NONCONCEPTUAL REPRESENTATIONS FOR ACTION

1. Introduction

The aim of this paper is to provide a defense of the notion of nonconceptual content in the domain of action.

A general defense of the notion of nonconceptual content involves showing that this notion is both coherent and explanatorily useful. To show that it is coherent, in other words to show that the phrase 'nonconceptual content' is not an oxymoron, one must provide a characterization of nonconceptual content that makes it clear both that it is distinct from conceptual content and that it is nevertheless worthy of the name content. To show that it is explanatorily useful is to show that use of this notion helps us give an account of certain phenomena that is more perspicuous than accounts that either appeal to conceptual content or do not appeal to content at all.

The notion of nonconceptual content is a relative newcomer on the philosophical scene. G. Evans (1982) seems to have been the first philosopher to use the terminology of nonconceptual content. One should note, though, that the way Dretske was already characterizing what he called non-epistemic seeing in *Seeing and Knowing* (1969) presents striking similarities with the characterization Evans later offered of nonconceptual content. Anyway, the idea has been gaining ground in recent years, thus compelling an increasing number of philosophers to take sides on the issue whether the notion of nonconceptual content is philosophically sound and useful. As is clear in both Dretske's and Evans' writings, the original impetus for the development of this idea came from the study of perceptual experience. Both wanted to describe a way of perceiving that they saw as belief-independent and as more primitive in a sense than a way of perceiving connected with judgement and belief. This connection between talk of nonconceptual content and perceptual experience is still very strong today. Most arguments pro and contra nonconceptual content have been formulated in the context of discussions of perceptual experience (Peacocke, 1992, 1998; McDowell, 1994, 1998; Crane, 1992; Brewer, 1999, Bermúdez and Macpherson, 1998). Visual experience has been the main focus of interest, but nonconceptual content has also been appealed to in discussions of musical experience, kinesthetic and proprioceptive experience and the experience of pain. The purview of nonconceptual content has also recently been extended to accounts of primitive forms of self-consciousness (Hurley, 1997; Bermúdez, 1998). Still another area in which the notion of nonconceptual content has been put to use is that of subpersonal computational states. Yet, with but a few exceptions (O'Shaughnessy, 1980; Proust, 1999), nonconceptual content has played very little role in discussions of the nature and content of the representations of action. As a consequence, attempts to motivate positing nonconceptual content have almost never adduced considerations arising from the study of action.

A specific defense of the notion of nonconceptual content as applying to representations of action involves in turn two interrelated tasks. One is to show that the notion of nonconceptual content does useful explanatory work in the domain of action. The second is to give a positive account of the nonconceptual content of representations of action. Of course, a successful

defense of the notion of nonconceptual content in the domain of action should allay the general worries raised by philosophers who are skeptic about the very coherence of this notion. There are also more specific worries it should come to grip with. Even philosophers well disposed towards the idea of nonconceptual content in the domain of perception may refuse to grant the existence of such content in the domain of action. One possible explanation for this negative attitude comes from the traditional asymmetry in the philosophical treatment of the input vs. the output side of cognitive systems. Whereas perceptual states and processes are commonly treated as falling within the scope of psychological or mental states and processes, the states and processes involved in action preparation have typically been conceived as falling outside the realm of the mental and as belonging to the sole province of physiology. Moreover, even philosophers willing to grant the existence of nonconceptual representations in action preparation may balk at the idea that these nonconceptual representations form a *sui generis* kind. Since representations for action are bound to include information about the environment in which the action takes place as well as about the agent himself, it may indeed be thought that the nonconceptual ingredients in action representations are simply species of nonconceptual perceptual representations (e.g., visual, tactile, kinesthetic or proprioceptive).

Thus, our task is to show not just that representations with nonconceptual content are explanatorily useful in the domain of action, but also that these representations are distinct from the kinds of nonconceptual representations thought to account for certain aspects of perceptual experience. Note that since nonconceptual content is first characterized negatively in terms of the features definitory of conceptual content that it lacks, this general characterization leaves room for the possibility that representations differing along a number of dimensions satisfy the criteria for being nonconceptual. What I mean here is not just that we can expect nonconceptual representations to differ with respect to the realms of objects and properties they might represent. It is indeed obvious enough and rather trivial that if there exist nonconceptual representations of, say, auditory events, visual scenes or actions, there will only be a partial overlap among the objects and properties they will represent. More interestingly, I think, nonconceptual representations in various domains may also be structured in different ways and allow for different forms of (proto-)compositionality and cognitive integration. The more interesting possibility I am hinting at is that the representational formats of nonconceptual representations in various domains may present idiosyncratic features. In particular, there may well be several distinct ways of satisfying the criteria for being states with nonconceptual content.

I will proceed as follows. First, I shall offer a brief characterization of the distinction between conceptual and nonconceptual content. Second, I shall argue that standard models of action explanation are incomplete, that they leave a gap between the mental antecedent of an action and the specific act it causes, and that in order to fill this gap we need to introduce representations of the movements that will constitute the intended actions. The next step in the argument will involve showing that these representations cannot be conceptual on pain of vicious circularity. This will provide the first part of our defense of nonconceptual content in the domain of action. The second part of this defense will involve giving a positive characterization of these nonconceptual representations. This will allow us in turn both to show that they are indeed worthy of the name 'representation' and to see how they differ from the nonconceptual representations thought to account for perceptual experience.

2. Conceptual vs. nonconceptual content

Let me start by acknowledging that the debate surrounding the existence of nonconceptual content is in part a verbal dispute. It is a verbal dispute in the sense that different philosophers

have more or less stringent criteria for concepthood. Yet verbal does not necessarily mean illusory. It may well be that the criteria that some take to be sufficient for concepthood are deemed necessary but not yet sufficient by others. Insofar as this is the case, there is an element of stipulation in the distinction between conceptual and nonconceptual content. But as long as there are actually criteria that permit interesting discriminations among kinds of contents, the distinction is not unfounded. I have no original view to offer as to what the distinction amounts to. Rather, I will avail myself of criteria proposed by Bermúdez (1998) for drawing the line between the conceptual and the nonconceptual. His strategy is as follows. He starts by proposing a number of criteria that mental states should satisfy to qualify as contentful, where a characterization of contentful states by means of these criteria remains neutral as between conceptual content and nonconceptual content. He then offers further criteria for content to be conceptual. The difference between the two kinds of content is then that nonconceptual content must only meet the first set of criteria whereas conceptual content must meet both.

Let me start with the first set of criteria Bermúdez proposes and the motivations he offers for them. One first criterion one might want to set for a state to be representational is that it have what Peacocke calls *correctness conditions*: A state presents the world as being in a certain way if and only if there is a condition or sets of conditions under which it does so correctly and the content of the state is defined in terms of these conditions. As Bermúdez notes, this condition, although necessary, is certainly not sufficient. For instance, it does not allow one to distinguish between genuine representational states and what Dretske calls information-carrying states, where a state is said to carry information about another state if and only if there is a nomic covariation between the two kinds of states. One classical example is the relationship between the number of rings on a tree and the age of a tree. If having correctness conditions were sufficient for being representational, the number of rings could be said to represent the age of the tree. An even more extreme consequence would be that everything could be said to be self-representational. For instance, my wearing a hat would be deemed to represent that I am wearing a hat insofar as we could always give the following correctness condition: "My wearing a hat represents correctly that I am wearing a hat if and only if I am wearing a hat if and only if I am wearing a hat." One further condition we want to impose on states to count as representations is that they leave open the possibility of *misrepresentation*. In other words, it should be possible for a state to have correction conditions that are not always satisfied. This rules out automatic self-representations as well as less trivial instances of nomic correlations as genuine cases of representations.

