



Representing Without Representations

Mark Rowlands

Department of Philosophy
University of Miami, USA

Abstract

There is a problem of representation and an apparatus of representations that was devised to solve this problem. This paper has two purposes. First, it will show why the problem of representation outstrips the apparatus of representations in the sense that the problem survives the demise of the apparatus. Secondly, it will argue that the question of whether cognition does or not involve representations is a poorly defined question, and far too crude to be helpful in understanding the nature of cognitive processes.

Keywords: action; cognition; intentionality; normativity; representation.

1. Introduction

The claim that cognition is to be explained in terms of operations performed on mental representations has stumbled upon hard times in recent years; at least in certain circles. The 'circles' in question comprise a loose and, according to some, unholy alliance of researchers in the fields of robotics, embodied approaches to cognition, enactive approaches to perception, dynamicists, and vehicle externalists, to name but a few. To a considerable extent, the tribulations of the apparatus of mental representations were predictable. The idea of a mental representation has come to mean too many different things to too many different people. As an erstwhile, if distinctly revisionist, defender of mental representations, I strongly suspect that the problem is not with the idea of mental representations as such – if only people could be persuaded to use the term in the right way. But that, unfortunately, is the tricky part. And so, even if the theoretical utility of the idea of mental representations remains intact, its practical utility must be questioned.

This paper, therefore, begins with a strategic decision: Let us refuse to quantify over mental representations. Let us eliminate all talk of mental representations from our vocabulary. Let us cease to use the concept of representation as an explanatory concept in the development of our theories of cognitive processes. There: at a stroke it is done! And this is all absolutely fine with me. But even if we

do this, we are still left with a problem of *representation*. The problem of representation is a problem that the apparatus of representations was, in effect, introduced to solve. But the problem of representation is not exhausted by the apparatus of mental representations. It outstrips this apparatus in crucial ways. Therefore, even if we eschew the apparatus, the problem remains. Of course, without the apparatus, we might no longer call it the problem of representation. And, indeed, before we have progressed more than few pages, I shall be calling it something else. But a rose by any other name smells as sweet. And the problem remains irrespective of its appellation.

This paper has two purposes. First, it will show why the problem of representation outstrips the apparatus of representations in the sense that the problem survives the demise of the apparatus. Secondly, it will argue that the question of whether cognition does or not involve representations is a poorly defined question, and far too crude to be helpful in understanding the nature of cognitive processes. My stalking horse in this paper will be provided by Shaun Gallagher's interesting but, as I shall try to show, ultimately ill-advised attempt to take issue with the position I defended in my book, *Body Language*. I examine Gallagher's critique not only because I think it is ineffective as a critique of the position defended in that book but, more importantly for present purposes, as a means of getting to the primary destination of this paper: the ill-defined character of the concept of representation.

2. Representation and Normative Grip

Cognitive states make a *normative* claim on the world. Such a claim is essential to their possessing content. This point is a familiar one, and is strongly associated with McDowell (1994), through Sellars (1956/1997) all the way back to Kant. The point is usually made in connection with belief, the most obvious example of a cognitive state. If I have a belief with the content that *p*, then the world *should* be *p*. If the world is not *p* then something has gone wrong. Let us call this the *normativity condition*. This condition is closely associated with another: the *misrepresentation condition*. The possibility of believing, in general, entails the possibility that at least some of one's beliefs are wrong. The normativity constraint grounds – i.e. makes true – the misrepresentation constraint. The normativity constraint is, therefore, more basic.

To see this, suppose your beliefs were simply about whatever caused them. Then they would be infallible. Their content simply consists in whatever produces them, so there would be a perfect match between content and extrinsic state-of-affairs. The normativity condition forces us to distinguish between what actually causally produces a belief and what should produce it. Not everything that actually *does* produce my belief that there is a horse in front of me *should* produce this belief: the donkey that is in front of me and is, in fact, causally producing my belief should not do so. It does produce the belief, but it shouldn't. And this is why my

belief is false. So beliefs can misrepresent – be false – only because beliefs make a normative claim on the world. The normativity condition is, therefore, basic.

