

Intentions and Motor Representations: the Interface Challenge

Myrto Mylopoulos^{1,2} · Elisabeth Pacherie¹

© Springer Science+Business Media Dordrecht 2016

Abstract A full account of purposive action must appeal not only to propositional attitude states like beliefs, desires, and intentions, but also to motor representations, i.e., non-propositional states that are thought to represent, among other things, action outcomes as well as detailed kinematic features of bodily movements. This raises the puzzle of how it is that these two distinct types of state successfully coordinate. We examine this so-called “Interface Problem”. First, we clarify and expand on the nature and role of motor representations in explaining intentional action. Next, we characterize the respective functions of intentions and motor representations, the differences in representational format and content that these imply, and the interface challenge these differences in turn raise. We then evaluate Butterfill and Sinigaglia’s (2014) recent answer to this interface challenge, according to which intentions refer to action outcomes by way of demonstrative deference to motor representations. We present some worries for this proposal, arguing that, among other things, it implicitly presupposes a solution to the problem, and so cannot help to resolve it. Finally, we suggest that we may make some progress on this puzzle by positing a “content-preserving causal process” taking place between intentions and motor representations, and we offer a proposal for how this might work.

1 Introduction

Standard philosophical accounts of action appeal to intentions to explain the purposiveness of action. Intentions are conceived as mediating between beliefs and desires, on the one hand, and purposive behaviour on the other. By contrast, psychologists and

✉ Myrto Mylopoulos
myrto.mylopoulos@gmail.com

¹ Institut Jean Nicod, CNRS UMR 8129, Institut d’Etude de la Cognition, École Normale Supérieure & PSL Research University, Paris, France

² Department of Philosophy and Institute of Cognitive Science, Carleton University, Ottawa, Canada

neuroscientists working on motor cognition often focus on what happens downstream of intention, and emphasize the importance of motor representations in an explanation of purposive behaviour (Jeannerod 1997, 2006). They are primarily interested in understanding how such states interact with sensory states in order to allow for the smooth execution of action.

A promising conciliatory approach that has become popular recently (e.g., Nanay 2013), has it that a full explanation of the purposiveness of action requires an appeal to both intentions, as on the standard account, and motor representations, as is typical in the science of action. However, this approach is not without its difficulties. In a recent paper, Butterfill and Sinigaglia (2014) point out that the functional roles of intentions and motor representations are distinct but complementary, and that to play these distinct but complementary roles they must be harnessed together and pull in the same direction. The problem of explaining what ensures that they pull in the same direction is what they call “the interface problem”. This problem will be the focus of this paper.

Butterfill and Sinigaglia’s (2014) solution to the interface problem appeals to concepts deployed by intentions that refer to actions by deferring to motor representations. We concur with Butterfill and Sinigaglia and others in thinking that a complete explanation of the purposiveness of action must appeal to more than just intentions and other propositional attitudes. However, we are skeptical that the deferral solution that Butterfill and Sinigaglia propose to the specific interface problem they identify is successful. We think a different strategy is more promising.

In section 2, we address the question of why we need to appeal to motor representations in addition to beliefs, desires and intentions in order to explain the purposiveness of action and why this is philosophically important. In section 3, we characterize the respective functions of intentions and motor representations, the differences in representational format and content that these imply, and the interface challenge these differences in turn raise. In section 4, we discuss the solution to the interface problem proposed by Butterfill and Sinigaglia (2014) and explain why we think it fails. Finally, in section 5, we propose that a full solution to the interface problem should involve an appeal to motor schemas as the intermediary link between intentions and motor representations, as well as a special class of action concepts—what we call “executable action concepts”—and we sketch how such a solution might look.

2 Motor Representations and their Philosophical Relevance

In philosophical action theory it is customary to explain actions in terms of the beliefs, desires and intentions that motivate the agent to act as she does. Typically, a motivating complex is thought to include a pro-attitude toward actions with a certain outcome *O* (e.g., an intention that represents *O*), orienting beliefs of the agent regarding the circumstances she is in and instrumental beliefs regarding the effects of her actions in circumstances *C*. Thus, if an agent intends to bring about outcome *O*, believes she is in circumstances *C* and believes that performing action *A* in *C* will yield *O*, then she will *A*. Typically also, the success or failure of the action to bring about the intended outcome is thought to depend on the truth or falsity of the orienting and instrumental beliefs in the motivating complex. However, this

standard philosophical account of action explanation is faced with several difficulties. We will consider here only two of them.

The first difficulty is what Israel et al. (1993) have called “the problem of the wrong movement”. 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 is conceived of in the standard model. They illustrate this with the following example: Suppose Brutus intends 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. In other words, the truth of the beliefs figuring in the motivating complex as it is traditionally conceived does not guarantee that the bodily movement performed by the agent is appropriate.

The motivating complex as it is characterized in the standard account leaves a gap to be filled between the motivating cognitions and the action itself. When an agent’s orienting and instrumental beliefs are correct, what ultimately accounts for the success or failure of an intended action are the bodily movements actually performed and the effects they have. If we take it that the explanandum in a theory of action explanation is the action itself, not just the attempt or volition, we should be ready to include in the explanans representational states pertaining to movements and their effects. Motor representations are highly plausible candidates for such states.

A second problem the standard philosophical account is confronted with is the long-standing problem of antecedent causal deviance.¹ We can illustrate it with the following vignette adapted from Davis (1994).

The Marriage Proposal: Suppose John wants and intends to get down on his knees to propose marriage. Arriving in front of his beloved one, he is so overcome with emotion that he suddenly feels weak and sinks to his knees.

Here, as pointed out by Davis, John’s sinking to his knees is not an action even though it is caused by his desire and intention to get down on his knees. As the vignette illustrates, the existence of a causal relation between an intention representing a certain outcome and a corresponding outcome does not guarantee that the production of this outcome is a purposive action. The challenge is then to specify what causal connection must hold between the intention and the resultant bodily behavior for the latter to qualify as a purposive action and for the former to explain its purposiveness (see Shepherd 2014 for one such account).