One further condition Bermúdez sets for representations is that they should be explanatorily useful. In other words, it is legitimate to give explanations that appeal to representational states (intentional explanations, for short) only if what is thus explained could not be accounted for by explanations that did not appeal to contentful states. His suggestion is that we need to appeal to representations in situations where the behavior to be explained cannot be accounted for in terms of invariant relations between sensory input and behavioral output. There are in turn several reasons why the lawful correlation between stimulus and response may break down. One connects back with the idea of misrepresentation. A stimulus may be present and yet fail to trigger a response if it is not represented in the way suited to generate the response. Conversely, if a response is produced in the absence of the appropriate stimulus, one may say that it was brought about by an incorrect representation of the presence of the stimulus. Second, the lack of lawful correlations between stimulus and response can reflect, more positively, the plasticity and flexibility with which an organism responds to his environment. This plasticity and flexibility can be explained as the result of complex interactions between the internal states of the organism.

Thus, if representational states are to play a role in the explanation of plastic behavior, they must also admit of *cognitive integration*. Bermúdez further distinguishes two aspects of cognitive integration. First, it means that representational states have to interact with other representational states, and therefore that there must be pathways enabling them to connect up with one another. Second, cognitive integration demands that a creature representing the environment be capable of registering when it is relevantly similar over time and to register in what respect a given situation is similar to a previously encountered one in order to adapt its response.

The demands of cognitive integration make it necessary in turn that representational states be structured and exhibit some form of *compositionality* to allow recognition of partial similarities and to allow at least some primitive forms of inference. Regarding compositionality, Bermúdez stresses two important points. First, compositionality is a graded notion, going from partially recombinable content constituents to fully recombinable content constituents. Second, he also insists that, although for there to be structure at the level of content, there must be structure at the level of vehicle, compositionality at the level of content need not entail strong compositionality at the level of the vehicle of content of the kind envisaged by Fodor and Pylyshyn (1988).

To recap, Bermúdez proposes the following criteria for a state to qualify as representational: (1) it should have correctness conditions and allow for the possibility of misrepresentation; (2) it should play a role in the explanation of behavior that cannot be accounted for in terms of invariant relations between sensory input and behavioral output; (3) it should admit of cognitive integration; and (4) it should be compositionally structured.

What more is needed for states to have conceptual content? First, note that both cognitive integration and compositionality are graded notions. So, one way of drawing the distinction between conceptual and nonconceptual representations would be to say that conceptual representations must satisfy more stringent criteria of full cognitive integration and full compositionality. Indeed, Bermúdez suggests that the distinction between conceptual and nonconceptual content may in part be a matter of degree of compositionality and cognitive integration. Evans's Generality Constraint for concepts may be seen as a version of this claim. Bermúdez also insists on a further criterion that appeals to a capacity for critical reflection on one's own thinking. Underlying the distinction between thoughts with conceptual content and thoughts with nonconceptual content is the distinction between having thoughts and making transitions between thoughts, on the one hand, and being able to reflect on the thoughts in question and the transitions between them on the other. Thus, Bermúdez' contention is that there is a constitutive link between a capacity for conceptual thought and a capacity for genuine inference, where having a capacity for genuine inference is linked to an ability to appreciate the rational grounds for, and thus to justify, one's inferences. This in turn leads Bermúdez to the further claim — what he calls the Priority Thesis — that there is a constitutive link between language mastery and concept mastery. His main reason for that claim is that providing justification is a paradigmatically linguistic ability.

One may, following Levine (2001), have reservations regarding the two claims that a capacity for genuine inferences is constitutively linked with a capacity to justify these inferences and that a capacity for justification requires language mastery. As Levine points out, the first claim seems to be based on the assumption that in order to be justified in an inference the subject must not only possess the concepts employed in the inference but also the concept of justification. However, this internalist and personal-level conception of justification cannot just be assumed, it must also be argued for. Second, one may concede that it is difficult to see how, without language, one could justify one's judgements to someone else, but deny that the

claim that without language one cannot convey justifications to others automatically entails the further claim that without language there can be no justification. However, acceptance of this controversial criterion of justificatory competence and of the Priority Thesis is not necessary for our purpose. We may say that what sets conceptual content apart from nonconceptual content are simply the requirements of full compositionality and full cognitive integration. It may also be suggested that what sets apart the full-blown inferential capacities linked to concept possession from the more limited sensitivity to the truth of inferential transitions already present in nonconceptual representations is not a matter of justification in Bermúdez' sense, but rather something akin to what Fodor calls isotropy: justifiers and defeaters for certain inferred conclusions can come from anywhere in the system. A capacity for genuine inference in this latter sense is a consequence of full cognitive integration and full compositionality rather than being a separate characteristic of conceptual thought. Finally, although one may not be convinced by Bermúdez defense of the Priority Thesis, it may be possible to provide a different defense of this claim and to argue that there can't be concepts without language, not because justification requires language, but because full cognitive integration and compositionality require language. I won't take a stand of this issue here, I will just assume that the distinction between full and partial cognitive integration and compositionality suffices to motivate the distinction between conceptual and nonconceptual thought.

What I will now try to show is that there exists in the domain of action a specific explanatory job that standard models of action explanation do not undertake, that we must provide such an explanation if action is to be made fully intelligible, that it is necessary to appeal to a special type of mental representations to provide this explanation, and, finally, that these mental representations are not conceptual.

3. Standard models of action explanation and the problem of the wrong movement

According to the classical version of the Causal Theory of Actions[1], what distinguishes actions from mere happenings is the nature of their causal antecedents. Genuine actions are events with a distinctive mental antecedent. The relevant causal antecedent is conceived as a complex of some of the subject's beliefs and desires. According to Davidson (1980: Essay 1), for instance, the causal antecedent of an action is a combination of a pro-attitude toward actions with a certain property P — such as bringing about a certain result or state of affairs — and a belief that a certain action A has that property. This analysis can be somewhat refined by further distinguishing between an orienting belief of the agent — a belief that he is in circumstances C — and an instrumental belief that action A in C has property P.

Several features of this standard version of the causal theory may be held responsible for the attraction it has exerted. First, it has, one may say, the advantage of killing two birds with one stone, by simultaneously offering an account of the nature of action and an account of the explanation of action. What distinguishes actions from non-actions is the fact that they have a distinctive antecedent, namely a belief-desire complex, but this belief-desire complex also provides a causal explanation of the action. Second, and relatedly, the theory brings into line the justificatory and the explanatory role of reasons by insisting that in cases where reasons genuinely explain, the reason-providing intentional states cause the actions for which they provide reasons. The belief-desire complex is not just the causal antecedent of the action; it also provides the material for a reason-giving explanation of the action. The structure of such an explanation can be made readily apparent by presenting it as a form of practical reasoning, where the elements in the desire-belief complex provide the premises for the reasoning and where the action is its conclusion. Thus, the standard theory fosters the hope of narrowing the

gap between the normative realm of explanations by reasons and the natural realm of causal explanation.[2] Third, it may be commended for its ontological parsimony. It does not postulate any special type of mental events such as willings, volitions, acts of will, settings of oneself to act, tryings, etc. and thus seems to avoid the charge of obscurantism that has been raised for theories hypothesizing such entities. According to the theory, to say that somebody acted with a certain intention is just to say that his actions stood in the appropriate relations to his desires and beliefs.

