entities (Gadsby, forthcoming).

Another issue of cognitive ontology is that of nomenclature: once the number of representations are agreed upon, and their boundaries delineated, what set of terms do we use to refer to them? Though this second issue is important, it is somewhat subservient to the first; we can't decide what representations ought to be named before we decide on how many there are, how to delineate them, and what role they play within our broader cognitive economy. As such, it'll be the first issue of cognitive ontology which I focus on most in this paper, though I will also briefly touch on the issue of nomenclature.

In a recent paper, Pitron and colleagues (2018) propose a new model of the long-term body image and long-term body schema, whereby the representations build one another up, through a process of reciprocal interaction. At the heart of this model are two crucial assumptions. The first, is that, during tool use, tools are not only incorporated into the long-term body schema—as is standardly assumed—but the long-term body image also. This first assumption is justified with reference to evidence that tool use enacts changes in the results of both tactile distance estimation, and body part localisation tasks (p. 353; see: Cardinal et al., 2009; Sposito et al., 2012). The second assumption is that Alice in wonderland syndrome involves distortion of the long-term body image. This assumption stems from the syndrome being characterised as change in *experienced* body size.

Here, I challenge both these assumptions. First, I argue that a distinct representation of bodily dimensions underpins tactile distance estimation and body part localisation tasks—*the body model* (Longo & Haggard, 2010). I survey a number of ways in which the long-term body image and body model might be interrelated, or even co-extensive, ultimately arguing that they are distinct. Given this, incorporating tools is a feature of the body model, not the long-term body image. Finally, I discuss how distortion of the body model, rather than long-term body image, is a more plausible explanation of Alice in wonderland syndrome.

2. To what does the term 'long-term body image' refer?

The long-term body image is standardly thought of as a *representation* of the spatial dimensions of the body. In speaking of representations, we are speaking of *vehicles* with particular *contents*, in virtue of which they represent particular *targets*. The vehicle is the internal state of the information processing system, the target is the property or object of the world which the state is about, and the content is how the target is represented to be. The distinction between content and target is important, as it allows for misrepresentation. While two different maps may represent the geographic layout of Melbourne, their contents might differ in terms of *where* they represent Flinders St station (e.g. as either next to Federation square, or next to Brunswick McDonalds). The map whose content more closely matches the target properties (the actual geographic layout of Melbourne) is most accurate.

Another important feature of representations is the relationship between elementary and complex representations (Fodor, 1975). Consider the term "The cat that lives in Steve's house"—a vehicle which has as its target the cat that lives in Steve's house. We can recognise this as a complex representation made up of more elementary representations—vehicles with their own targets and content (e.g. "cat", "Steve's house"). Similarly, two maps representing two halves of Melbourne may be combined to form a map of Melbourne as a whole. As such, a representation can be made up of more elementary representational components, which themselves can serve as representations independently.

We aren't yet sure how to identify what the relevant vehicles are in terms of body representation, or the *format* of these vehicles (see: Haugeland, 2013). It might be that body representations function in a way similar to linguistic systems, through ordered manipulation of symbolic-type structures; or it might be that they function using a map-like format, whereby structural relationships between internal states are exploited as stand-ins for spatial properties of the body (cf. Williams & Colling, 2018). There might even be differences between the kinds of representational format employed by the different body representation systems. For example, the long-term body image may be akin to a pictorial representation (cf. Pearson & Kosslyn, 2015), while the long-term body schema may be map-like; interesting questions would then follow regarding how these formats interact. Though it's too early to make strong claims regarding the format of body representations, we are, I believe, in a position to debate how many representations

there are, what their contents are, and what functions they serve—crucial issues of cognitive ontology.

2.1 Self-Recognition

The most paradigmatic psychological capacity which the long-term body image can be assumed to underpin is self-recognition—specifically, self-recognition based on bodily dimensions. Though Pitron and colleagues don't specifically mention this as a function of the long-term body image, they do agree that anorexia nervosa patients exhibit long-term body image distortion. It's worth taking a moment then to discuss the evidence for this.

It isn't simply that anorexia nervosa patients proclaim to *feel* overweight that researchers assume they exhibit long-term body image distortion, rather, it's due to the results of a particular kind of self-recognition task: *body size estimate* (BSE) tasks (Smeets, 1997). The depictive variation of these tasks—perhaps the most paradigmatic measure of the long-term body image—involves matching a visual stimulus (e.g. a silhouette) to the size of one's body (Longo, 2015a). The underlying assumptions of such tasks is that comparative judgments about size are achieved through a comparator process: the perceptual content (derived from the stimulus) is compared against the relevant content stored in the long-term body image. Performance in such tasks is thus indicative of the accuracy of the long-term body image: consistent misjudgement indicates the long-term body image is distorted (Smeets, 1997). I thus assume that Pitron and colleagues, following anorexia nervosa and other body representation researchers (see below), would concede that these depictive BSE tasks are underpinned by a representation befitting of the title of the long-term body image.

This very assumption has been used to investigate the content of the long-term body image's representation of hand surfaces (Longo & Haggard 2012; Longo, 2015a). Researchers found that participants were accurate at matching hand shape and size for both sides of the hand (dorsal and palmar), suggesting that the long-term body image represents the shape of both sides of the hand fairly accurately. Further, the small amount of distortion they did find correlated between both sides of the hand. These researchers thus additionally concluded that the (long-term) body image represents the hand as a *coherent 3-D object* (Longo & Haggard, 2012).