Davidson (1980; Essay 3: 79) admitted despairing of spelling out the way in which an agent’s attitudes must cause action if they are to rationalize them. His despair, however, might have been premature, as an appeal to motor representations may help us with the problem of antecedent causal deviance. If a behavior can fail to qualify as purposive despite being caused by an intention with a seemingly appropriately related content, one may suspect that other representations must also contribute to explaining

¹ Brand (1984) draws a useful distinction between two problems of causal deviance, or, as he calls it, causal waywardness: antecedent and consequential. The first problem concerns the causal connection between the antecedent mental events and the initiation of bodily behavior; the second concerns the consequences of the activity once initiated. Our focus here is on the former.

the purposiveness of action. In many if not all instances of antecedent causal deviance, the problem may be diagnosed as arising from the fact that the causal chain leading from intention to behavior fails to include motor representations. Thus, causing “in the right way” may be a matter of intentions causing behavior via the instantiation of appropriate motor representations.

We have offered two reasons to think that a full account of purposive action must appeal not only to propositional attitude states like beliefs, desires, and intentions, but also to motor representations and that these representations are philosophically relevant. First, the standard philosophical account cannot fully explain what happens when we make action errors; motor representations may be needed to fill this explanatory gap. Second, motor representations may help us with the problem of antecedent causal deviance. But what exactly are motor representations and how do they differ from intentions?

3 Intentions Vs. Motor Representations

Butterfill and Sinigaglia (2014) argue that while both intentions and motor representations are representations of action outcomes, motor representations differ from intentions with respect to their representational format. Intentions have a propositional format while motor representations have a non-propositional, motoric format. This difference is, in their view, the source of an interface problem. The most straightforward way to arrive at the claim that they have different formats and therefore that an interface problem arises is, we think, by considering the different but complementary functions each of these types of states is thought to serve. While both motor representations and intentions integrate information from various sources, they differ with respect to the kind of information they integrate, the purpose of this integration, and hence also the constraints this integration obeys.

Intentions are standardly understood as propositional attitude states with a distinctive role in practical reasoning. In addition, in the last decades, a number of philosophers of action have developed dual theories of intention. For instance, Searle (1983) distinguishes between prior intentions and intentions-in-action, Bratman (1987) between future-directed and present-directed intentions, Brand (1984) between prospective and immediate intentions, Bach (1978) between intentions and executive representations, and Mele (1992) and Pacherie (2006, 2008) between distal and proximal intentions. Leaving aside slight differences of emphasis, this distinction is motivated by an analysis of their different and complementary functional roles, of the contents they must have to serve these functions and of their respective temporal scales.

Thus, Bratman (1987) stresses three functions of distal, or as he calls them future-directed, intentions: (i) they are terminators of practical reasoning about ends, (ii) they are prompters of practical reasoning about means and plans, and (iii) they have a coordinating function at both the intra-personal and the inter-personal level. In serving these three functions, intentions are subject to strong rationality constraints. Practical reasoning about ends should lead one to select the course of action one judges best given all the reasons that the agent takes to be relevant to this assessment, where these reasons include both the agent’s preferences and beliefs about the feasibility of various courses of action.

A proximal intention often inherits an action plan from a distal intention. Its function is then to anchor this plan in the situation of action. The temporal anchoring, the decision to start acting now is but one aspect of this process. Once the agent has established a perceptual information-link to the situation of action, she must ensure that the action plan is implemented in that situation. The job of proximal intentions is thus also to integrate conceptual information about the intended action as specified in the distal intention with relevant perceptual information about the current environment in order to yield a more specific situated representation of the action to be performed. Thus, in contrast to the content of distal intentions, the content of proximal intentions cannot be purely descriptive, it must be at least in part indexical and include pointers to elements of the environment.

The rationality constraints that bear on both distal and proximal intentions require the presence of a network of inferential relations among intentions, instrumental and perceptual beliefs, and desires. The existence of these inferential relations is what makes possible the internal and external consistency of an action plan, its consistency with the other activities of the agent or with the activities of others and finally its consistency with the situation of action. It is generally admitted that concepts are the inferentially relevant constituents of intentional states and that their sharing a common conceptual representational format is what makes possible a form of global consistency, at the personal level, of our desires, beliefs, intentions and other propositional attitudes. If we accept this common view, what follows is that for intentions to satisfy the rationality constraints they are subject to, they must have a propositional format and conceptual content.

The main functions associated with motor representations involve planning the bodily movements needed to perform an action and guiding and controlling their execution. Motor representations must therefore specify a motor program for an action and the exact spatial and temporal parameters of its constituent elements. This involves the multimodal functional integration of the external (e.g., visual, tactile, etc.) and internal (e.g. proprioceptive) sensory information needed to specify these parameters. The main characteristics of motor representations are as follows (for full descriptions, see Jeannerod 1997, 2006).

First, objects and situations are represented in terms of those of their properties that are immediately relevant for action, namely information about where relevant objects are located relative to the agent (e.g., position) and information concerning how to interact with these objects (e.g., size, shape, etc.). This is thought to involve a form of visual processing different from the visual processing supporting conscious visual perception. According to the influential Two Visual Systems Theory developed by Milner and Goodale (Milner and Goodale 1995; see also Ungerleider and Mishkin 1982, and the discussion in Jacob and Jeannerod 2003) there exist two visual systems, dedicated respectively to action and to the identification and recognition of objects and scenes. The vision for action system extracts from visual stimuli information about the properties of objects and situations that is immediately relevant for action, and uses this information to build motor representations used in effecting rapid visuo-motor transformations. Thus, the attributes of objects and situations are represented in a format useful for the specification of the relevant parameters of the selected motor program and motor representations involve egocentric frames of reference: centered on the agent's own body and effectors, structured by the agent's body schema and motor

skills. For instance, if one wants to grab an object, its spatial position will be represented in terms of the movements needed to reach for it and its shape and size in terms of the type of hand grip it affords.

Second, these representations of the movements to be effected reflect an implicit knowledge of the biomechanical constraints and the kinematic rules governing the motor system. For instance, bodily movements as represented in motor representations respect the isochrony principle (the tangential velocity of movements is scaled to their amplitude), Fitt's law (the time required to rapidly move to a target area is a function of the ratio between the distance to the target and the width of the target) or the two-third power law between curvature and velocity. Similarly, the movements of the effectors will be programmed so as to avoid awkward or uncomfortable limb positions and to minimize the time spent in extreme joint angles.