However, as a number of critics have pointed out, the standard causal model of action explanation is faced with several difficulties. I will here consider only one of them, pointed out by Israel, Perry and Tutiya (1993) and by Dokic (1999)[3]. Israel *et al.* call this difficulty the 'problem of the wrong movement'[4]. The problem is that the failure of an action cannot always be traced back to the falsity of some belief figuring in the motivating complex as it conceived of in the standard model. In other words, the truth of the beliefs figuring in the belief-desire complex does not guarantee that the bodily movement performed by the agent is appropriate. For instance, if we take Israel *et al.*'s original example, Brutus might intend to kill Caesar by stabbing him. His orienting belief that Caesar is to his left and his instrumental belief that stabbing Caesar in the chest would kill him may both be true and yet Brutus may fail to kill Caesar because he makes the wrong movement and misses Caesar completely. His movement could be wrong in two different ways. It could be the case that he simply does not make the movement he was trying to make or it could be that he makes the movement he was trying to make but that that movement is not the right one — in other words, Brutus is wrong to think that that movement would constitute a stabbing of the person to his left. Israel *et al.*'s conclusion is that something is missing in the traditional account: "Brutus' motivating complex needs to reflect which movement he is trying to make, and what he thinks its effect will be" (1993: 528).

The motivating complex as it is conceived in the standard account leaves a gap to be filled between the motivating cognitions and the act itself. If we consider, as I think we should, that the explanandum in a theory of action explanation is the act itself, not just the attempt or volition, we should be ready to include in the explanans cognitions pertaining to movements. When the agent's orienting and instrumental beliefs are correct, what ultimately accounts for the success or failure of an intended action are the bodily movements performed. Thus, if to explain an action is to make it intelligible, then full intelligibility requires that the bodily movements themselves be made intelligible by the cognitions that motivate them. I will call movement representations these cognitions that are meant to fill the gap between the motivating cognitions and the act they cause and thus render intelligible the fact that the agent performs one movement rather than another.

On Israel *et al.*'s model, when beliefs and desires motivate acts they do so by causing what they call a volition to execute a movement of a certain type, where the volition is something akin to what Searle (1983) calls an intention-in-action and is in turn the proximate cause of the movement. For this to be possible the motivating complex must include not just a desire for a certain result, an orienting belief, and an instrumental belief, but also what they call a belief-how. A belief-how is different from an instrumental belief. An instrumental belief is a belief that performing a certain action is a way of performing some other action. For instance, stabbing is a way of killing. By contrast a belief-how concerns the relations between movements and action. A belief-how is a belief that performing a movement of a certain kind in certain circumstances is a way of bringing about a certain result (hence of performing a certain action) in those circumstances. Moreover, for an agent to have such a belief-how, not just any idea of the type of movement in question will do, the agent must have an executable idea of it. That is, he must know or think he knows how to perform movements of this type.

It is clear that in this model the gap between the motivating complex and the act it motivates if filled by a belief-how and the volition it causes. Yet, what remains somewhat unclear is how exactly Israel *et al.* conceive of the movement representations involved in the beliefs-how and the volitions they cause. They could have in mind a detailed representation of a very specific bodily movement (or sequence thereof) where this would also require of the agents that they have very detailed beliefs about the circumstances, since exactly what movements should be performed depends on what exactly the circumstances are. Or they could have in mind a more schematic representation of a type of bodily movement. As Dokic points out, Israel *et al.* appear to be faced with a dilemma. The first option is not very plausible. The volition is thought to cause the movement and thus must precede it. But requiring that the agent form a very detailed representation of the movement before its execution threatens to impose on him an excessive cognitive burden, especially if the movement is somewhat complex. On the other hand, if we follow the second option and let the representation of the movement be schematic, we will not be in a position to account for the specific conditions of success of the action. Suppose, to take Dokic's example, that the content of my volition is simply that I perform a movement of a type that will result in my seizing the glass of water in front of me. The problem is that the success of my action depends not just on its result but also on the way this result is achieved. There is an indefinite number of bodily movements that would result in my seizing the glass, but not all of them would constitute what I would consider as the natural, harmonious, successful unfolding of my action. If my hand were to waver in space for a full minute before finally reaching the glass and grasping it awkwardly with the thumb and little finger, I would not in most circumstances consider my action as fully successful.

However, there is one way out of the dilemma. The dilemma depends on the assumption that all the representational work has to be done and be over with before the execution starts. But nothing forces us to accept this assumption. Actually it is quite implausible. Representations do not just trigger the action; they also guide and control it until its completion. One way to account for the specific conditions of success of my action, while avoiding the problem of cognitive overload, is to consider that the representations that guide and control the movement are not fully specified before its onset but are dynamical and relational and serve to adjust the movement to the changing context in which it unfolds. As another way to put it, the degree of specification of the bodily movements depends on the degree of specification of the circumstances. It is implausible to suppose that an agent represents all the circumstances in advance of the action. The point here is not just Israel's *et al*. point that we are attuned to a limited range of environments and that facts that are uniform or stable in those environments need not be explicitly represented because they are already reflected in the way agents are built both physically and cognitively. Rather, the point is that even if we limit ourselves to circumstances that are variable and relevant to the success of the action, we need not explicitly represent them in all their details before starting to act. It is enough that the agent represents the circumstances to the extent that he can see that they allow for an action that would bring out a certain result to be performed and also see that executing movements of a certain type would in those circumstances be a mode of accomplishing this action. In other words the agent is trusting in two different ways. He implicitly trusts it to be the case that the uniformities of nature will remain uniform, but he also trusts it to be the case that he will be able as the action unfolds to pick up information about the details of his circumstances to further guide and adjust his movements. Very general and stable circumstances don't have to be explicitly represented at the level of the motivating complex, because they are already reflected in the way we are built. Very specific and transitory circumstances don't have to be represented at the level of the motivating complex because they will be represented as the need arises in the motor representations that will guide the action. My suggestion then is that

the variable circumstances that have to be represented for an action to be executed divide into macro-circumstances, represented in advance of action and at the level of the motivating complex and into micro-circumstances that are represented as the action proceeds. I further suggest that macro-circumstances are represented in a conceptual format, whereas micro-circumstances are nonconceptually represented.

This in turn allows us to give a better account of the problem of the wrong movement. Recall that this problem is twofold. The agent may be wrong in thinking that performing a movement of a certain type would bring about some intended result or, although he is right in that respect, the movement he performs is not of that type. Israel *et al.* have a ready explanation for wrong movements of the first category. Here, the agent's belief-how explains why he makes the movement he does and the incorrectness of his belief-how explains why his movement is wrong. But the only thing they can say about the second category of wrong movement is that the volition to execute a movement of a certain type does not cause a movement of that type. The wrong movement is in no way rationalized, what we have is only the brute fact of a causal failure. If instead of just considering the volition that causes the onset of the movement, we also consider the dynamical representations internal to the execution of the action, failures of the second kind may also be made intelligible— at least in some cases. [5]

The first step in this argument in favor of nonconceptual representations of actions involved establishing the existence of a specific explanatory job. Our discussion of the problem of the wrong movement was meant to show that ultimately what action explanation must make intelligible is why such and such bodily movements are performed. This requires that we consider cognitions pertaining to movements. Furthermore, as Israel *et al.* acknowledge, the problem of the wrong movement is twofold, with two different categories of wrong movements. I argued that although Israel *et al.* can account for wrong movements in the first category in terms of the kind of movement-related cognitions they postulate (beliefs-how), they give no proper explanation for wrong movements in the second category. I also argued that we need to appeal to a further level of movement representations in order to account for wrong movements in this second category. Let me call representations at this second level motor representations.