As promised, two issues of cognitive ontology are already evident. The first pertains to the entities of cognition: a representation exists, and this representation is what facilitates our capacity for self-recognition. Second, the question of nomenclature: this representation ought to be labelled the long-term body image. I take it that the first of these issues is fairly uncontroversial. Given that we don't have direct consistent sensory access to the spatial characteristics of our bodies, it's a computational requirement of self-recognition that this information be stored somewhere, thus we are justified in inferring that such a representation exists, embedded in some broader system of operations which maintain, manipulate, and employ it for self-recognition purposes. Though the issue of nomenclature may seem more contentious—for example, Longo & Haggard (2010) refer to this representation as "the conscious body image", and I have referred to it before as "the perceptual body image" (Gadsby, 2017b)—for now I'll

simply follow Pitron and colleagues in using the term "long-term body image", returning to the issue of nomenclature at a later point (see: conclusion).

2.2 Distinct body size relevant capacities

Thus far we have identified a psychological capacity (self-recognition based on body size) as underpinned by a certain cognitive entity (a representation of bodily dimensions), and we have settled on calling that entity the long-term body image. This assumption has facilitated certain inferences: that the long-term body image's content, in neurotypical individuals, is (fairly) accurate, cohering to the principles of three-dimensional representation. Yet in their paper, Pitron and colleagues appear to make a further assumption: that the long-term body image underpins other psychological phenomena: tactile distance estimation (TDE) and body part localisation (BPL). Just as Longo & Haggard use behavioural results from depictive self-recognition tasks to infer the content of the long-term body image, Pitron and colleagues (2018) use behavioural results from TDE and BPL tasks to infer the content of the long-term body image (p. 353). Specifically, they use evidence of *changes* in these behavioural results after tool use as evidence that the long-term body image's content changes during tools use.

This inference—from the results of TDE and BPL tasks to the content of the long-term body image—is inappropriate, as the available evidence suggests that TDE and BPL are underpinned by a distinct representation. As in the case of BSE tasks, significant research has gone into measuring the representational content which underpins TDE and BPL tasks in the domain of hand size. For example, in neurotypical participants, Longo & Haggard (2011) found that the distance between two tactile stimuli are perceived as 40% larger when running across, rather than along, the hand dorsum. From this, they infer representational content: that the dorsal surface of the hand is represented as fat and squat (cf. Haggard et al., 2018). Longo & Haggard (2010) also sought to measure the representation of body size underpinning BPL, by asking participants to estimate the location of visually occluded hand landmarks. By mapping the identified landmarks in relation to one another, they constructed maps of how the hand's dorsal surface is represented, concluding that its content was "massively distorted in a highly consistent and stereotyped way across participants" (Longo, 2015b, p. 10-11; cf. Medina & Duckett, 2017).

Judging by the language they use, these researchers appear to take the outcomes of these tasks to be measures of representational content. Yet, given this, the content measured appears to be quite different between BSE, TDE, and BPL tasks. The standard interpretation of these findings is that *different* representations underpin the different tasks—as Longo puts it "where qualitatively distinct patterns of distortion are found, different representations may be involved" (2017, p. 385). Based on this principle, Longo (2017) suggests an accurate representation—the body image—underpins depictive BSE, and a distorted representation—the body model—underpins TDE and BPL.[1] According to this cognitive ontology, the evidence Pitron and colleagues cite suggests that tools are incorporated into the body model, not the long-term body image.

[1] For present purposes, I'll simply follow Longo in assuming that the same representation (the body model) underpins both TDE and BPL, rather than two distinct representations (see: conclusion).

3. What is the long-term body image?

At first glance, the aforementioned evidence seems to point to the following conclusion: there are important differences between the characteristics of, on the one hand BSE, and on the other, BPL and TDE tasks, which likely derive from differences in the representational content which underpins them. How this finding bears on our assumptions of cognitive ontology ultimately depends on what those assumptions are. Specifically, to what does the term 'long-term body image' refer? And, given this, what relationship does it have to the body model? Though I will ultimately argue for one interpretation, two anonymous reviewers have implored me to rigorously survey multiple possible interpretations. This was a good suggestion, as considering these various interpretations will be useful for drawing out some more general principles for research into body representation.

3.1 The Long-term Body Image as Unique Body Representation

The first—and, I argue, most plausible—interpretation is that the term 'long-term body image' is a unique representation, *independent* to the body model (see: fig. 1). Thinking back to the vehicle/content/target distinction, this entails that the long-term body image refers to a specific representational vehicle, with determinate contents, targeting the spatial dimensions of the body. This vehicle is distinct from the vehicle(s) of the body model, which represents the same target—the spatial dimensions of the body. Though the long-term body image and body model *could* have the same content (e.g. both representing their target with complete accuracy), the available evidence suggests that they don't (one is more accurate than the other).