Third, a motor representation normally codes for transitive movements, where the goal of the action determines the global organization of the motor sequence. For instance, the type of grip chosen for a given object is a function not just of its intrinsic characteristics (its shape and size) but also of the subsequent use one wants to make of it. The same cup will be seized in different ways depending on whether one wants to carry it to one's lips or turn it upside down.

Motor representations also guide and control movements as they unfold. In order to do so, they must anticipate their effects but also allow for adjustments during execution. In recent decades, theories of motor control have emphasized the role of internal forward models. These models capture the causal relationships between actions and their sensory consequences and can be used by the motor system to estimate the effects of the motor commands sent to the effectors, compare these predicted effects with the desired effects, and make adjustments if needed. Actual feedback when it becomes available can also be used for corrections as well as to update the forward models (for full descriptions of these models, see Desmurget and Grafton 2000; Wolpert et al. 1995; Wolpert and Kawato, 1998).

As pointed out by Butterfill and Sinigaglia, it makes good computational sense that motor representations should represent action outcomes (i.e., the goal of the movements) rather than simply patterns of joint displacement or muscle activation, since doing so can considerably simplify the computations needed for motor planning and control. In addition, to ensure fluent bodily behavior and the smooth online corrections and adjustments of movements, these sensorimotor computations should be extremely rapid. Motor control, in other words, is subject to very tight temporal constraints. The time-scale at which motor representations must operate suggests that their functioning must be automatic and that they are typically sub-personal and thus not directly accessible to consciousness. For instance, several pointing experiments (e.g., Castiello et al. 1991; Goodale et al. 1986) have shown that in a task where subjects have to point to a target, they can do so accurately even on trials where the target is suddenly displaced by several degrees and they have to adjust their trajectories. Moreover, they can do so while remaining completely unaware of both the displacement of the target and their own corrections.

These considerations together suggest that motor representations must have a proprietary representational format, distinct from the format of intentions, and adapted to the functions they serve. Rather than being the inputs or the outputs of practical reasoning processes, they are the inputs and outputs of rapid sensorimotor

computations. Rather than being subject to norms of practical reasoning, they are subject to a set of biomechanical constraints and motor rules. Rather than being personal-level representations, they are subpersonal representations. Rather than functioning under conscious control, they function largely automatically.

Butterfill & Sinigaglia characterize the interface problem as the problem of explaining how it is that intentions and motor representations, having as they do different representational formats, are able to coordinate such that the action outcomes that they specify “non-accidentally match”. We typically take this match for granted. However, both mundane cases of action slips and pathological conditions, such as apraxia or the anarchic hand syndrome (AHS), illustrate the existence of an interface problem.

Anarchic hand syndrome is an especially striking example of a breakdown in interfacing between intentions and motor representations. Typically the result of lesion to SMA or anterior corpus callosum, AHS is a condition in which patients perform complex, goal-oriented movements with their cross-lesional limb that they feel unable to directly inhibit or control. The limb is often disproportionately reactive to environmental stimuli, carrying out habitual behaviors that are inappropriate to the context, e.g., grabbing food from a dinner companion’s plate (Della Sala 2005, 606). It is clear from many of the behaviors observed in these cases that the anarchic limb fails to “hook up” with the agent’s intentions. This is especially apparent in cases where the limb behaves at cross-purposes with those intentions, as in the following case described by Banks et al. (1989):

While playing checkers on one occasion, the left hand made a move he did not wish to make, and he corrected the move with the right hand; however, the left hand, to the patient’s frustration, repeated the false move. On other occasions, he turned the pages of the book with one hand while the other tried to close it; he shaved with the right hand while the left one unzipped his jacket; he tried to soap a washcloth while the left hand kept putting the soap back in the dish; and he tried to open a closet with the right hand while the left one closed it (457).

While anarchic hand syndrome indicates a general disruption in interfacing between the agent’s intentions and the motor representations that actually guide the movements of the anarchic limb, there are other pathological cases that suggest more localized disruptions, with respect to either the target that one intends to act upon or the means required for acting upon that target. Apraxia, typically the result of damage to parietal regions, is characterized broadly as “the consequence of a disruption of the normal mechanisms for action representations” (Jeannerod 2006, 12). Apraxic patients often display difficulties manipulating tools in the correct way. Sometimes these difficulties are reflected in failures to successfully execute a pantomime task. For example, one study found that apraxic patients were unable to pantomime the action of slicing bread, reflecting improper orientation of their limbs and impaired spatiotemporal organization (Clark et al. 1994; see also Ochipa et al. 1997 and De Renzi and Lucchelli 1988 for similar findings). And one recent study involved a patient who was unable to form correct hand postures for tool use, but was able to correctly identify the required hand postures in observational tasks, as well as verbally describe how to use the relevant tools without any difficulty (Hayakawa et al. 2015). This suggests that the deficit is not

one related to semantic knowledge. Such cases of apraxia can thus be seen as illustrating a failure to interface between intentions and motor representations with respect to the means for manipulating a tool.

As for more everyday breakdowns in interfacing, specifically related to the target of an action, we are all familiar with experiences like pouring orange juice rather than milk onto the cereal, or pressing ‘send’ on an email rather than ‘delete’. We will have more to say about the source of these “action slips” later in the discussion, but for now we merely flag that they indicate an inability to interface intentions with motor representations with respect to the particular target that is to be acted upon.

Still, despite these breakdowns, intentions and motor representations typically do manage to coordinate. But how? We turn now to evaluating Butterfill and Sinigaglia’s (2014) proposed solution to the interface problem.

4 Butterfill and Sinigaglia’s (2014) “Deferral Solution” to the Interface Problem

Recall that Butterfill and Sinigaglia characterize the interface problem as the problem of explaining how it is that intentions and motor representations, having as they do different representational formats, are able to coordinate such that the action outcomes that they specify “non-accidentally match”. In other words, the basic question here concerns how it is that intentions and motor representations are able to successfully align with one another. Before examining the specific solution that Butterfill and Sinigaglia propose, it will be useful to get a firmer grip on this notion of alignment.