We shall now turn to the second step of the argument, that involves showing that at least some of these movement representations that account for the movements performed, whether successful or not, are not conceptual. Since we have two levels of movement representations, there are two possible options for a nonconceptualist. He or she may either want to argue that representations at both levels are nonconceptual or that representations are conceptual at the level of beliefs-how but nonconceptual at the level of motor representations. I'll take the second option.

I am willing to grant that the representations of bodily movements involved in belief-how are conceptual in kind. One reason for thinking they should be is that beliefs-how can occur together with instrumental beliefs, orienting beliefs and desires as premisses in practical reasoning. And assuming, as philosophers normally do, that these other beliefs and desires have conceptual content, it would be difficult to explain how beliefs-hows could be inferentially related to these other beliefs and desires if they did not themselves have conceptual content. Now, a belief-how is supposed to involve a representation of a type of movement seen as a mode of a bringing about a certain result. Moreover, not just any representation of the movement will do, it must be what Tutiya *et al.* call an executable idea, something that could guide the formation of a volition. So we are now looking for representations of movements that meet the two conditions of being conceptual and being

executable. What I will try to show is that the concepts of movements that would fit the bill can only be possessed by agents that already have executable nonconceptual representations of movements.

First we should note that it is typically very difficult to articulate what the precise characteristics of a type of movement are. When asked to explain what kind of movement you would make in order to reach for the glass in front of you, your only way of answering the request might be to say something like 'doing like this', while showing the movement in question. As Dokic (1999) notes, in most cases we seem unable to offer in place of the demonstrative 'this', that designates a movement of a certain type, a purely discursive definition. If one were to require that conceptual representations be fully discursive representations, our quest would be doomed to failure. Given that conceptual representations of movements cannot be fully discursive, the only possibility that remains open is that these representations have as constituents demonstrative concepts of movements. So what we have to do is give an account of what the possession of an executable demonstrative concept of a movement involves, an account that avoids both circularity and regress. Now, my contention is that the only way to avoid circularity or regress is to admit the existence of nonconceptual representations of action. Here is the argument. For an agent to be said to possess a demonstrative concept of a movement, he must be able to demonstrate the movement in question, that is, to execute it. Therefore, to be able to entertain a demonstrative concept of a type of movement, one must already be able to perform movements of that type. But this capacity to perform movements of that type requires in turn that the agent already have executable representations of those movements. If there were only one level of representation of movement and this level was conceptual, the circularity would be blatant. Being capable of executing a movement of a certain type would be a precondition of possessing a demonstrative concept of movements of that type. But possessing this demonstrative concept would in turn be a precondition on being able to perform the movement in question, since for one to be able to perform the movement, one must have an executable representation of it and this representation would have to have among its constituents a demonstrative concept of the movement. Acknowledging two different levels of representations of movements while maintaining that both of them are conceptual will not help, for then the choice will be between circularity and regress. If the same movement concepts are thought to figure in both types of representations, the circularity remains. If it is thought that different, perhaps more primitive, concepts of movements figure at the lower level, we must give an account of the possession of these concepts. Since there is no reason to think that we can have at the lower level fully discursive representations of the movements we can't discursively define at the higher level, we will once more have to think that these concepts are demonstrative, and then we will be engaged in a regress, since we will have to appeal to some third level of representation of movement.

But, the conceptualist may well retort, why insist that the demonstrative pick up a movement of the agent himself and why insist that performance of this movement involve an executive representation of it? The answer is that it is our second condition, namely that the concepts be executive concepts of movements that makes it necessary. First, note that the suggestion that the demonstrative component of the demonstrative concept need not designate a movement performed by the agent himself will not help the conceptualist. If one could indeed always acquire an appropriate concept of a movement by simply watching someone else perform a movement, thus using 'like this' to designate the observed movement, the circularity would be avoided. But we must distinguish between observational concepts of movements and executable concepts of movements. If I am a spectator at an ice-skating competition, I may indeed form demonstrative concepts of the movements performed by the skaters based on my

observation of them. But the concepts in question may well fail to be executable concepts, where an executable concept of a movement would be one that figures in an executive representation that can cause me to perform the movement in question. That I formed a concept of a certain movement by watching an ice-skater perform a triple-axel in no way guarantee that I will be able to perform this movement myself. Similarly, it would not be enough that the demonstrative component designate some movement my body makes, if the movement is simply a movement I can observe my body make but not a movement I can effect. If what is designated by the indexical component is simply a bodily happening and not a movement the agent can effect, the indexical concept will not be an executable concept. Thus, a genuinely executable concept must hook up to movements that are already under the control of motor representations. An executable concept is a concept of movements already executable by the agent, not a concept of movements that in some mysterious way would make them executable. Conceptualizing movements is not to make them executable, rather the point of this conceptualization of movements is to allow us to connect a capacity for executing movements of a given type to what Israel *et al.* call an accomplishment, i.e. the bringing about of a certain result.

Let me take stock. I have argued that for an agent to have an executable concept of a movement is to have a demonstrative concept the indexical component of which a movement of a kind that is already executable by the agent. Executable movements are movements that can be guided and controlled by the agent, i. e. movements that are under the control of motor representations. Thus to have an executable concept of a movement, an agent must already have representations of movements. On pain or circularity or regress, the latter representations cannot themselves already involve concepts of movements. In other words, the only way to make sense of executable concepts of movements is to acknowledge the existence of nonconceptual representations of movements. We will now turn to the second part of this defense of nonconceptual representations of actions and this will involve giving a positive characterization of these nonconceptual representations.

4. Format and content of motor representations

Philosophical characterizations of what exactly motor representations consist of and of how they fulfil their role are usually rather tentative and less than fully specific. Yet, by combining philosophical analyses with neurophysiological work on motor representations, one can obtain a reasonably precise characterization of these representations[6].

Here, I will briefly draw on Jeannerod's (1997) work on motor representations[7]. Based on neurophysiological evidence, Jeannerod argues in favor of the following theses. First, he claims that actions are driven by an internal representation of a goal rather than directly by the external world. He warns us against an artificial separation of movement representations, assumed to pertain to a physiological approach, and action representations, assumed to pertain to a psychological approach. His claim is that there is no such dichotomy but rather a continuum. More precisely, there is a hierarchy of motor representations such that the goals and parameters of the actions coded for at the higher levels act as constraints on the lower levels of motor representations

Second, Jeannerod contends that the motor representations that drive the action have a specific content, involving two main aspects: a representation of the body in action as a generator of forces and a representation of a goal of action encoded in a 'pragmatic' mode, distinct from 'semantic' modes of representations. With respect to the first aspect, Jeannerod insists that the motor representation is a representation of the acting self that involves a representation of the body as a generator of acting forces, not just a representation of the effects of these forces on the external world. Experimental studies reviewed by Jeannerod

(Decety *et al.* 1993; Gandevia 1982, 1987; Gandevia and McCloskey, 1977; McCloskey *et al.* 1983) suggest that the amount of force needed to produce the desired motor effect is encoded in this component of the representation. Moreover, experiments with completely, or partially, paralyzed patients (Gandevia 1982; Jeannerod 1994; Scheerer 1987) suggest that the programming of force has a subjective correlate — the sensation of effort. Empirical evidence also suggests that the central representation of action encodes certain parameters of movement execution dictated by kinematic rules (Decety and Michel 1989; Georgopoulos and Massey 1987; Georgopoulos *et al.* 1989; Viviani and McCollum 1983) and biomechanical constraints (Rosenbaum *et al.* 1990; Rosenbaum and Jorgensen 1992; Shiffrar and Freyd 1990).