Satisfying the normativity condition is an absolute desideratum on any satisfactory account of cognitive states. The apparatus of mental representations was introduced precisely to do this. In its original seventeenth century guise, mental representations were divided into impressions and ideas, corresponding roughly to the distinction between perceptual representations and subsequent memory representations (and representations involved in various processes of ratiocination). This apparatus was introduced precisely to register the epistemic gap between the contents of the mind and the contents of the world. That is, it was introduced precisely to acknowledge the possibility that, given the contents of the world, the contents of the mind might not be what they should be. This was, obviously, a theme that had come to prominence with Descartes.

Contemporary representationalism, at least in cognitive science, views representations not as items of which we are aware in the having of experiences, but as neural states, typically individuated by their higher-order physical or functional properties. But the same desire to safeguard the normativity of representation, and the resulting possibility of error, can still clearly be discerned. According to this account, cognitive states are relations to representations. Representations are understood as discrete, identifiable, internal states of a subject: typically brain states individuated by way of their higher-order properties. Crucially, they can be instantiated independently of what is going on in the outside world. Thus, to use an example of Fodor's, consider the HORSE representation. If this representation is instantiated when the world does not, in an appropriate way, contain a horse, or is not instantiated when the world does, in the relevant sense, contain a horse, then something has gone wrong. Both the normativity and the misrepresentation conditions are satisfied as a consequence of the independence of mental representation and worldly fact. It is this independence that makes possible the distinction between the worldly fact or item that *should* produce the mental representation and what actually *does* produce it.

3. The Persistence of the Problem of Representation

The problem of representation is, ultimately, a problem of explaining our normative *grip* on the world. That is, it is the problem of explaining how we are able to interact with the world in a way that both reflects and respects the normativity condition. The apparatus of mental representations – items sufficiently independent of the world for us to draw a distinction between what in the world does cause them and what in the world should cause them – is one way of explaining our normative grip on the world. But eschewing this apparatus will not make this requirement go away. Consider, for example, the following passage from Dreyfus (quoted from Gallagher 2008). Here, Dreyfus is arguing that skilful intentional action does not require representations:

A phenomenology of skill acquisition confirms that, as one acquires expertise, the acquired know-how is experienced as fine and finer discriminations of situations paired with the appropriate response to each. Maximal grip [Merleau-Ponty] names the body's tendency to refine its responses so as to bring the current situation closer to an optimal gestalt. Thus, successful learning and action do not require propositional mental representations. They do not require semantically interpretable brain representations either (2002: 367).

It is clear that, in this passage, Dreyfus *presupposes* that we have a normative grip on the world. But he does not (in this passage) say how this is possible. Maximal grip, we are told, names the body's tendency to refine its responses so as to bring the current situation closer to an optimal gestalt. But just how, one might ask, does it do this? This is precisely the problem of normative grip. But the problem of representation essentially is the problem of explaining our normative grip on the world: or, at least, once you've explained normative grip, you have done most of the hard work required for explaining representation. Therefore, one cannot simply *presuppose* that we have a normative grip on the world and then claim to have *obviated* the need for representation. (This, it goes without saying, is not intended as a point against Dreyfus, but against the more general assumption that normative grip on the world is not something that requires explanation).

Our normative grip on the world is, I have argued, precisely what the traditional apparatus of representations was supposed to explain. The type of explanation this apparatus proffers is one that, as we might say, goes from the inside out. The true locus of normativity is to be found in internal states of a subject. These then go on and shape that subject's behaviour, and so account for its ability to interact with the world in ways that respects and reflect the normativity condition. The guiding idea behind my book, *Body Language* was that the way to reject the traditional apparatus of representations, while at the same time respecting the normativity condition, is to reject this direction of explanation. The true locus of normativity does not, or does not necessarily, reside on the inside, in the form of inner representations. Rather, it is also to be found on the outside. Our behaviour is infused with a form of normativity that is *sui generis* and does not derive from the inner states of a subject.