Suppose an intention specifies action outcome A and a motor representation specifies action outcome B. On Butterfill and Sinigaglia’s view, action outcomes A and B align just in case, relative to a particular context, the occurrence of B either constitutes (at least partly) or causes (at least partly) the occurrence of A or vice versa. So, if A is <take a sip of water > and B is <reach>, we may say that these two outcomes align because reaching is partly constitutive of taking a sip of water. Likewise, if A is <ring the alarm > and B is <press the button>, then these two outcomes align insofar as the occurrence of a button-press is a partial cause of the occurrence of an alarm-ringing. In some more tightly aligned cases, the action outcome specified by A may just be the action outcome specified by the motor representation. So I may intend to reach for an object and also form a motor representation that specifies that I reach for an object. Given their different representational formats, as emphasized in the previous section, an explanation is called for as to how such alignment is achieved. This is the interface problem.

Butterfill and Sinigaglia consider three solutions to the interface problem, the third of which they endorse. These are (i) the common cause solution, (ii) the content-respecting causal processes solution, and (iii) the deferral solution. We shall look at each of these in turn, starting with the common cause solution.

The common cause solution proposes that intentions and motor representations coordinate in virtue of sharing a common cause that triggers them both. By analogy, imagine two healthy observers looking at the same red apple under standard lighting conditions, and thereby entering into the same type of visual state. We might ask ourselves how it is that these two visual states are aligned, that is, how it is that they

have matching content. And the answer will plainly not appeal to any causal interaction between those two states, but, rather, to their independent causal interactions with the same stimulus: the red apple. Similarly, the idea here is that a sensory state of the agent (e.g., a perception of a coffee mug) or an environmental stimulus (e.g., a coffee mug) triggers both an intention and a motor representation with aligned contents relating to the grasping of the mug. An advantage of this solution is that the difference in formats between these two representations does not raise any difficulties, since it is not in virtue of a causal interaction between them that they align.

However, Butterfill and Sinigaglia are skeptical that a full solution to the interface problem is available here. The reason they cite is that neither intentions nor motor representations are always triggered by environmental causes. Intentions are often the result of deliberation or planning, and motor representations are frequently keyed to intentions rather than stimuli in the environment or an agent's sensory states. But if so, then we must still find an explanation for the many remaining instances in which they manage to coordinate.

Another way that Butterfill and Sinigaglia propose the interface problem might be solved is via what they call "content-respecting causal processes" holding between intentions and motor representations. This would seem more promising as a comprehensive solution to the problem, given that content-preserving causal relations do, as a matter of fact, regularly obtain between intentions and motor representations. Butterfill and Sinigaglia maintain, however, that given that intentions and motor representations have a different representational format, any content-preserving causal process would require a translation process between the two states.

Butterfill and Sinigaglia find the current prospects for establishing the nature of this translation unacceptably weak. They write: "The difficulty is that nothing at all is known about this hypothetical translation between intention and motor representation, nor about how it might be achieved, nor even about how it might be investigated. Of course this doesn't show that we couldn't fully explain matching by appeal to content-respecting causal processes. But it does show that no such explanation is currently available" (133). In our view, Butterfill and Sinigaglia's dismissal here is too quick. In the next section, we will propose a content-preserving causal process solution to the interface problem that does not rely on a mysterious translation process.

Before getting to this, however, let us look at Butterfill and Sinigaglia's own solution to the interface problem. They write: "The key to solving this problem, we shall suggest, is to recognise that intentions can have constituents which refer to outcomes by deferring to motor representations of those outcomes" (120). In order to get clear on what they have in mind, let us examine the notion of deferral at play. Butterfill and Sinigaglia seem to have a view on which representation A defers to representation B just in case its reference is determined by the reference of B. They thus follow Levine (2010) in endorsing the idea of an intentionally-mediated mechanism, i.e., "a mechanism for securing reference that depends essentially on the content of another representation." (173). An ordinary example of such a mechanism at work might be the case of someone drawing a lottery prize winner and saying, "The person whose name appears on the piece of paper I am holding is the winner of the prize." In this case, the reference of 'the person whose name appears on the piece of paper I am holding' in the individual's speech act is determined by the reference of the name on the piece of paper.

Butterfill and Sinigaglia maintain that intentions refer to actions via demonstrative deference to motor representations. Demonstratives are a species of indexical, i.e., linguistic or mental items the reference of which shifts as a function of the context within which they occur. Unlike pure indexicals, like ‘I’, ‘here’, and ‘now’, for which context alone suffices to determine their reference, however, it is widely agreed that demonstratives require additional factors for this. When it comes to demonstratives deployed in speech acts, for example, many hold that a demonstrative gesture is required to fix the reference. So if I am at the market buying an apple, I may say to the merchant, “I’d like that one, please”, and physically point to the apple I intend as the referent of the demonstrative ‘that’ in my utterance.

To illustrate how demonstrative deferral between intentions and motor representations might operate, Butterfill and Sinigaglia provide the analogy of a visual map and a spoken utterance, each referring to the same route. The map refers to the route in the customary way, by depicting a line connecting two spatial points. The spoken utterance is simply “Follow that route!” Butterfill and Sinigaglia suggest that the utterance may refer to the route that the line on the map represents by deferring to that line by way of deploying the demonstrative ‘that’. Applying this picture to intentions and motor representations, the idea is that our intentions sometimes deploy demonstrative concepts that defer to motor representations specifying certain action outcomes, and thereby refer to those action outcomes, without any need for translation. So, on this proposal, we can consider the content of an intention to be “Do that!” and the demonstrative ‘that’ would defer to a motor representation referring to the relevant action. As Butterfill and Sinigaglia put it: “These demonstrative concepts would be concepts of actions not of motor representations, but they would succeed in being concepts of actions by deferring to motor representations. For any such concept, it is a motor representation which ultimately determines what it is a concept of” (134).

As we see it, there are at least four problems with this solution. We will now articulate each of these in turn.

4.1 The Lack of Example Issue

The first difficulty that Butterfill and Sinigaglia’s solution faces is that there are no clear examples of intention deferring to motor representation, which makes it dubious that this ever takes place, let alone that it very commonly does. Butterfill and Sinigaglia attempt to provide an example involving mental pantomime. They ask the reader to.