The second essential aspect of a motor representation is a representation of the goal of action. According to Jeannerod, this representation includes a representation of both the external object toward which it is directed, and the final state of the organism when that object has been reached. In simple, object-oriented actions (i. e., when objects are goals for actions), the visual attributes of those objects are represented in a specific, 'pragmatic' mode used for the selection of appropriate movements and distinct from other modes of representation used for other aspects of object-oriented behavior (categorization, recognition, etc.).[8] Jeannerod suggests that a motor representation of a goal object includes both a visuo-spatial component pertaining to its spatial location and an object-centered component determining how to deal with it. He also suggests that the function of those representations 'falls between' a sensory function (extracting from the environment attributes of objects or situations relevant to a given action) and a motor one (encoding certain aspects of that action). Pragmatic representations involve a rapid transformation of sensory input into motor commands. Object attributes are represented to the extent that they trigger specific motor patterns. Pragmatic representations thus specify how to deal with the object. In other words, in a pragmatic representation, object attributes are treated in a causally indexical way (Campbell, 1993, 1994)[9], or to use a different terminology as 'affordances' (Gibson 1979)[10], activating predetermined motor patterns.

Although for expository purposes I distinguished between two aspects of the content of motor representations, it would be mistaken to assume that they correspond to two separate components of the content. Rather, motor representations should be viewed as relational models, with the body and the goal functioning as the terms of the relation. What the motor representation represents are neither states of the body *per se* nor states of the environment *per se*, but rather dynamic relations between body and goal. To use a different formulation, we could say that the goal is given under a specific mode of presentation, it is represented in terms of the motor patterns that it affords to the agent.

Another important aspect of motor representations is their dynamical character. Motor representations are involved in the guiding and monitoring of the action as long as it unfolds. How is this role reflected in their content? In order for the motor representation to guide the action, it must anticipate the consequences of the movements. In order to control it, it must allow for adjustments during execution. Feedback (be it visual, kinesthetic, or proprioceptive) will be needed to make adjustments. The content of the representation is dynamical in the sense both that it gets elaborated over time — it becomes more determinate through feed-back— and that the motor representation is itself responsible for making available the information that will make the content more determinate. The only way to gain the extra information (the feedback information) is for motor representation to create the context in which the information will be available. Perhaps we could say here that an important feature of the content of intentions in action is its dynamical indexicality: indexicality because the exact value of certain constituents of the representation can only be fixed relative to a context;

dynamical indexicality because the context itself must be brought into existence by the representation.

Jeannerod suggests that the same general framework used for simple object-oriented actions remains applicable to higher-order representations encoding long-term action plans. The lower-level object-oriented motor representations should be considered as constituents of more complex action representations that will also need to involve higher-order schemas for controlling the selection, the activation, the inhibition, and the sequencing of elementary motor schemas.

To better see the characteristics of motor representations, let us consider an example of a particular type of action, say, placing an object in a box, and see what it involves.[11] First, this action, despite being relatively simple, involves a number of subcomponents, namely reaching, grasping, lifting, transporting and releasing the object. Reaching means transporting the hand to some location within extrapersonal personal space. It involves a computation of the direction and distance of the object with respect to the body (as well as if the object is moving, a computation of its direction and velocity of motion) and a visuomotor transformation where the position of the object in space is converted from its representation in the system of coordinates used for vision into a representation using a proprioceptive system of coordinates and where the direction and amplitude of the movement to be performed, its kinematics, and the amount of force to be programmed are specified. Grasping involves a different type of visuomotor transformation: what must be computed are intrinsic properties of the object, such as size and shape, in order to determine the shape the hand must form in order to deal with object. Here the transformation to be effected is from the an object-centered representation of the shape and size of the target-object to a representation in the system of coordinates used for representing the movements of hand and fingers and the relative posture of the thumb and the other fingers. Moreover, grip is not determined solely by the shape and size of the object, it is also a function of intended activity. For instance, the same object may be held with a precision grip or with a power grip depending on whether I intend to put it in a large box or to insert it in a tight-fitting container or it may be held with an overhand or an underhand grip depending on its future use. Lifting the object involves using both visual and tactile information about size, texture, compliance, hardness and weight to compute the appropriate load force. But load force is also a function of grip force; the grip force/load force ratio must be such as to exceed the slip ratio, itself determined by the coefficient of friction between the skin and the object surface. Transporting the object to the box will in turn involve computing the distance between the object and the box as well as the egocentric location of the box and transforming this information into a representation of the movement to be executed and of its kinematics. In particular, the kinematics obeys Fitt's law, a law that states that the speed of movement is inversely proportional to the difficulty of the task. Thus, if the box where the object is to be put is large, the movement should be fast, if the box is small, the movement will be slower.[12]

This brief description is enough to show that motor representations have constituent units, falling under a number of categories such as reaching, grasping or lifting. These constituents units can be combined in various ways. Although, in the example I took, the sequence follows a definite order (reaching, grasping, lifting, transporting, releasing), it is clear that depending on both intrinsic (shape, size, texture, weight, etc.) and extrinsic (distance and location with respect to the agent's body) properties of the target object, the representation could combine various reaching movements with various types of grip and manipulations. Different actions will involve different combinations of these and other categories of units and, at a higher level of organization, more complex actions will in turn involve combinations of relatively simple

actions such as putting something into a box. For instance, this action could be a recurrent element in the complex action of packing my suitcase before a trip.

Motor representations do not just have a lexicon; they also have what may be called a 'grammar' for assembling the constituent units into a coherent pattern. There are spatial, temporal and motor (kinematic and biomechanical) constraints on the coordination of action that must be reflected at the level of grammar. The coordination of reaching and grasping may serve as an illustration. First, the combination of reaching and grasping units must obey certain spatial constraints. Reaching is mostly achieved by the proximal joints of the arm and it makes use of an egocentric or body-centered system of representation of locations. Grasping on the other hand is a function of the intrinsic shape and size of the target object; it involves a transformation of visual information encoded in allocentric, object-centered coordinates into motor information encoded in the system of coordinates used to define the prehension space. Yet reaching and grasping must be spatially compatible. In particular, reaching must take into account not just the location of the object but also its orientation, so that the final position of the arm is compatible with the correct position of the hand and fingers for grasping the object. Second, reaching and grasping must also be temporally coordinated. Empirical evidence shows that their temporal coordination goes beyond mere succession. The fingers begin to shape during transportation of the hand at the object location. Preshaping first involves a progressive opening of the grip with straightening of the fingers, followed by a closure of the grip until it matches object size. The point in time where grip size is the largest occurs within about 60 per cent to 70 per cent of the duration of the reach, well before the fingers come in contact with the object. Moreover, the amplitude of grip aperture during grip formation covaries with object size. Paulignan *et al.* (1991) have shown that when the target object is suddenly displaced at the onset of the reach-to-grasp movement, an untrained subject is perfectly able to correct for this perturbation and to accurately grasp the displaced object. However, the correction results in prolonging the duration of the reach by about 100 ms; meanwhile the opening of the grip is interrupted, grip size decreases before increasing again until it reaches its peak aperture. The two actions or reaching and grasping are thus temporally coordinated in a very precise way. To account for these data, Hoff and Arbib (1993) have postulated the existence of a coordinating schema that receives information about the time needed for each constituent action and ensures their temporal integration. Finally, let me illustrate the existence of biomechanical constraints on action coordination with the following example. Rosenbaum *et al.* (1990) and Rosenbaum and Jorgensen (1992) found that the decision to form a given type of grip for grasping an object is made according to the subsequent use of the object so that awkward or uncomfortable hand positions are avoided and the time spent in extreme joint angles is minimized. They used a horizontally positioned bar that the subjects had to position vertically on either its right end or its left end. They observed that subjects consistently made overhand grasps when they had to place the bar on its right end and underhand grasps when they had to place it on its left end.