4. *Body Language*: The Thesis

The central thrust of *Body Language* was, in effect, to argue that the possibility of satisfying the normativity condition (and other assorted constraints) can be divorced from the traditional apparatus of mental representations. What is crucial, I argued, is explaining the normative grip that an agent has on the world. Whether we do so by way of an appeal to mental representations is an optional extra – one that may, ultimately, turn out to be misguided. To explain the normative grip an agent has on the world is to explain how it behaves in ways that reflect and respect the normativity condition. Nothing I said in *Body Language* precluded the possibili-

ty of trying to satisfy this condition by appeal to the traditional apparatus of representations. That issue was bracketed. However, I did argue that our normative grip does not consist simply in the production within us of mental representations. It consists also in certain things we do in and to the world around us. These *deeds*, I argued, have a normative – and indeed representational – status that is independent of that possessed by any internal states. Deeds are normative, and they are so independently of mental representations. Their normativity is *sui generis*.

The argument can be divided into two parts. First, there is the identification of a sub-category of things we do that I called *deeds* or *pre-intentional* acts. Second, there is the argument that these deeds satisfy most typical criteria of representation. Here is an example of a deed (or rather several) that I employed quite extensively in *Body Language*:

- You are trying to catch a ball in a high velocity sport such as cricket. You have less than half a second before the ball reaches you. However, to compound difficulties, the ball is moving towards you at a tricky height: lower chest height. To complete the catch you have an awkward decision to make – whether to point your fingers up or down.

Pointing your fingers up or down is, in this case, an example of a deed. So too are:

- All the additional on-line, feedback modulated, adjustments of the fingers you will have to make in order to successfully receive the ball.

In general, deeds are *hierarchically structured*: deeds have other deeds as components, and the order of performance is a function of the task. Deeds, I argued, have three defining features:

1. They are things we *do* rather than things that *happen to us*. In catching the ball it is not as if we discover our hands and fingers moving of their own accord – that would be a very alien and unnerving experience.
2. They fall short of intentional *action* in the strict sense. Like actions, intentional states may play a role in explaining the *status* of a deed as something we do. I move my fingers because, ultimately, I have a general antecedent intention to catch the ball. However, unlike actions, general antecedent intentions are not sufficient to *individuate* deeds. An entire array of feedback-modulated adjustments may go in to satisfying one general antecedent intention.
3. They are distinct from sub-intentional acts in O'Shaughnessy's (1980) sense. They are not at all like random tongue-wagglings or toe-tappings where, as O'Shaughnessy puts it, reason plays neither a positive nor negative role in their genesis. Deeds are things we do precisely because we have general antecedent intentions we wish to satisfy.

Deeds are, in effect, pre-intentional acts: items that fall somewhere in between actions in the traditional philosophical sense and sub-intentional acts of the sort made famous by O'Shaughnessy. The fact that they are not individuated by prior intentional states – intentions, volitions, etc – has, I argued, one crucial consequence: it entails that if the deeds performed by a subject have normative status, this cannot have been inherited from prior intentional states of that subject. This makes deeds crucially dissimilar to actions – for which normative status clearly is inherited from prior intentional, hence normative, states.

The second part of the argument was concerned to show that deeds satisfied most – but not all – of the commonly accepted criteria of representation: the conditions that an item must satisfy in order to qualify as capable of representing an item distinct from it. These criteria are:

- (1) *Informational constraint*: An item r qualifies as a representational item only if it carries information about some state-of-affairs s that is extrinsic to it.
- (2) *Teleological constraint*: An item r qualifies as representational only if it has the *proper function* either of tracking the feature or state-of affairs s that *produces* it, or of enabling an organism or other representational *consumer* to achieve some (beneficial) task in virtue of tracking s .
- (3) *Misrepresentation Constraint*: Item r qualifies as representing state-of-affairs s only if it is capable of misrepresenting s .
- (4) *Decouplability Constraint*: Item r qualifies as representing state-of-affairs s only if r is, in an appropriate sense, decouplable from s .
- (5) *Combinatorial constraint*: For any item r to qualify as representational, it must occur not in isolation but only as part of a more general representational framework.

Applied to the simple binary example of a deed described above – pointing one's fingers up or down in order to catch a ball moving with high velocity – the story went something like this. The position of the fingers carries information about the trajectory of the ball; at least it does so to no lesser extent than traditional mental representations carry information about things extrinsic to them. It has the function of tracking the ball's trajectory, or of enabling the catcher to do something in virtue of such tracking. The finger position can misrepresent the ball's trajectory. It can form part of a larger combinatorial system. For example, pointing the fingers up then makes possible the more subtle modulation of fingers necessary to receive the ball. It makes no sense to modulate your fingers in this way unless and until the fingers are pointed up.