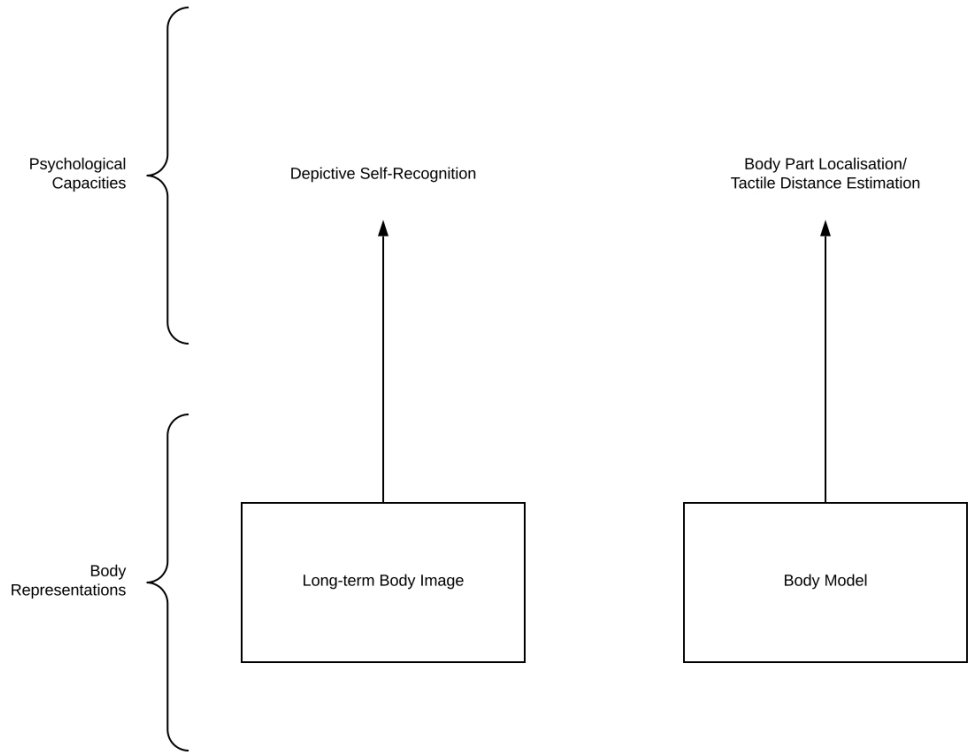

Fig. 1

Though I've claimed the long-term body image is an independent representation, it might also be complex, comprising of a number of more elementary representations. For example, these more elementary representations may target body *parts* (e.g. right hand, left hand), ultimately forming a coherent representation of the whole body. However, what this interpretation doesn't allow is that the long-term body image consists of multiple vehicles with the same target: within the long-term body image there are no two vehicles which represent the spatial dimensions of the left hand (either with the same, or different content). Because of this restriction, we cannot claim that distinct long-term body image sub-representations underpin distinct (body size relevant) psychological capacities. If two distinct cognitive functions rely on the spatial dimensions of the left hand, each cannot use different *parts* of the long-term body image.

If we follow this assumption, and we also assume (following Longo's suggestion) that behavioural results are indicative of underlying representational content, then assuming BSE, TDE and BSP tasks are all underpinned by the long-term body image is unwarranted. Different content (about the same target) underpins these tasks and thus different representations must. And if different representations underpin each task, then the outcome of one task cannot be used to infer the content of the representation underpinning a different task—results from TDE or BSP tasks cannot be used to infer the content of the long-term body image. If this is the case, we have no reason to assume the long-term body image incorporates tools—the evidence only suggests tool incorporation as a feature of the body model.

3.2 A representational hierarchy

It's worth taking a moment to see how this proposal coheres with a hypothesis put forward by Longo (2015b, p. 11-12), who claims that rather than being thought of as "entirely distinct representations emerging from different sensory modalities, the body image arising through vision and distorted implicit representations [i.e. the body model] through somatosensation", these representations may lie at "opposite ends of a continuum of body representations". He describes this hierarchical continuum like so:

> At one end are primary somatosensory maps, representing the body surface as a mosaic of individual receptive fields, each constituting a single "pixel." At the other end is our conscious experience of our body as a volumetric object in the world. In between these extremes may be 2-D maps of individual skin surfaces (such as I have argued may underlie tactile localization and tactile size perception), and 2.5-D representations (such as I have argued underlies position sense). (ibid. p. 12)

Perhaps then, the long-term body image might be said to exhibit some internal hierarchical structure, whereby lower levels represent more basic spatial features, and higher levels more complex ones. Each layer could then be thought of as an elementary representation, all combining to form something more complex. If this is the case, then

Longo's proposal is a way in which the long-term body image and body model might be elementary sub-components of a single representation, contradicting the above proposal.

I think this position is untenable. Firstly, to clarify, though Longo contrasts his view with one where distinct representation emerge from distinct sensory modalities, that is not a necessary feature of the proposed model (section 3.1.). It seems most likely that the long-term body image and body model are both sensitive to input from multiple modalities—treating them as unique does not entail treating them as tied to a single modality. Further, they may bear some special (perhaps hierarchical) relationship to one another, in terms of the exchange of content—i.e. the content of one might be used to 'build up' the content of another—just as is with Marr's (1982) visual representations, and Pitron and de Vignemont's (2017) proposal regarding the relationship between long-term body image and long-term body schema.

That the long-term body image and body model stand in a hierarchical relationship to one another does not entail, however, that they be thought of as elementary sub-components of a single representation. We ought to resist that interpretation for the following reason: there is good reason to assume that our cognitive systems treat the long-term body image (as defined above) as a unique representation, employing it in a comparator process to facilitate depictive self-recognition. Given that the system treats it as a unique representation, our cognitive ontology ought to reflect that—assuming these representations as elementary sub-components leads to overlooking the crucial differences (in terms of content and functional characteristics) between them.[2] I expound on this point further below.

3.3 The Long-term Body Image as Representational Category

One anonymous reviewer suggests a contradictory interpretation to the one discussed. They claim that Pitron and colleagues' assume the long-term body image is simply a *category* of representation. So the long-term body image presumably refers to a number of distinct representations, all belonging to the same category. Unlike in the previous interpretation, this category could consist of multiple representations of the same target (e.g. two distinct vehicles, which each represent the dimensions of the left hand). Given that these representations independently represent the spatial properties of the body, each could underpin distinct psychological capacities that rely on these properties.