I have focussed here on the properties of motor representations for relatively simple actions. Although this discussion was far from exhaustive, I hope it gave a sense of the features that are characteristic of motor representations. In the last section of this paper, I will try to show both that motor representations thus characterized are indeed worthy of being called representations — that is they satisfy Bermúdez' criteria for representationality — and that they have distinctive features that set them apart from nonconceptual perceptual representations.

5. The distinctiveness of nonconceptual motor representations

Let me consider first the representational status of motor representations and see how they fare with respect to Bermúdez' criteria for representationality.

The first criterion for a state to be representational is that it have correctness conditions. One important property of motor representations is that they code simultaneously for things that are coded separately at the (conceptual) level of the antecedent motivating complex. The antecedent motivating complex for an action includes both conative and doxastic elements, namely, on the conative side, a desire that a certain result *R* obtain, and on the doxastic side, an orienting belief that one is in circumstances *C*, an instrumental belief that in *C* action *A* brings about *R*., and a belief-how that executing a movement of type *M* is a mode of *A*ing in *C*. Situation, goal and means are thus represented separately. There is no such dissociation at the level of motor representations. The situation is coded in terms of a goal it affords and the goal itself is coded in terms of the means — i.e. the motor commands —towards its achievement. (Proto-)doxastic and (proto-)conative elements are thus inextricably intertwined. One consequence of this intertwining is that the classical distinction between states with a mind-to-word direction of fit and states with a world-to-mind direction of fit gets blurred at the level of motor representations. A motor representation represents a situation as affording a certain goal, and it does so by representing the motoric means by which the goal is to be achieved. For instance, it represents an object as reachable by representing how the reaching is to be effected. The representation may be said to be correct if the object is (1) within reach and (2) the motoric means prescribed will allow one to reach it. It is incorrect if either the object is out of reach or, although it is within reach, the movement is not the correct one. One may want to say that the first condition involves a mind-to-world direction of fit — a motor representation is correct if it represents a situation as affording an action it actually affords, incorrect otherwise — and that the second condition involves a world-to-mind direction of fit — it is correct if the movements represented bring about a change in the world such that the intended result obtains, incorrect otherwise. Yet, the distinction remains somewhat artificial, for the two conditions I distinguished might as well be described as just one, namely 'reachable by such motoric means'. More generally, it may be suggested that although Searle's distinction between mind-to-world and world-to-mind directions of fit is useful as a way of contrasting beliefs and desires, its application to motor representations is not fully perspicuous. We can perhaps bypass this issue, if we characterize the correctness conditions of a motor representation as follows. A motor representation *M* specifies a type of dynamical pattern of motor relations between bodily circumstances and world circumstances. A tokening of *M* by an agent *A* at a certain place and time is correct if it causes a dynamical pattern of relations between bodily circumstances and world circumstances that instantiates the type specified by *M*; it is incorrect otherwise. This characterization of the correctness conditions of motor representations also makes sense of the idea that the success of an action does not just depend on the fact that a certain result is achieved, but also on the specific way in which the result is achieved. For a given motor representation to be correct, it is not sufficient that it causes some series of changes in the relations between body and world, where the last element in the series corresponds to some desired result, the changes must also conform to a certain dynamical pattern.

The discussion in the previous section should also have made it clear that motor representations have structure and exhibit some form of compositionality. That is, they have identifiable constituent units (such as reaching, grasping or lifting components) that can be combined in various ways but whose combinations should obey certain spatial, temporal and motor constraints. Motor representations also admit of cognitive integration, both with other representational states and with motivational states. We have seen, for instance, that how an

object is grasped is a function not just of its size, shape and orientation, but also of what we intend to do with it. Similarly, the way we interact with an object does not just depend on its shape or size but also on its function, where the function may not be visually salient. Thus, motor representations will be influenced by knowledge of function. More generally, our motor interactions with an object will often be determined not only by sensory information immediately available to the agent but also by her stored beliefs and knowledge regarding certain attributes and properties of the object. Motor representations also connect up with other motivational states. They exploit the motor affordances present in the agent's environment. But of course, except perhaps in very impoverished settings, the environment will present not just one but many affordances for action. Yet, we do not respond to all the solicitations for action that the environment provides. Which motor representations are formed and acted upon is not just a function of environmental saliencies, it can be determined in part by the agent's motivational states, her higher-order goals and intentions or her emotional states (Pacherie, 2001). Motor representations are thus cognitively penetrable to a certain extent and can be influenced by information coming from other sources.

Moreover, it may be argued that the cognitive integration of motor representations is not just a matter of motor representations being influenced by other cognitive states. The influence can also work in the other direction. In particular, there is evidence that our interpretation of the actions we observe others perform is influenced by our knowledge of sensory-motor patterns, a knowledge that is stored in our system of motor representations. Motor representations allow for certain primitive forms of inferences. For instance, the motor representations that are built when preparing for an action allow us to anticipate the consequences of certain motor commands (by predicting both the next state of the action and the sensory feedback). Neuroscientists here talk of forward models of action. It seems though, that these forward models can be used not just to predict the consequences of our own actions, but also to interpret observed movements and to infer their goal. Thus, motor representations may contribute certain premisses to cognitive systems engaged in the interpretation of intentional behavior.

The last criterion to be considered is explanatory usefulness. I argued in section 3 for the existence of a particular explanatory task — namely explaining why in his pursuit of a certain result an agent is making certain bodily movements rather then others — and I further argued that an explanation appealing only to conceptual representations of movements would not be sufficient. For nonconceptual motor representations to be vindicated, it must also be demonstrated that a purely mechanical explanation of the motor behavior would not do. According to Bermúdez, the need for explanations appealing to contentful states arises in situations where the behavior to be explained cannot be accounted for in terms of invariant relations between sensory input and behavioral output. Our discussion of the influence of cognitive and motivational factors on the construction of motor representations makes it clear, I think, that the motor behavior they are meant to explain could not be explained in terms of a lawful correlation between sensory stimulus and behavioral response. For instance, the same sensory stimulus (a horizontal bar in front of the agent) will be responded to with either an overhand or an underhand grip depending on what the agent intends to do. A mechanistic explanation may perhaps be enough to account for reflexes, but the movements we want to explain are relationally characterized movements — movements related to a certain goal - and as Bermúdez (1998: 86) suggests, for such movements we need intentional explanations.

As I pointed out in the introduction, a successful defense of nonconceptual content in the domain of actions involves showing not just that certain representations that play a role in the explanation of action are nonconceptual, but also that these nonconceptual representations are not simply nonconceptual perceptual representations and that their content is not the same as

the content of perceptual representations. Here we should distinguish between the issue of mode and that of content. One could argue that the nonconceptual representational states involved in the explanation of action are nonconceptual perceptual states pure and simple. Or one could argue for the more subtle claim that they are indeed states with a distinctive mode — different from the perceptual mode— but that they share with perceptual state the same type of nonconceptual content. I want to reject both claims and to that effect I will stress two important differences between nonconceptual motor representations as I characterized them here and nonconceptual perceptual representations as characterized by Peacocke (1992) in terms of scenarios.