This list of criteria was formulated, in part, with a view to inclusiveness. Thus, while some views place the idea of information at the heart of our understanding of the idea of representation, others (for example, the teleosemantic view developed in Millikan 1984) do not. In this paper, I shall focus on what I actually think is

the core of representation: the normativity condition. In *Body Language* I called this the *teleological* condition; a label that I now realize is inadequate. We need to explain our normative grip on the world. But to describe this as a ‘teleological constraint’ is too narrow, and pertains only to one way in which this normative grip might be implemented (a way associated with Millikan and others). The concerns of this paper do not require me to take a stand on whether this is the only way that the normativity condition might be satisfied. Accordingly, I shall take no stand. Henceforth, instead of talking of the teleological condition, I shall talk of the *normativity* condition.

5. Gallagher’s Critique

In an interesting and useful critique, Shaun Gallagher (2008) takes me, in *Body Language*, to be defending a form of *minimal representationalism*. According to this interpretation, I want to (a) cling to the claim that representations are implicated in cognition, and (b) offer a reinterpretation of what representations are along broadly minimalist lines. So, according to this interpretation, my purpose in *Body Language* was to argue that certain sorts of actions – pre-intentional acts or *deeds* – were representations. This would necessitate an expansion in our concept of what sort of thing can be a representation, and accordingly, a drift away from the traditional idea of representation as a discrete, identifiable, and above all internal state of the subject. Gallagher, then, argues that there is no real sense in which such items can be regarded as representations.

In fact, my view was almost the exact contraposition of the one Gallagher describes and attacks. There is, I agree, no real sense in which deeds can be regarded as representations. Therefore, what we need to do is understand that representing the world is an activity that does not require representations. While this may sound paradoxical, it is based on a distinction that is quite simple: the distinction between an item being a *representation* and its being *representational*. One reason – perhaps the principal reason – that Gallagher attributes to me the contraposition of my actual view is because of failure to observe this distinction. Thus, he writes:

Rowlands argues that such ‘deeds’ or pre-intentional acts are representational. Although he hesitates to call them representations per se, to call them representational suggests that they involve representations at some level, and in any case do fit his definition of representation (2008: 356).

However, it is not that I *hesitate* to call pre-intentional acts representations: I *deny* that they are representations. Gallagher claims that to call something representational ‘suggests that they involve representations at some level.’ But that is precisely what I deny. Thus, for example:

The claim to be defended in the rest of this book is that deeds can, in fact, possess representational status. This is not to claim that they are representations as such. As I indicated earlier, the claim that deeds are representations

is not so much false as disingenuous. The content of the concept of a representation is, I think, too closely bound up with the assimilation of representations to the category of the word. And as such, it would be better, for the purposes of stating the principal arguments developed in the following pages, to bracket the apparatus of representations. What are fundamental are not representations but the activity of representing. Deeds, I shall argue, often form a critical part of this activity. And, crucially, the part of the activity of representing constituted by deeds is no less a genuinely representational part than that constituted by internal representations traditionally construed (2006: 114).

To understand the distinction between an item qualifying as a representation and its qualifying as representational, let us first recall the conditions on representation cited earlier: the informational, normativity (formerly ‘teleological’), misrepresentation, decouplability and combinatorial conditions. Collectively, these conditions are an attempt to specify what it is for one item to be *about* another. That is, they are attempts to specify the conditions that must be met if an item is to possess *content*. One may I think, justifiably question whether they do succeed in this task. But that is not pertinent to present purposes. Naturalistic constraints on representation are, precisely, attempts to identify the conditions that must be satisfied if an item is to possess content: that is to count as *representational*. In *Body Language*, I argued that some deeds satisfy these conditions, and do so in the same way and for the same reasons that internal states of an organism might satisfy them.