Under this interpretation (see fig. 2), the representations so far discussed still maintain their relationship to the relevant behavioural capacities; however, they are also to be considered as part of a broader representational *category*: the long-term body image. It's then this category of representations which plays the relevant functional role in Pitron and colleagues' model. I suspect this (or some form of) category interpretation might be

[2] This also holds for our understanding of Marr's representational account of vision. It would be poor cognitive ontology to lump low, intermediate and high-level visual representations as simply primitive sub-components of a single representation: these representations bear important functional dissimilarities e.g. intermediate (but not low or high) level representations are suggested to underpin conscious perception, (see: Prinz, 2002). So, although they stand in a hierarchical relationship with one another (in terms of information processing), they are nevertheless distinct.

common amongst body representation researchers, so I will take a moment to spell out why it is problematic.

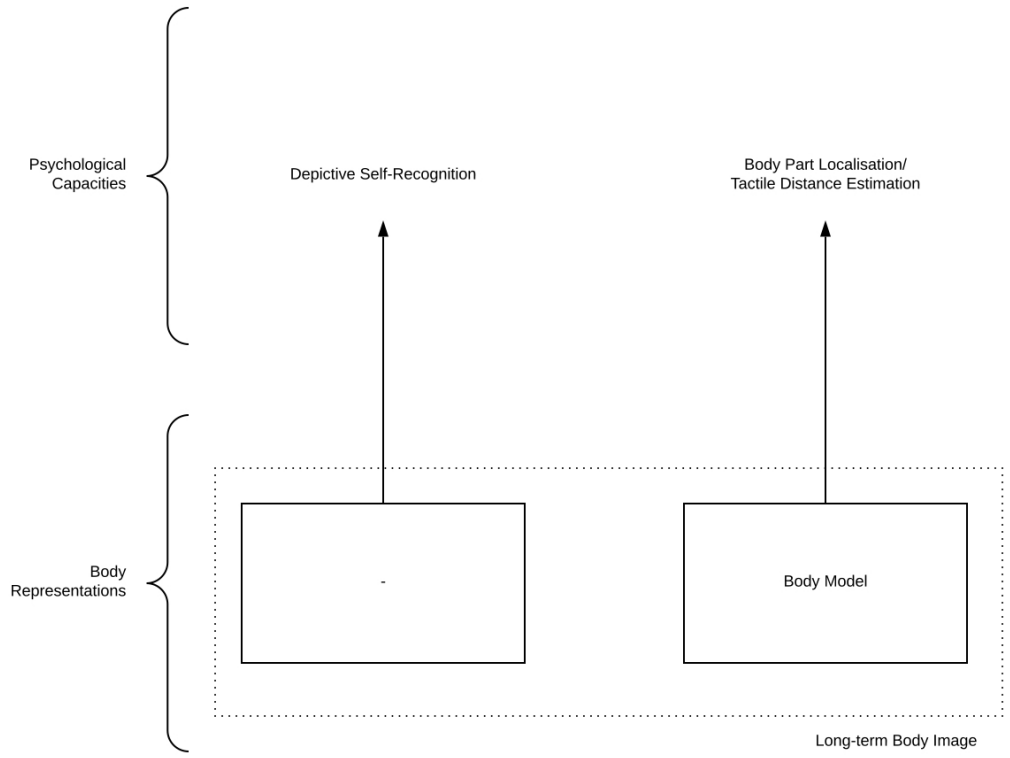

Fig. 2

3.3.1. Descriptive Representation and the Body Image

To be fair, it isn't *necessarily* problematic to group together representations and treat them as a single entity; this kind of taxonomic operation is undoubtedly an important feature of cognitive ontology. However, when taxonomic operations are conducted inappropriately, it leads to inferential problems (Clutton, Gadsby, Klein, 2017). What we require from groupings are projectable inferences—we want to infer properties that hold for *all the things* considered as part of the 'long-term body image'. Lumping representations which are too dissimilar into the same category robs us of that power, and leads researchers into making false inferences i.e. that what holds for one representation (e.g. incorporating tools) necessarily holds for another.

With that in mind, let's explore the category approach to the long-term body image. The most obvious question for this approach is: on what grounds do we draw the boundaries of this long-term body image category? Which representations qualify as belonging to the long-term body image? In line with comments made in the target article (p. 354), and de Vignemont's most recent work (2018), the anonymous reviewer claimed that the body image category is defined by the representations within it having descriptive content (in contrast to body schema representations which have both descriptive and prescriptive content).

One significant problem with this approach is that referring to descriptive content does not help us draw the boundaries in any useful way. Beliefs about body size have descriptive content, yet it would be unhelpful to group these together with the representations under discussion here, beliefs about bodies are *propositional attitudes*, which exhibit significantly different characteristics to other forms of representation (Fodor, 1983; see: Bermúdez, 2005, Chapter 9; Williams, 2018). Similarly so for the formation of purely visual representations of the body. Just as we form visual representations of all the objects we look at, it's possible to generate a visual representation of our body by looking at it. Though this visual representation contains descriptive content, we would not want to include it in the category of long-term body image: such representations are simply a feature of a more general capacity of the visual system, unrelated to any cognitive system whose function it is to track the body.

So drawing the boundaries of this category according to representations of body size which have descriptive content is inappropriate—we end up with a category which is too heterogenous to grant inferential power. From this, we can draw a more general point about cognitive ontology and body representations: if we are to lump together multiple representations into a single category, then the boundaries of that category must be defined in such a way as to prohibit excessive heterogeneity. As such, defining any representational category based on a broad-stroke principle like 'representations of bodily dimensions which are conscious/descriptive/perceptual', is inadequate.[3]

3.3.2. Grouping together the body image and body model

Even if we were to narrow down the long-term body image category such that it *only* included the two representations under discussion (fig. 2), this would be still be a poor taxonomic move. If different sub-representations hold entirely distinct (and inconsistent) content and also underpin entirely distinct psychological capacities, then what is to be gained by grouping them together as all belonging to the same category? Such a category is of very little use in terms of the kinds of psychological theorising myself and Pitron and colleagues are engaged in. Consider: if different subsections of the category can hold inconsistent content, operate entirely independently, be employed for significantly different processing chains, culminating in significantly different psychological capacities, then presumably these sub-sections can functionally interact in entirely different ways with other representations (or representational categories), such as the long-term body schema. This is problematic for models which claims that the long-term body image and long-term body schema functionally interact with one another.