Peacocke (1992) proposes to characterize the basic structure of perceptual content in terms of what he calls a scenario. Scenarios are defined as spatial types specifying the ways in which the space around the perceiver is filled out. Specifying a given scenario or spatial type involves specifying an origin and a set of axes as well as specifying for each point identified by its distance and direction from the origin, whether there is a surface there, and if so, what texture, hue, saturation, brightness, and so on it has at that point. The first important difference between motor representations and Peacocke's perceptual scenarios concerns their correctness conditions. According to Peacocke, a positioned scenario — a scenario together with an assignment to the origin and axes of real directions and places in the world and an assignment of time — is correct if the volume of the real world around the perceiver, with an origin and axes in the world fixed in accordance with the labeling of the scenario falls under the spatial type defined by the scenario. It is clear then that the correctness conditions of a scenario involve a mind-to-world direction of fit. Perceptual representations belong in the category of doxastic states taken in a wide (i.e., including proto-doxastic states). By contrast, as I discussed at the beginning of this section, motor representations are mixed conative-doxastic states with a two-way direction of fit. The second important difference between perceptual and motor representations concerns their structure. Perceptual content has spatial structure. The perceptual properties of the scene represented are coded in terms of their spatial location in an egocentric reference frame[13]. As we have seen, motor content also has spatial structure, but it has more than just spatial structure, it has temporal and motor structure as well. Whereas in the perceptual domain the constraints on the possible combinations of representational elements are mainly spatial — for instance, the same point in space cannot be represented as having two different colors at the same time —, in the motor domain possible combinations of representational units are also determined by temporal and motor constraints.

The difference in direction of fit is enough, I take it, to defeat the claim that nonconceptual motor states reduce to a species of nonconceptual perceptual states. Clearly, if the direction of fit is different, the mode is different. The difference in structure is enough in turn to defeat the claim that perceptual and motor states share the same type of nonconceptual content.

References

Bach, K. 1978. A Representational Theory of Action. Philosophical Studies 34, 361-379.

Bermúdez, J. L and F. Macpherson. 1998. Nonconceptual content and the nature of perceptual experience. *Electronic Journal of Philosophy*, Fall 1998, 77-85.

Bermúdez, J. L. 1998. *The paradox of self-consciousness*. Cambridge, MA: MIT Press.

Bratman, M. E. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Cambridge University Press.

Brewer, B. 1999. *Perception and Reason*. Oxford: Oxford University Press.

Campbell, J. 1993. The Role of Physical Objects in Spatial Thinking. *Spatial Representation*, eds. N. Eilan, R. McCarthy, and B. Brewer, 65-95. Oxford: Blackwell.

Campbell, J. 1994. *Past, Space and Self*. Cambridge, MA: MIT Press.

Crane, T. 1992. The non-conceptual content of experience. In T. Crane (ed.) *The Contents of Experience*. Cambridge, Cambridge University Press, 136-157.

Davidson, D. 1980. *Essays on Actions and Events*. Oxford University Press: Oxford.

Decety, J, and Jeannerod, M. 1996. Fitt's law in mentally simulated movements. *Behavioral and Brain Research*, 72: 127-134.

Decety, J. 1991. Motor information may be important for updating the cognitive processes involved in mental imagery of movement. *European Bulletin of Cognitive Psychology*, 4: 415-426.

Decety, J., and Michel, F. 1989. Comparative Analysis of Actual and Mental Movement Times in Two Graphic Tasks. *Brain and Cognition* 11: 87-97.

Decety, J., Jeannerod, M., Durozard, D, and Baverel, G. 1993: Central Activation of Autonomic Effectors during Mental Simulation of Motor Actions. *Journal of Physiology* 461: 549-563.

Dokic, J. 1999. L'Action Située et le Principe de Ramsey. *La Logique des Situations. Nouveaux Regards sur l'Écologie des Activités Sociales,* (Raisons Pratiques 10), eds. M. de Fornel and L. Quéré, 131-155. Paris: Éditions de l'Ecole des Hautes Études en Sciences Sociales.

Dretske, F. 1969. *Seeing and Knowing*. Chicago: The University of Chicago Press.

Evans, G. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.

Fodor, J. A. and Pylyshyn, Z. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Gandevia, S. C. 1982: The Perception of Motor Commands of Effort during Muscular Paralysis. *Brain* 105: 151-159.

Gandevia, S. C. 1987. Roles for Perceived Voluntary Commands in Motor Control. *Trends in Neuroscience* 10: 81-85.

Gandevia, S. C., and McCloskey, D. I. 1977. Changes in Motor Commands, as Shown by Changes in Perceived Heaviness, during Partial Curarization and Peripheral Anaesthesia in Man. *Journal of Physiology* 272: 673-689.

Georgopoulos, A. P., and Massey, J. T. 1987. Cognitive Spatial-Motor Processes. *Experimental Brain Research* 65: 361-70.

Georgopoulos, A. P., Crutcher, M. D., and Schwartz, A. B. 1989. Cognitive Spatial Motor Processes: 3. Motor Cortical Prediction of Movement Direction during an Instructed Delay Period. *Experimental Brain Research* 75: 183-194.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton-Mifflin.

Goldman, A. 1970. *A Theory of Human Action*. Englewood Cliffs, NJ: Prentice-Hall.

Hoff, B., and Arbib, M. A. 1993: Models of trajectory formation and temporal interaction of reach and grasp, *Journal of Motor Behavior*, 25: 175-192.

Hurley, S. 1997. Non-conceptual self-consciousness and agency; perspective and access. *Communication and Cognition*, 30, 3/4, 207-248.

Israel, D., Perry, J. & Tutiya, S. 1993. Executions, Motivations and Accomplishments. *The Philosophical Review* 102: 515-540.

Jeannerod, M. 1994. A Theory of Representation-Driven Actions. *The Perceived Self: Ecological and Interpersonal Sources of Self-Knowledge*, ed. U. Neisser. Cambridge: Cambridge University Press.

Jeannerod, M. 1997. *The Cognitive Neuroscience of Action*. Oxford: Blackwell.

Levine, J. 2001. The self, and what it is like to be one: Reviews of Bermúdez and Weiskrantz. *Mind and Language*, 16, 1, 108-119.

McCloskey, D. I., Colebatch, J. G., Potter, E. K., and Burke, D. 1983. Judgements about Onset of Rapid Voluntary Movements in Man. *Journal of Neurophysiology* 49: 851-63.

McDowell, J. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.

McDowell, J. 1998. Reply to commentators. *Philosophy and Phenomenological Research*, LVIII, 2, 403-431..

Milner, A. D. 1997: Vision without knowledge. *Philosophical Transactions of the Royal Society of London. Biological Sciences* 352: 1249-1256.

Milner, A. D. and Goodale, M. A. 1993. Visual Pathways to Perception and Action. *Progress in Brain Research*, eds. T. P Hicks, S. Molotchnikoff, and T. Ono, 317-337. Amsterdam: Elsevier.

Milner, A. D., and Goodale, M. A. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.

O'Shaughnessy, B. 1980. *The Will,* 2. Vol. Cambridge: Cambridge University Press.

Pacherie, E. 2000. The content of intentions. *Mind and Language* 15: 400-432.

Pacherie, E. 2001. The role of Emotions in the explanation of action. *European Review of Philosophy*, to appear.

Paulignan, Y., MacKenzie, C., Marteniuk, R., and jeannerod, M. 1991. Selective perturbation of visual input during prehension movements. I. The effects of changing object's position. *Experimental Brain Research,* 83: 502-512.

Peacocke, C. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.