... consider purely mental pantomime—that is, phenomenologically action-like imagination. One might use this kind of imagination to explore different ways of completing a task and then, having hit on a good solution, think to oneself ‘Do that!’. It seems possible that in some such cases the demonstrative refers by deferring to a motor representation of action involved in imagining acting (135).

The obvious problem here is that we need not suppose that the intention refers to the relevant action by way of deferring to a motor representation. What seems more likely is that it does so by deferring to the mental image of the relevant action. In order to get around this issue, Butterfill and Sinigaglia ask us to “Contrast two cases of phenomenologically action-like imagination, both involving a tool. In the first case, imagine

grasping an object with the tool; and in the second case, imagine releasing the same object with the tool” (135) and further, to “[l]et the movements and muscle contractions involved in both cases be as similar as possible: let the difference between grasping and releasing the object be a matter of how the tool is configured rather than of how your body moves” (135). They suppose, in addition, that the action imagery is not of *intending* to grasp or *intending* to release a particular object, but of simply grasping or releasing it. Their argument seems to be that, since the imagery remains the same in both cases, the only way for the intention to refer to one action rather than the other is via deferral to the relevant motor representation.

But this example is problematic. For clearly there must be a phenomenological difference between the two cases with respect to the displacement of the object—in the first case it is picked up, and in the second case it is dropped. If these features of the action are not represented in the imagery, then these cannot be genuine cases of imagining grasping versus imagining releasing an object. So we do not see how this example helps to bolster Butterfill and Sinigaglia’s case, and since it is the only example they offer of intentions deferring to motor representations, we are left without a reason to think that this putative psychological phenomenon occurs.

4.2 The Conscious Accessibility Issue

Another issue here relates to the specific properties of the demonstratives that are thought by Butterfill and Sinigaglia to underlie this special brand of deferral. In general, demonstratives require a perceptual-attentional link to their referents. In this case, as Butterfill and Sinigaglia have it, the referents of the demonstratives deployed in intentions are actions, but these references are secured by deferral to the motor representation. And, if we are to take seriously the analogy with the utterance and the visual map, it seems that the perceptual-attentional link enters in at this stage: the reference of the demonstrative is specified by establishing a perceptual-attentional link with the representation that determines it. But in the case of intention, there is no such perceptual-attentional link with the motor representation; these are not mental representations that we can attend to or perceive.

It is important to stress that we are not claiming here that no mental representations can be attended to or perceived. It is arguable that one sometimes attends directly to one’s conscious experiences. And if perceptual models of introspection or higher-order awareness (e.g., Lycan 1996) are correct, then one might also sometimes perceive such states. But whatever states one attends to or perceives must be ones that one is capable of being aware of in some way, and motor representations do not meet this criterion. They are subpersonal states that are inaccessible to awareness, thus precluding them from being states that we can perceive or attend to, as is required if intentions are to demonstratively refer to them.

4.3 The Error/Fallibility Issue

Yet another pressing issue for Butterfill and Sinigaglia’s view is that it seems not to have the resources to explain errors in action execution. On their proposal, presumably such errors would come about when the demonstrative defers to the wrong motor representation. In the case of demonstratives deployed in speech acts, some hold that

they can go wrong in virtue of the accompanying demonstration going awry. For example, undergoing a sudden spell of dizziness, in attempting to refer to the green apple in the heap as “That apple” by way of a pointing gesture, one may incorrectly point to a red apple. To accommodate such cases, some theorists appeal to directing intentions, i.e., the intentions that govern our demonstrative utterances, in order to determine the correct reference of such utterances. On this picture we can say that one successfully refers to the green apple in the heap with one’s utterance of “That apple”, even though one points to the red one, because one intends for one’s utterance to refer to the green one.

Unfortunately, though, no such resources would be available in the case of demonstratives deployed in the content of intentions. For one, it is unclear what it would be for such a demonstrative to misfire, i.e., to “point” to the wrong motor representation, as there is of course no actual pointing gesture taking place and, as we noted earlier, no attentional or perceptual link to a motor representation. In addition, the strategy of positing directing intentions, i.e., intentions to intend to defer to a particular motor representation, will not help here, since these intentions would also need to pick out the appropriate motor representation and so we’d be faced with the same issue. Matters are made even worse if one holds, as some do (e.g., Levine 2010) that mental demonstratives in general simply *cannot* go wrong due to these very reasons, for then the view that Butterfill and Sinigaglia put forward entails that we never make action execution errors as the result of engaging the wrong motor representations, which is clearly false.

4.4 The Selection Issue

This last issue brings us to the doorstep of what is the most serious difficulty for Butterfill and Sinigaglia’s view: it presupposes a translation process, the very thing that they were trying to avoid. In the case of a regular demonstrative utterance, the speaker has an independent way of picking out which object to refer to. If I want to demonstratively refer to the green apple in the heap, I will employ my perceptual resources in order to determine which item to point to. Similarly, in the case of demonstrative deferral in intention, the agent must have an independent grasp of which motor representation is the appropriate one to select via such deferral. But this would require a way of translating between the intention and the motor representation being picked out, in order to establish which motor representation correctly corresponds, and Butterfill and Sinigaglia have already argued that we know nothing about how this translation process would work.

Given these four difficulties, we think that Butterfill and Sinigaglia’s deferral solution to the interface problem cannot work. We turn now to sketch what we take to be a more promising strategy moving forward.

5 Towards a Solution to the Interface Problem

The solution we are proposing is a version of the content-preserving causal process approach that, as we shall show, does not require a translation process. It appeals to the notions of executable action concepts and motor schemas. We first introduce the notion of an executable action concept. We then argue that to have an executable concept for a

given type of action one must have a motor schema for actions of that type and explain how such motor schemas could be acquired. Finally, we outline the role of attention in helping to select the information needed to specify motor schemas so as to yield motor representations.