Pitron and colleagues' most recent model claims that the long-term body image and long-term body schema *reshape* one another, through a process of reciprocal influence. Yet if

[3] Unfortunately, the lumping of wildly disparate phenomena under the label 'body image' is common place in body representation research. For example, in the target article Pitron and colleagues claim a number of diverse phenomena—the rubber hand illusion, verbal suggestions during hypnosis, tendency for self-objectification—as all underpinned by "the body image". Yet in fact these phenomena are underpinned by an incredibly diverse collection of cognitive functions (multi-sensory processing, speech processing, propositional attitude formation). As de Vignemont herself recently put it, "How could all these aspects be part of one single category? Because of this apparent heterogeneity, there is a risk that the concept of body image becomes empty of meaning and of explanatory power" (2018, p. 141).

we assume that the long-term body image is a category of distinct representations, their claims become difficult to interpret. For example, which parts of the long-term body image category are involved in the operations they claim? What does it mean, for example, to say that the long-term body image acts as a prior for the long-term body schema (Pitron et al., 2018, p. 355)? How exactly would such a process play out? Given that sub-representations hold distinct and inconsistent content, there must be some internal mechanics that determine which particular sub-representation's content contributes to re-forming the body schema. Are both sub-representations weighted against one another according to some principle? Does one of the sub-representations take precedent when acting as prior for the long-term body schema? The representational category interpretation leaves too many questions unanswered—Pitron and colleague's model becomes nothing more than a vague suggestion.

One final issue for any category approach which groups the body model as part of the body image is that it seems plausible that the body model might simply refer to the same representation as 'the long-term body schema'. First, consider the functional role of the body model: a representation of body size and shape used to calculate body posture and location. This representation is then combined with afferent signals to computer current body position and location (Longo & Haggard, 2010). Represented body posture and location is a fundamental feature of the standard computational framework for motor control, as it is used to compute body appropriate motor commands (via inverse models) (Wolpert et al., 2000; see: de Vignemont. 2009). Indeed, given this role, this representation might be referred to simply as the "short-term body schema"—a representation of the short-term properties of the body (location and posture), employed in motor control (Gadsby & Williams, 2018) (see fig. 3).

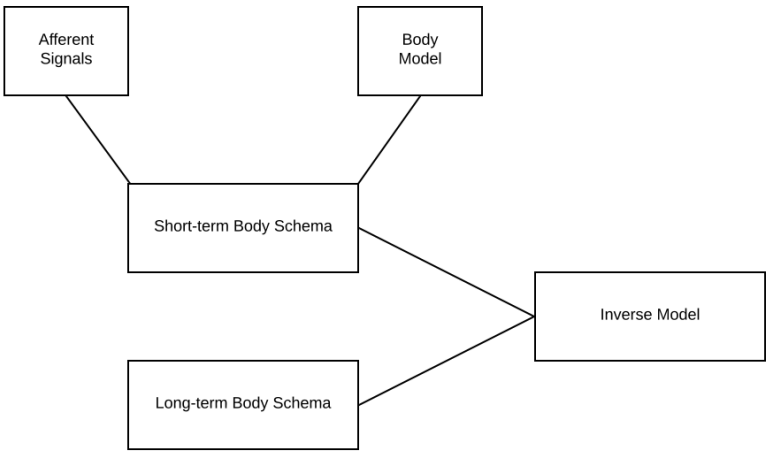

Fig. 3

I am, of course, not the first to suggest that the body model plays this role in driving action (Martel et al., 2016). In their original proposal, Longo & Haggard clearly recognise the body model's role in motor control, discussing different hypotheses for how its distorted content might be compatible with skilled manual action (Longo & Haggard, 2010, p. 11730). Similarly, Peviani and Bottini (2018) compared results from a BPL and movement task to assess whether the same (distorted) representation of hand dimensions was employed for both. Participants' movements appeared to be affected by the distorted

content of the body model, with Peviani and Bottini concluding that the body model does indeed "play a pivotal role in motor planning" (p. 890).

A crucial question now is: is it plausible to think there are two distinct representations of bodily dimensions at play in motor control—the long-term body schema and body model (as in fig. 3)—or is it more likely that our motor control system simply relies on a single representation of these properties. If the latter is the case, then "the body model" simply refers to the same representational entity as "the long-term body schema". This would mean that (contrary to Pitron and colleagues' claim) differences in BPL and TDE results after tool use is simply more evidence for the standard hypothesis: that tools are incorporated into the body schema. Further, it would mean that any cognitive ontology where the body model is included as part of the long-term body image is untenable, as it dissolves the central distinction between body image and body schema.

3.4 The Long-Term Body Image as Additional Representation

Another interpretation is one where "the long-term body image" is a distinct representation to those discussed thus far. Under this view, Pitron and colleagues can be seen to agree with Longo and Haggard's findings (or at least my interpretation of them)—that distinct representations (with distinct content) underpin the discussed abilities—they might simply claim that each of these tasks are underpinned by an *additional* representation (what *they* think of as the long-term body image) (see fig. 4).