Peacocke, C. 1998. Nonconceptual content defended. *Philosophy and Phenomenological Research*, LVIII, 2, 381-388.

Perry, J. 1993. *The Problem of the Essential Indexical*. Oxford: Oxford University Press.

Proust, J. (1999) Are there nonconceptual representations of actrion?. Manuscript.

Rosenbaum, D. A., and Jorgensen, M. J. 1992. Planning Macroscopic Aspects of Manual Control. *Human Movement Science* 11: 61-69.

Rosenbaum, D. A., Marchak, F., Barnes, H. J., Vaughan, J., Slotta, J. D., and Jorgensen, M. J. 1990. Constraints for Action Selection. Overhand versus Underhand Grips. *Attention and Performance XIII: Motor representation and Control*, ed. M. Jeannerod, 321-42. Hillsdale, NJ: Lawrence Erlbaum.

Scheerer, E. 1987. Muscle Sense and Innervation Feelings. A Chapter in the History of Perception and Action. *Perspectives on Perception and Action*, eds. H. Heuer and A. F. Sanders, 171-194. Hillsdale, N.-J.: Lawrence Erlbaum.

Searle, J. 1983. *Intentionality*. Cambridge: Cambridge University Press.

Searle, J. 1992. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.

Shiffrar, M., and Freyd, J. J. 1990. Apparent Motion of the Human Body. *Psychological Science* 1: 257-64.

Sirigu, A., Cohen, L., Duhamel, J. R., Pillon, B., Dubois, B., Agid, Y., and Pierrot-Deseiligny, C. 1995. Congruent unilateral impairments for real and imagined hand movements. *NeuroReport*, 6: 997-1001.

Viviani, P., and McCollum, G. 1983. The Relation between Linear Extent and Velocity in Drawing Movements. *Neuroscience* 10: 211-18.

NOTES

[1] The most prominent proponents of the Belief-Desire version of the Causal Theory are Davidson (1980) and Goldman (1970).

[2] It should be noted that although Davidson is known as the most famous advocate of the view that reasons are causes, his is not the strongest version of this claim. According to his thesis of the anomalism of the mental, there can be no empirical causal laws employing intentional vocabulary. But since he also maintains that two events related as cause and effect must fall under a strict causal law, it follows that the causal laws required to support the singular causal links between reasons and actions must advert to non-intentional characteristics of these events. Most causal theorists, however, are willing to give a stronger reading to the claim that reasons are causes. They contend that when some of an agent's beliefs and desires can legitimately be said to explain why the agent acted in a certain way, it must be the case that certain aspects of the content of those states were causally relevant in explaining why the agent acted as he did. In other words, contrary to Davidson, they think that the causal laws needed to support the singular causal links between reasons and actions must employ intentional vocabulary.

[3] Note, however, that this difficulty is not specific to the belief/desire versions of the causal theory but is also by versions of the causal theory such as Bratman's (1987) or Davidson's (1980: essay 5) that countenance states of intending as distinctive states not reducible to combinations of beliefs and desires.

[4] Other difficulties include the notorious problem of causal deviance as well as problems posed by the status of automatic or spontaneous actions.

[5] I don't want to claim that the success or failure of an action can always totally be explained and thus made intelligble by reference to the correctness or incorrectness of causally relevant representations of the agent. It may well be that some of the conditions necessary for the success of an action are not normally represented by the agent and that success or failure always depend in part on what Perry (1993: Essay 9) calls a certain benevolence on the part of Mother Nature. My contention here is simply that Israel, Perry, and Tutiya may leave too much to the benevolence of Mother Nature and that more than they think can be explained in terms of the agent's representations.

[6] See Pacherie (2000) for a more detailed discussion of the format and content of motor representations

[7] Neuroscientists often make a more liberal use of the term 'representation' than philosophers are wont to do, thus making the philosophers wary that what the neuroscientists call representations really qualify as mental representations. The philosophers' worries are of two kinds. First, one may wonder whether these so-called representations are really *representations*.. I will come back to this first worry in the next section. Second, one may wonder whether the neuroscientist's 'representations' qualify as *mental* representations. I won't launch here into a full-length discussion of the criteria for mental representations. Let me just offer one consideration that should allay this second worry. One very strong criterion of what can qualify as mental is the connection principle proposed by Searle (1992). According to this principle, no state qualifies as a mental state unless it is in principle accessible to consciousness. I personally think that this requirement is too strong. Although being accessible to consciousness may be taken as a sufficient condition for something to qualify as mental, I don't think this condition is necessary. I am willing to countenance mental states and mental representations that are essentially subpersonal. The point though is that is that Jeannerod's motor representations, although usually non-conscious, are in principle accessible to consciousness at least in part. Motor images are motor representations that have reached consciousness. The hedge I introduced — accessible to consciousness at least in part — reflects Jeannerod's idea that motor representations form a hierarchy, such that the aspects of an action encoded at higher levels of motor representations will be more easily accessible to consciousness than aspects encoded at lower levels. Executive representations therefore satisfy the connection principle and qualify as mental.

[8] Jeannerod (1977) distinguishes between two visual processing systems, what he calls the "what" system that derives semantic representations used for identification, categorization and recognition tasks and what he calls the "how" system that derives 'pragmatic' representations encoding information about objects used for visually guided action. Milner and Goodale (1993) proposed a general distinction between two visual processing systems very similar to Jeannerod's, though different in detail. Their account is further elaborated in Milner and Goodale (1995) and Milner (1997).

[9] It should be noted that Campbell discusses causal indexicality at the level of linguistic predicates, whereas the format of Jeannerod's representations is sensory-motor and independent of capacities for verbalization. Campbell points out that many notions are causally significant insofar as judgements made using them have some significance for the ways in which the world will behave, and for how it would behave in various possible circumstances. A subclass of those notions has the further characteristic that grasp of their causal significance consists in one's practical grasp of their immediate implications for one's own actions. Notions in this subclass are what Campbell calls causally indexical notions. Predicates such as 'is a weight I can easily lift', 'is too hot for me to handle' or 'is within my reach' are offered by Campbell as examples of causally indexical predicates. He

notes, however, that although these examples make use of the first person and use notions of weight and temperature, use of indexical terms need not depend upon self-consciousness or grasp of non-indexical notions. Unstructured uses of, say, 'is heavy', 'is hot' or 'is within reach' may be taken as more primitive examples of causally indexical terms insofar as they have immediate implications for the subject's actions. What I want to suggest here is that the notion of causal indexicality has application not just at the level of linguistic predicates or concepts but also at the level of nonconceptual content.

[10] Gibson's theory of affordances is taken by many as controversial. We should distinguish, however, between two different aspects in Gibson's theory. When I say that Jeannerod's pragmatic representations are akin to Gibson's affordances, what I am interested in is Gibson's notion of an affordance as a property of an object determined jointly by the physical attributes of the objects and by the sensory and motor capacities of an organism. What is controversial in Gibson's theory is his idea that affordances are directly perceived, that the nervous system somehow resonates to them and that here is no need for computation and information processing to detect those invariants. Jeannerod most certainly doesn't share this latter view and neither do I.

[11] The description that follows draws heavily on Jeannerod (1997).

[12] Evidence that this speed accuracy trade off is indeed reflected in the content of motor representations comes from experiments on mental simulation of actions, where the same relation between speed and difficulty can be observed (Decety, 1991; Decety and Jeannerod, 1996; Sirigu *et al.*, 1995).

[13] It may be suggested that at the level of protopropositional content (Peacocke's name for a higher-level of nonconceptual content), allocentric frames of reference are also used, yet the structure remains spatial.