We must start by distinguishing between two different types of action concept. First, there are those that do not hook up with corresponding motor representations. For instance, most of us have a concept of ‘tail wagging’ that we can deploy when we judge, for instance, that Julius the dog is wagging his tail. Or if you are not convinced that tail wagging constitutes purposive behavior, consider the action concept ‘tail swinging’, as in ‘cows constantly swing their tails to flick away flies’. Why we do not have motor representations corresponding to these concepts of ‘tail wagging’ or ‘tail swinging’ is obvious enough not to need spelling out. But even limiting ourselves to human actions, examples of action concepts that are not linked to motor representations abound. Some humans can voluntarily wiggle their ears, others cannot. Yet, even if you belong to the latter category, you can still have a concept of ‘ear wiggling’. In addition, we also have concepts for actions we do not know how to perform, although we could in principle learn (e.g., cartwheeling). Such action concepts are observational, in that they are formed on the basis of third-personal perceptual representations, i.e., visual representations of others performing such actions. These concepts, like observational concepts more generally, allow us to accurately categorize instances of action that fall under them.

But a second class of action concepts does more than this. These concepts allow us to activate appropriate motor representations, in part because they are formed on the basis of such representations, rather than third-personal perceptual information. We call these *executable action concepts* (Pacherie 2011). If one has an executable action concept, then one has the ability to perform the action in question. Like Butterfill and Sinigaglia, we agree that the key to solving the interface problem is to recognize that intentions can have as constituents specific types of action concepts. We disagree, however, that these should be identified with demonstrative deferential concepts. Rather, we propose that in order to properly interface with motor representations, proximal intentions must have as constituents executable action concepts.²

The empirical evidence for a bidirectional link between the processing of linguistic items pertaining to action concepts and the activation of motor representations is well documented. First, passively reading action verbs has been found to somatotopically activate areas of the motor and premotor cortex associated with the relevant body parts needed to carry out the specified actions (Hauk et al. 2004). For example, there is evidence that reading the word ‘kick’ activates the dorsal part of the motor cortex, whereas reading the words ‘pick’ and ‘lick’ activate lateral and ventral regions. This pattern of activation overlaps significantly with the actual activation that takes place when carrying out these actions with the relevant effectors of foot, hand, and mouth respectively. Second, stimulation of the motor system has been found to affect the linguistic processing of action concepts. For instance, one study found that applying

² We leave it open here whether the action concepts that feature in the contents of distal intentions should also be executable action concepts or perhaps superordinate executable action concepts. For instance, one might argue that Frédérique can rationally form the distal intention to cartwheel around her house to celebrate her next birthday, even though at the moment she is forming this intention she has no executable concept of cartwheeling, provided that it is part of her plan to acquire this motor skill before her birthday.

TMS to hand and foot areas of the motor cortex improved the recognition of hand-related ('pick') and foot-related ('kick') action verbs respectively in lexical decision tasks (Pulvermüller et al. 2005; see also Kiefer and Pulvermüller 2012). Another study found that participants who were right-hand dominant were faster at lexical decision tasks involving manual action-verbs following stimulation of the left premotor cortex vs. stimulation of the right premotor cortex. The same effect was not found, however, for participants performing lexical decision tasks involving nonmanual verbs, or concrete vs. nonconcrete nouns (Willems et al. 2011).

But while this is evidence that action concepts and motor representations are indeed "hooked up", this leaves us with a gaping question: how is it that they get connected in the first place? And how do we come to acquire action concepts?

The solution we outline appeals to the notion of a motor schema. Motor schemas are more abstract and stable representations of actions than motor representations. They form a bridge between such representations and the executable action concepts that are deployed in the contents of proximal intentions. Our view is that to have an executable concept for a given type of action one must have a motor schema for actions of that type. While some elementary motor schemas (motor primitives) may be innate, motor schemas can also be acquired in a bottom-up fashion through a process of inductive generalization from sets of motor representations or from sets of already extant motor schemas.

According to schema theory (Schmidt 1975, 2003; Arbib 1981, 2003; Jeannerod 1997) motor schemas are both repositories of information and control structures. They are internal models or stored representations that represent generic knowledge about a certain pattern of action and are implicated in the production and control of action. For instance, in the influential Motor Schema Theory proposed by Richard Schmidt (1975, 2003), a motor schema involves a generalized motor program, together with corresponding 'recall' and 'recognition' schemas. The generalized motor program is thought to contain an abstract representation defining the general form or pattern of an action, that is the organization and structure common to a set of motor acts (e.g., invariant features pertaining to the order of events, their spatial configuration, their relative timing and the relative force with which they are produced). This generalized motor program has parameters that control it. In order to determine how an action should be performed on a given occasion, parameter values adapted to the situation must be specified. Thus, a motor schema also includes a rule or system of rules describing the relationships between initial conditions, parameter values and outcomes and allowing us to perform the action over a large range of conditions (the 'recall schema' in Schmidt's terminology). Finally, the motor schema also includes a rule or system of rules describing the relationships between initial conditions, exteroceptive and proprioceptive sensory feedback during an action, and action outcome (a recognition schema), allowing agents to know when they have made an error – i.e., the action does not have the sensory consequences it is expected to have – and to correct for it. Where Schmidt talks of 'recall' and 'recognition' schemas, more recent versions of schema theory often talk of internal inverse and forward models.

As we understand motor schema theory, the motor representations that guide specific action are instantiations of motor schemas where the values of the parameters that control the action are specified and then updated depending on sensory information and feedback. Motor schemas are thus more abstract and enduring representations of

actions that store knowledge about the invariant aspects and the general form of an action.³ How then can motor schemas be learned in a bottom-up fashion? Here we tentatively suggest that motor schema learning might be a form of Bayesian learning. We emphasize that our account is not committed to a Bayesian framework. We are merely suggesting one possible model of how the details of the motor schema acquisition process might be worked out. We note that though the Bayesian approach has many attractive features, which we discuss in what follows, more evidence and theoretical refinement are still needed before it can be adopted with any certainty.