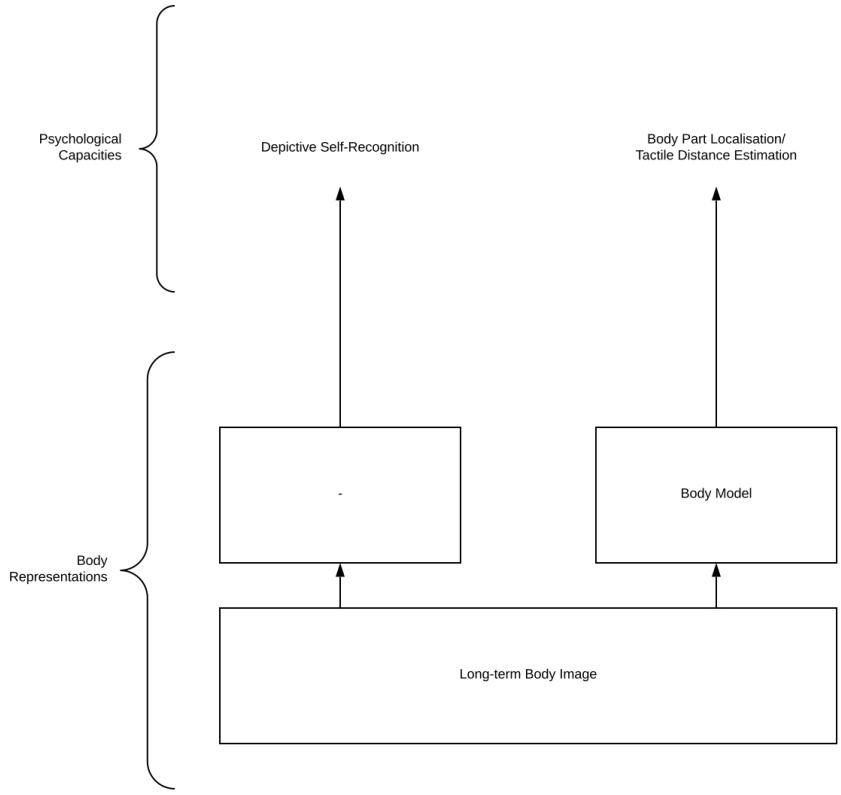

Fig. 4

While this interpretation is plausible enough, it isn't particularly parsimonious. It's not clear if there is any reason why we ought to posit this additional representation: it certainly doesn't appear to offer any explanatory benefit over and above the more parsimonious ontology. Further, this interpretation undermines the kinds of inferences at the heart of the issue: that one can assume a change in long-term body image content from the results of BSE, TDE, and BPL tasks. Such an inference can only be made if such tasks are *exclusively* underpinned by the long-term body image. If, for example, TDE was underpinned by two representations of bodily spatial properties—the body model and long-term body image—then a change in TDE would underdetermine a change in long-term body image content, as changes in behavioural results might simply arise from changes in the body model alone. So not only does this interpretation lack parsimony, it undermines the justification for Pitron and colleagues claim that tools are incorporated into the long-term body image.

3.5 Behavioural Changes as Content Irrelevant

There is one final interpretational option available—again proposed as most intuitive and charitable by an anonymous reviewer—which might be of some help in rescuing Pitron and colleagues' model. Under this view, there is a single representation: the long-term body image, however, this single representation is employed in significantly different processing chains which underpin the different psychological capacities. Content-irrelevant differences (e.g. in distinct information processing operations, cognitive resources etc.) then explain the evident differences in BSE, BPL, and TDE data. At the heart of this interpretation is a challenge to the standard inferential route—that we can infer representational content from the analysis of behavioural task results. I find this proposal the most interesting of all—also the most disastrous.

Firstly, we are faced with the issue of underdetermination again. Given that, ex hypothesi, behavioural patterns that *appear* reflective of content are simply now assumed to emerge from distinct cognitive factors, these same patterns are no longer a reliable measure of representational content. If they are no longer a reliable measure of representational content, however, Pitron and colleagues cannot use the results of BPL and TDE to claim that tools are incorporated into body representations.

Some variation of this interpretation could be maintained without falling victim to the problem of underdetermination, if we were to adopt an interpretation whereby the term "long-term body image" refers not only to a representation but also a broader system of information processing operations. We might then interpret "incorporated into the long-term body image" as meaning "either the content *or* the operations of this system have changed in ways which are sensitive to the spatial properties of the tool". As long as we drew the boundaries of such a system appropriately, this would allow us to again infer changes in 'the long-term body image' based on the observation of behavioural patterns.

Such an interpretation would only be useful for Pitron and colleagues if the "long-term body image" in their model referred not only to a representation but a representational system, involving a much broader set of information processing operations. However, this interpretation leads us into a different one of the previously discussed issues—the

claims of their model become difficult to understand. If the long-term body image refers to a representation plus a broad set of information processing operations (which can shift in ways that fool researchers into thinking changes in content have occurred), then these operations and the effect that they have on the long-term body schema must be explained.

The assumption that we can infer characteristics of representational content from the analysis of behavioural output is fairly essential to theorising in the realm of body representation. Nevertheless, it is certainly worth keeping in mind that many of the behavioural patterns we observe may arise through non-content relevant factors. Indeed, we ought to look to the contexts in which behavioural tasks are conducted and what non-content specific effects these contexts might prompt—as researchers are currently doing (Haggard et al., 2017; Medina & Duckett, 2017; Tamè, et al., 2017). Explaining differences in behavioural output as emerging from differences in non-content specific factors helps to avoid the unnecessary proliferation of body representations (de Vignemont, 2007).