However, these conditions do not exhaust the generally accepted naturalistic constraints on an item qualifying as a representation. There is, it is generally accepted, a further *explanatory* constraint: if an item is to qualify as a representation then it must play a role in explaining (or producing, or guiding) the behaviour of an organism (see, for example, Cummins 1983). Clearly deeds do not satisfy this further constraint. They do not play a role in explaining (producing, guiding) behaviour because they *are* behaviour. It is, however, easy to identify the difference between this constraint and the other constraints introduced earlier. The informational, normativity, misrepresentation, decouplability and combinatorial constraints pertain to the relation between the proposed representation and the worldly items that it is supposed to be about. The explanatory constraint concerns the relation between representation and behaviour. The first five constraints are in the business of explaining the *aboutness* of an item. The explanatory constraint is in the business of specifying the relation between an item and the behaviour it causes. This is, perhaps, some sort of causal relation. But whatever is true of it, it is not a relation of aboutness.

In *Body Language*, I argued that deeds satisfy the naturalistic constraints pertaining to the aboutness of an item. For this reason they qualify as *representational*. They do not satisfy the further explanatory constraint, and for this reason they do not qualify as *representations*. This, then, is the content of the distinction between what is representational and what is a representation. Satisfying the information-

al, normativity, misrepresentation, decouplability and combinatorial conditions qualifies – at least arguably – an item as representational. Satisfying these and, in addition, the explanatory condition, qualifies an item as a representation. Therefore, deeds are representational but not representations.

In this paper, my focus will be on the middle three conditions: normativity, misrepresentation and decouplability, since it is these that are germane to Gallagher's critique. Of these conditions, normativity is basic. As we have seen, if you can satisfy the normativity condition, the misrepresentation condition comes for free. The misrepresentation condition is entailed by the normativity condition. In *Body Language* I also argued that the decouplability condition was derivative upon the misrepresentation condition in the sense that it is simply a misleading way of stating that condition. That is, I argued that all the decouplability we can reasonably require can be found in misrepresentation, if this is properly understood. If this is correct, then satisfying the normativity condition gives you both of the other conditions for free. The argument for this relation between decouplability and misrepresentation is a novel one, may conceivably be mistaken, and it is here, I think, that the argument needs to be attacked.

Instead, however, Gallagher contents himself with attacking the idea that deeds could be decouplable on intuitive grounds. Implicated in this attack is, I think, a mistaken conception of decouplability. He writes:

But once we do decouple a pre-intentional act from x (the ball, the piano keys, the painting) I suggest that we are no longer talking about action in the same sense. Indeed, it is difficult to see how pre-intentional acts can be decoupled from x (the ball, the piano keys, the painting) or the context without becoming something entirely different from an element of the action at stake' (2008: 357).

How 'difficult to see' it is depends on the situation one is trying to imagine. Waving one fingers around in a context where there is no piano is, of course, a very different action – or, as I would prefer to put it, succession of deeds – from that that of playing the piano. But simply hitting the wrong key during the playing of a difficult passage does not make it a different sort of action at all. Similarly, pointing one's fingers up or down in a context where there is no cricket ball hurtling towards you is a very different type of action from that of performing the same movement in the heat of the game. But simply getting it wrong – pointing your fingers up when you should have pointed them down, and so dropping the ball, for example – does not make it a different sort of action. Gallagher seems to be using the first type of scenario for his model of decouplability, whereas I understand decouplability in the second way.

The key to Gallagher's thinking in this regard is perhaps indicated by the sentences that immediately follow the passage quoted above:

Off-line cognition, imagining, remembering, or even re-enacting an action decoupled from its original context and absent x may (or may not) require re-presentation – but this says nothing at all about representation in action (2008: 357).

I certainly agree with the last clause, of course. If we want to think of decouplability as a general condition on representation, then we cannot assimilate all decouplability to a model supplied by off-line reasoning contexts such as planning, imagining, and re-enacting. If one were the sort of person who believed in perceptual representations, for example, and one wanted them to be decouplable (since one regarded this as a general constraint on representation), then we would require a broader notion of decouplability that acknowledges the relevant differences between perceptual and imaginative representation. If you take perceptual representations out of a context where there is something causally producing them, then you are no longer dealing with perception at all. But this makes it mystifying why Gallagher thinks decoupling the action from the ball or the piano makes it an entirely different sort of action. This would be true only if we were thinking of decouplability in terms appropriate to imagination or re-enactment: a waving of fingers in the absence of a piano. But since this cannot be regarded as a general model of decouplability, what reason is there for holding deeds up to this bizarre and unmotivated standard?