Bayesian models of cognitive development and learning (Perfors et al. 2011; Tenenbaum et al. 2011; Hohwy 2013) provide a general purpose computational framework for understanding how a learner might make the inductive generalizations involved in learning structure and grasping causal relations, explaining them as forms of Bayesian or probabilistic inference. In a nutshell, the mathematics of Bayesian statistics allow one to compute which of a set of hypotheses (an hypothesis space) best explains observed data. The central claim of these models is that human minds learn and reason according to Bayesian principles. In recent years, evidence has accumulated that Bayesian inference and learning models can serve as powerful and biologically plausible models and successfully explain a range of phenomena including perceptual learning (Orbán et al. 2008; Maloney and Mamassian 2009) and causal learning in adults and infants (Tenenbaum and Griffiths 2001; Gopnik et al. 2004; Lucas et al. 2014). This Bayesian framework has also been applied to motor learning, contributing to important progress in our understanding of the computational aspects of such learning (e.g., Wolpert et al. 2011; Orbán and Wolpert 2011). In particular, there is evidence that parameter learning and structural learning in motor tasks can be best described as forms of Bayesian inference and modelling (Braun et al. 2009; Braun et al. 2010).

Conceiving of motor learning as a form of Bayesian probabilistic learning offers several advantages. Firstly, acquiring a motor schema for a certain type of action involves learning the general form or structure of the action, learning mapping rules between initial conditions, outcomes and parameter values (corresponding to inverse models) and learning mapping rules between these and exteroceptive and proprioceptive sensory consequences (predictive or forward models).⁴ The Bayesian approach offers useful tools for explaining how prototypes and causal relations can be extracted from data and thus for explaining both how the general form or structure of an action type and internal inverse and forward models of actions of that type (i.e., the mappings between sensory and motor variables) could be learned. In addition, hierarchical Bayesian modelling, by positing not just a single level of hypotheses to explain the data but multiple levels, allows actions to be described in more and more abstract terms and thus makes it possible to capture commonalities between sets of basic action schemas. It can also explain how novel schemas can be acquired not just from scratch,

³ Note that this suggests an alternative view of motor imagery. If the general form of an action is stored at the level of motor schemas, imagining performing an action may not involve, contra Butterfill and Sinigaglia's view, forming a fully specified motor representation of an action but rather activating the corresponding action schema without specifying parameters at all or while providing only rough estimates of these parameters rather than the values that would be needed for successful action (see, e.g., Arbib 2008).

⁴ Though we include mention of them here, we note that it is currently a matter of some controversy whether inverse models should be included in Bayesian accounts (see, e.g., Friston 2011).

by trial-and-error, but also by combining extant schemas or by abstracting from sets of extant schemas.

Secondly, according to the Bayesian approach, motor schemas are acquired through processes of probabilistic inductive generalization. Bayesian learning is thus, *prima facie*, consistent with a bottom-up approach to the interface problem. The Bayesian approach also maintains that these inductive processes are constrained to some degree by priors. It has been proposed that in the motor domain inductive generalization is constrained by motor primitives (Wolpert et al. 2011). As these authors explain, “motor primitives can be thought of as neural control modules that can be flexibly combined to generate a large repertoire of behaviours” (745). The set of motor primitives would constrain the learning process, for instance, by making a new action for which the motor system has many primitives easier to learn than a new action that cannot straightforwardly be approximated by sets of motor primitives.⁵ It may be thought that in acknowledging the role of such constraints on the acquisition of motor schemas and hence of executable action concepts, we are introducing a top-down element in our story. This, however, does not make our solution to the interface problem a top-down solution in the sense that Butterfill and Sinigaglia’s solution is. Their solution is indeed top-down insofar as they view the relevant action concepts as demonstrative deferential concepts and insofar as demonstrative reference fixing typically requires a top-down perceptual-attentional link to the intended referent. In contrast, we postulate no such perceptual-attentional link; rather we simply acknowledge the existence of built-in constraints on motor learning, constraints that may apply automatically and quite independently of top-down attentional processes.

Thirdly, the Bayesian approach to the acquisition of motor schemas we propose has the resources needed to explain errors in action execution, resources that, as we noted in section 4, Butterfill and Sinigaglia’s view appears to lack. Bayesian learning is probabilistic: acquiring a motor schema involves learning probability distributions on variables at the level below (e.g. learning that, for a given state of the world and a given state of the agent, a given motor command leads to a certain outcome with higher probability than to others). Thus, for most if not all actions, an intention with a certain executable action concept as part of its content triggering an appropriate motor schema is no absolute guarantee that action execution will be successful, it only guarantees a reasonably high probability of success. For instance an archer with extensive practice shooting at a target from a distance of ten yards may have acquired the motor schema and the mappings allowing her to hit the inner ring of the target nearly 100 % of the time. Yet, on the Bayesian approach, learning may be partial in the sense that the agent’s probability distributions remain coarse-grained or that she has acquired relevant mappings only for a certain range of circumstances rather than the full range of circumstances in which the action could be performed. For instance, our archer might have acquired very precise mappings for shooting targets 10 yards away explaining her near 100 % success rate at this distance, less precise mappings for targets 20 yards away and a correspondingly lower rate of success, and lack mappings altogether for targets 50 yards away, leaving her performance at chance level.

⁵ These constraints set by motor primitives on motor learning may be considered a form of what Clark (2013) calls systemic priors, i.e., priors built-in the motor system rather than empirical priors.

Finally, error in action execution might have yet another source. Action execution demands not just that an agent have acquired an executable action concept and be in a position to activate an appropriate motor schema with mappings defined over a range of situations that encompasses the situation at hand; it also demands that the information needed to set the value of the schema's parameters be selected and encoded in a format readily exploitable by the motor system. For instance, the parameters that control a hand action may include the size, shape, and location of the object to be grasped. Incorrect specification of this information may also lead to action failure. This means that a complete solution to the interface challenge also demands an explanation of how we select the object or objects we intend to act upon and the properties of these target objects relevant for setting the parameters of motor schemas. We now turn to this issue.

Here we can appeal to the role of attention in helping the motor system fill in the parameters of action schemas. In particular, as others have discussed (e.g., Campbell 2002; Wu 2015), attention plays a crucial role in helping us act upon targets in our environment in appropriate ways. First, often guided by an agent's intention, it identifies the object or objects to be acted upon from among competing objects. But in addition to this basic target identification, attention also isolates the relevant visual properties of the target(s) that are then used to help program an appropriate motor representation. To illustrate, consider the situation of an agent who is confronted with an array of objects: a coffee mug, a fork, and a plate. Not only will attention select for which object to act upon, depending on whether the agent intends to drink coffee, eat food, or pick up a plate, but it will also select for visual properties of the target that are relevant for informing the appropriate motor response, such as its size, structure, and orientation. As Wu (2015) puts it, "What attention is providing to the agent is the relevant target, but this involves providing content to systems that process that target in a way that sets parameters for and programs an appropriate response" (13).