4. **Alice in Wonderland Syndrome**

I have argued for a distinction between a long-term body image which underpins self-recognition and a body model which underpins body part localisation and tactile distance estimation. Given this ontology of body representation, we might ask which representation seems most plausible as distorted in Alice in Wonderland syndrome (AIW). Given that AIW is characterised by reports of a change in the *experience* of body size, it does not seem necessary to assume that this is associated with a change in long-term body image content—self-recognition is distinct from bodily experience.[4] Rather, in light of the preceding discussion of the different posited body representations and the kinds of experiences and behavioural tasks they underpin, a more plausible hypothesis is that AIW involves distortion of the body model.

The body model underpins our experience of body location. As such, if it were to represent one's body inaccurately, this would affect the experienced position of one's body. Turning to descriptions from those who have suffered from the syndrome, if the body model represented one as shorter than reality, this might plausibly cause one to claim "I have a very peculiar feeling of being very close to the ground as I walk along" (Pitron & de Vignemont, 2017, p. 118). Alternatively, if the body model represented one as taller than reality, this might cause one to report "A feeling that I was very tall. When walking down the street I would think I would be able to look down on the tops of others' heads" (*ibid.*). So while distortion of the (long-term) body image—when properly defined—doesn't account for AIW, distortion of the body model might.[5]

5. **Conclusion**

[4] Of course, the two are related. A change in the experience of body size could, and does, enact a change in the recognition of body size (see: Piryankova et al., 2014).

[5] Of course, such a hypothesis opens up a whole suite of further issues, such as why the vast majority AIW sufferers do not appear to suffer from any form of deficit in motor ability. Though perhaps, as has been suggested before, the online sensory feedback present during action (usually) makes up for the inaccuracy of the body model's content (Longo & Haggard, 2010, p. 11730; Peviani & Bottini, 2018).

Admittedly, the picture I have painted regarding the mapping between body representations and psychological capacities has been simple—basic one-to-one mappings between body representations and psychological capacities needn't hold. For example, participants are less accurate in BSE variants where they are required to make comparative judgments using metric (i.e. non-body), rather than depictive stimuli. Longo & Haggard (2012) suggested that these tasks are influenced by not only the long-term body image, but some *combination* of representations. This is representative of the principle that the free-handed addition of new body representations to the ontology, based on observed behavioural differences is not always the right move—we needn't assume some additional body representation underpins metric BSE tasks, as the present ontology has sufficient resources to explain the results (cf. de Vignemont, 2007).

Similarly, the principles I've been discussing might push us towards assuming that distinct body representations underpin TDE and BPL tasks—after all, the behavioural characteristics are different, suggesting distinct content is involved (Longo & Marcom, 2016). Nevertheless, that both tasks seem susceptible to influence by tool use, suggests that *if* different representations are involved they both bear a similarity in functional properties—both incorporate the spatial characteristics of tools. Similarly, Longo justifies the single body model assumption with reference to other functional similarities (both TDE and BPL tasks require dimensional content to be integrated with afferent signals) and content similarities—though both TDE and BSE indicate distortion, the distortion is qualitatively similar (Longo, 2017 p. 385). Establishing whether a single or dual body model ontology holds is a crucial issue for future research.

Finally, in regards to nomenclature, I have argued that the representation which underpins depictive self-recognition is the most befitting of the label 'long-term body image'. However, there is some question about this choice of term. For example, as I have defined it, there is no such thing as a 'short-term body image' (O'Shaughnessy, 1980). Recognition of one's current bodily position does not seem distinct from being aware of one's current bodily position which, ex hypothesi, is underpinned by the body model, in combination with afferent signals. As discussed, referring to *this* collection as the short-term body image would undermine the distinction between body image and body schema. Nevertheless, we ought to distinguish the representation used for self-recognition from the broader body image concept, regularly—but perhaps unfortunately—used to refer to propositional and affective attitudes related to the body (Gallagher, 2005). What terminology we adopt to do so is an open question; perhaps, as has been suggested, we ought to jettison the body image term entirely (de Vignemont, 2018).

These issues aside, the proposed split—between long-term body image and body model— is an important starting point in appreciating the differences between the representations that facilitate depictive self-recognition and those that facilitate TDE and BPL. Given this ontology, it is currently unwarranted to assume that the long-term body image incorporates the spatial properties of tools—rather, this seems like a feature of the body model. Furthermore, when the evidence from Alice in Wonderland syndrome is

reconsidered in light of this, the most plausible hypothesis is that it involves distortion of the body model, not the long-term body image.

This also opens up the debate regarding the relationship between long-term body image and long-term body schema. In a previous article (Gadsby, 2017) I claimed that the evidence from Alice in Wonderland syndrome seems to count against a 'shared prior' model of the relationship between long-term body image and long-term body schema. However, now that we have properly defined and delineated the long-term body image, and shown that it is not implicated in Alice in Wonderland syndrome, perhaps some form of shared prior model is worth reconsidering.

Acknowledgments

Thanks to Andy McKilliam, Iwan Williams, and two anonymous reviewers for feedback on earlier drafts. Thanks also to Manja Engel and Kayla Stone for helpful discussion. This research was supported by an Australian Government Research Training Program (RTP) Scholarship.

References

Bermúdez, J. (2005). *Philosophy of Psychology. A contemporary introduction*. Routledge

Cardinali, L., Frassinetti, F., Brozzoli, C., Urquizar, C., Roy, A. C., & Farnè, A. (2009). Tool-use induces morphological updating of the body schema. *Current Biology*, 19(13), 1157.

Clutton, P., Gadsby, S., & Klein, C. (2017). Taxonomising delusions: content or aetiology?. *Cognitive neuropsychiatry*, *22*(6), 508-527.

Fodor, J. (1975). *The Language of Thought*, Cambridge, Massachusetts: Harvard University Press.