In *Body Language* I argued that if we want decouplability to be a general constraint on representation, and thus a feature of all representations, then all we can reasonably expect from it is a way of typing tokens of mental representations independently of the character of their immediate environment. Thus, the immediate environment provides neither a physical nor logical constraint on the occurrence of the relevant representation. It is in this sense that the decouplability constraint, if understood as a general feature of all representations, collapses into the misrepresentation constraint. But on this way of understanding decouplability, deeds are indeed decouplable from their environment.

6. How to (and How Not to) Think About Representation in Cognition

My response to Gallagher, therefore, turned on the distinction between an item qualifying as *representational* (in virtue of satisfying informational, normativity, misrepresentation, decouplable, and combinatorial constraints) and qualifying as a *representation* (in virtue of it satisfying those constraints plus an additional explanatory constraint). I now suspect that, while this is a step in the right direction, it does not go far enough. The idea of a representation can be analysed into distinguishable components that can be expressed as constraints: normativity, informational, misrepresentation, decouplability, combinatorial and explanatory constraints (and perhaps others not identified here). The impetus for *Body Language* was, in effect, provided by the idea that the explanatory constraint can come apart

from the other ones: that an item might satisfy the other constraints but fail to satisfy the explanatory constraint. This, I argued, was true of deeds.

However, once we acknowledge that these components are distinguishable, there is no reason in principle why this situation might not apply more generally. Of course I have argued for a close connection between the middle three conditions – normativity, misrepresentation and decouplability. The first condition subsumes the remaining two. However, there is no reason, in principle, why an item should not satisfy say, the normativity condition and not the informational (in effect, an implication of teleosemantic approaches), or vice versa (in effect, an unfortunate implication of pure informational approaches). There is no reason why an item should not satisfy the misrepresentation condition, but fail to satisfy the combinatorial condition. Once we identify distinguishable components of the concept of representation, it is always an open possibility that an item may satisfy some components but not others.

There is a quintessential philosophical move that occurs when a certain phenomenon – let us call it ‘X’ – is identified as comprising conceptually distinguishable elements. An item is shown to satisfy some of these elements but not others. Then we are presented with the almost inevitable question: But is it or is it not an example of X? Underlying this question is the philosopher’s desire for generality. We want to know not merely whether it satisfies some or other specific constraints, but whether it falls under the more general category X. But this question, I am suggesting, may have no answer – and I suspect that is the case for the category of representation. Does cognition require representation? Well, some items that seem to be involved in cognitive processes may satisfy some of the conditions of representation, others may satisfy others of those constraints. And that is all we can say of the question of whether cognition requires representation.

This paper began with a piece of tactical advice: let us do away with the idea of representation. It ends with an understanding of what this tactical advice amounts to. We can focus on the extent to which an element of a cognitive processes stands in a normative relation to the environment, the extent to which it carries information about the environment, the extent to which it is capable of misrepresenting the environment, is decouplable from the environment, forms part of a larger combinatorial structure, plays a role in explaining (producing, guiding) behaviour. There may well be other constraints to be identified. But what we cannot after we have examined these issues is ask ourselves another question: Is this element of a cognitive process a representation or is it not? To the extent this question means anything at all it is reducible to the questions that we have already and (hopefully) answered.

References

- Cummins, R. 1983. *The Nature of Psychological Explanation*. Cambridge: MIT Press.
- Dreyfus, H. 2002. Intelligence without representation: Merleau-Ponty's critique of mental representation. *Phenomenology and the Cognitive Sciences*, 1(4): 367–83.
- Gallagher, S. 2008. Are minimal representations still representations? *International Journal of Philosophical Studies*, 16(3): 351–69.
- McDowell, J. 1994. *Mind and World*. Cambridge: Harvard University Press.
- Millikan, R. 1984. *Language, Thought and Other Biological Categories*. Massachusetts: MIT Press.
- Rowlands, M. 2006. *Body Language: Representation in Action*. Cambridge: MIT Press.
- O'Shaughnessy, B. 1980. *The Will (2 Volumes)*. Cambridge: Cambridge University Press.
- Sellars, W. 1956/1997. *Empiricism and the Philosophy of Mind*. Ed. R. Brandom. Cambridge: Harvard University Press.