Fodor, J. (1983). *The modularity of mind*. Cambridge: The MIT Press.

Gadsby, S. (2017a). Distorted body representations in anorexia nervosa. *Consciousness and Cognition*, 51, 17–33.

Gadsby, S. (2017b). Explaining body size beliefs in anorexia. *Cognitive neuropsychiatry*, *22*(6), 495-507.

Gadsby, S. (2018). How are the spatial characteristics of the body represented? A reply to Pitron & de Vignemont. *Consciousness and Cognition*, 62, 163–168.

Gadsby, S. (*forthcoming*) Manipulating body representations with virtual reality: clinical implications for anorexia nervosa. *Philosophical Psychology*

Gadsby, S., & Williams, D. (2018). Action, affordances, and anorexia: body representation and basic cognition. *Synthese*, *195*(12), 5297-5317.

Gallagher, S. (1986). Body image and body schema: A conceptual clarification. *Journal of Mind and Behavior*, 7, 541–554.

Gallagher, S. (2005). *How the body shapes the mind*. New York: Oxford University Press.

Haggard, P., Cheng, T., Beck, B., & Fardo, F. (2017). Spatial perception and the sense of touch. *The Subject's Matter: Self-consciousness and the Body*, 97-114.

Haugeland, J. (2013). Representational genera. In *Philosophy and connectionist theory* (pp. 75-104). Psychology Press.

Janssen, A., Klein, C., & Slors, M. (2017). What is a cognitive ontology, anyway?. *Philosophical Explorations*, 20:2

Longo, M. R. (2015a). Three-dimensional coherence of the conscious body image. *Quarterly Journal of Experimental Psychology*, *68*(6), 1116-1123.

Longo, M. R. (2015b). Implicit and explicit body representations. *European Psychologist*.

Longo, M. R. (2017). Distorted Body Representations in Healthy Cognition. *Quarterly Journal of Experimental Psychology*, *70*(3), 378–388.

Longo, M. R., & Haggard, P. (2010). An implicit body representation underlying human position sense. *Proceedings of the National Academy of Sciences of the United States of America*,107, 11727–11732.

Longo, M. R., & Haggard, P. (2011). Weber's illusion and body shape: anisotropy of tactile size perception on the hand. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(3), 720.

Longo, M. R., & Haggard, P. (2012). Implicit body representations and the conscious body image. *Acta Psychologica*, 141(2), 164-168.

Longo, M. R., & Morcom, R. (2016). No correlation between distorted body representations underlying tactile distance perception and position sense. *Frontiers in human neuroscience*, *10*, 593.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of Visual Information*, New York, W.H. Freeman and Company. pp. 302–309.

Martel, M., Cardinali, L., Roy, A. C., & Farnè, A. (2016). Tool-use: An open window into body representation and its plasticity. *Cognitive Neuropsychology*, *33*(1-2), 82-101.

Medina, J., & Duckett, C. (2017). Domain-general biases in spatial localization: Evidence against a distorted body model hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(7), 1430.

O'Shaughnessy, B. (1980). The will, a dual aspect theory, volume I. Cambridge: Cambridge University Press.

Pearson, J., & Kosslyn, S. M. (2015). The heterogeneity of mental representation: ending the imagery debate. *Proceedings of the National Academy of Sciences*, *112*(33), 10089-10092.

Peviani, V., & Bottini, G. (2018). The distorted hand metric representation serves both perception and action. *Journal of Cognitive Psychology*, 30(8), 880-893.

Piryankova, I. V., Wong, H. Y., Linkenauger, S. A., Stinson, C., Longo, M. R., Bülthoff, H. H., & Mohler, B. J. (2014). Owning an overweight or underweight body: distinguishing the physical, experienced and virtual body. PloS one, 9(8), e103428.

Pitron, V., & de Vignemont, F. (2017). Beyond differences between the body schema and the body image: Insights from body hallucinations. *Consciousness and Cognition*, 53, 115–121.

Price, C. J., and K. J. Friston. 2005. "Functional Ontologies for Cognition: The Systematic Definition of Structure and Function." Cognitive Neuropsychology 22 (3): 262–275.

Prinz, J. J. (2004). *Furnishing the mind: Concepts and their perceptual basis*. MIT press.

Ritchie, J. B., & Carruthers, P. (2015). The bodily senses. In *The Oxford Handbook of Philosophy of Perception* (pp. 353-370). Oxford University Press.

Smeets, M. A. (1997). The rise and fall of body size estimation research in anorexia nervosa: A review and reconceptualization. *European Eating Disorders Review: The Professional Journal of the Eating Disorders Association*, *5*(2), 75-95.

Sposito, A., Bolognini, N., Vallar, G., & Maravita, A. (2012). Extension of perceived arm length following tool-use: Clues to plasticity of body metrics. Neuropsychologia, 50(9), 2187–2194.

Tamè, L., Bumpus, N., Linkenauger, S. A., & Longo, M. R. (2017). Distorted body representations are robust to differences in experimental instructions. *Attention, Perception, & Psychophysics*, *79*(4), 1204-1216.

de Vignemont, F. (2007), How many representations of the body? *Behavioural and Brain Sciences*, 30(2), 204-5.

de Vignemont, F. (2010). Body schema and body image-Pros and cons. *Neuropsychologia*.

de Vignemont, F. (2018). *Mind the body*. Oxford University Press.

Williams, D. (2018). Predictive coding and thought. *Synthese*, 1-27.

Williams, D., & Colling, L. (2018). From symbols to icons: The return of resemblance in the cognitive neuroscience revolution. *Synthese*, *195*(5), 1941-1967.