The role of attention in interfacing comes out especially clearly when one considers certain types of action error. One such type, mentioned earlier, is that which Reason (1990) identifies as an "absent-minded slip", and is characterized by failures of execution arising in habitual or highly-learned action sequences. Reason points out that such errors occur when limited attentional resources are "captured" by irrelevant stimuli, thus substituting the wrong object for the correct one in the appropriate action sequence. A significant causal factor in pouring the orange juice onto the cereal rather than the milk is that this is the object to which one (absent-mindedly) attends, and so irrelevant information is processed and the parameters of the motor schema are incorrectly filled in. In addition, because habitual actions are not directly guided by intention each step of the way, but unfold automatically, attentional resources are more likely to get incorrectly allocated, and may be more strongly influenced by environmental factors.

6 Conclusion

In this article, we have argued, in agreement with Butterfill and Sinigaglia, that a philosophical theory of action cannot provide a full account of purposive action without appealing to motor representations and without explaining how intentions interface with motor representations. We have also pointed out, however, important difficulties

with their deferral solution to the interface problem. We have tried to explore instead a version of the content-respecting causal processes solution, an approach that, in our view, they too quickly dismiss as relying on a mysterious translation process. Our solution appeals to the notions of executable action concepts and motor schemas. We propose that in order to properly interface with motor representations, proximal intentions must have as constituents executable action concepts, where to have an executable concept for a given type of action one must have a motor schema for actions of that type. Motor schemas are more abstract and enduring representations than motor representations. They store knowledge about the invariant aspects and the general form of an action and are implicated in the production and control of action. On the one hand, they can be acquired through processes of probabilistic inductive generalization from motor representations or from already extant schemas. On the other hand, the activation of a motor schema once learned will yield a motor representation, when the information needed to specify its parameters is provided, typically via attentional processes. Motor schemas are thus, we submit, what bridges the gap between intentions and motor representations, ensuring proper, content-preserving coordination without requiring any mysterious translation process.

Acknowledgments Elisabeth Pacherie's work was supported by grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL*.

Myrto Mylopoulos would like to thank audiences at the Institut Jean Nicod PaCS workshop and the Carleton University philosophy colloquium for useful discussion. Elisabeth Pacherie would like to thank the audience at the Spring School on Action in Tübingen.

Both authors are grateful to Daniel Burnston for helpful comments on an earlier draft.

References

- Arbib, M.A. 1981. Perceptual structures and distributed motor control. In *Handbook of Physiology – The Nervous System II*, ed. V.B. Brooks, 1449–1480. American Physiological Society: Motor Control.
- Arbib, M. A. 2003. Schema theory. *The Handbook of Brain Theory and Neural Networks* (second ed.), MIT Press, Cambridge, MA, pp. 993–998.
- Arbib, M.A. 2008. From grasp to language: embodied concepts and the challenge of abstraction. *Journal of Physiology, Paris* 102(1): 4–20.
- Bach, K. 1978. A representational theory of action. *Philosophical Studies* 34(4): 361–379.
- Banks, G., P. Short, J. Martinez, R. Latchaw, G. Ratcliff, and F. Boller. 1989. The alien hand syndrome: clinical and postmortem findings. *Archives of Neurology* 46: 456–459.
- Brand, M. 1984. *Intending and acting*. Cambridge, MA: MIT Press.
- Bratman, M. 1987. *Intention, plans, and practical reason*. Cambridge, MA: Cambridge University Press.
- Braun, D.A., A. Aertsen, D.M. Wolpert, and C. Mehring. 2009. Motor task variation induces structural learning. *Current Biology* 19(4): 352–357.
- Braun, D.A., C. Mehring, and D.M. Wolpert. 2010. Structure learning in action. *Behavioural Brain Research* 206(2): 157–165.
- Butterfill, S., and C. Sinigaglia. 2014. Intention and motor representation in purposive action. *Philosophy and Phenomenological Research* 88(1): 119–145.
- Campbell, J. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Castiello, U., Y. Paulignan, and M. Jeannerod. 1991. Temporal dissociation of motor responses and subjective awareness. A study in normal subjects. *Brain* 114(Pt 6): 2639–2655.
- Clark, A. 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36(03): 181–204.

- Clark, M.A., A.S. Merians, A. Kothari, H. Poizner, B. Macauley, R.L.J. Gonzalez, and K.M. Heilman. 1994. Spatial planning deficits in limb apraxia. *Brain* 117: 1093–1106.
- Davidson, D. 1980. *Essays on actions and events*. Oxford: Oxford University Press.
- Davis, L.H. 1994. Action. In *A companion to the philosophy of mind*, ed. S. Guttenplan, 111–117. Oxford: Blackwell.
- De Renzi, E., and F. Lucchelli. 1988. Ideational apraxia. *Brain* 111(5): 1173–1185.
- Della Sala, S. 2005. The anarchic hand. *The Psychologist* 18(10): 606–609.
- Desmurget, M., and S. Grafton. 2000. Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Science* 4: 423–431.
- Friston, K. 2011. What is optimal about motor control? *Neuron* 72(3): 488–498.
- Goodale, M.A., D. Péllisson, and C. Prablanc. 1986. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature* 320: 748–750.
- Gopnik, A., C. Glymour, D.M. Sobel, L.E. Schulz, T. Kushnir, and D. Danks. 2004. A theory of causal learning in children: causal maps and Bayes nets. *Psychological Review* 111(1): 3.
- Hauk, O., I. Johnsrude, and F. Pulvermüller. 2004. Somatotopic representation of action words in human motor and premotor cortex. *Neuron* 41(2): 301–307.
- Hayakawa, Y., T. Fujii, A. Yamadori, K. Meguro, and K. Suzuki. 2015. A case with apraxia of tool use: selective inability to form a hand posture for a tool. *Brain and Nerve* 67(3): 311–316.
- Hohwy, J. 2013. *The predictive mind*. Oxford University Press.
- Israel, D., J. Perry, and S. Tutiya. 1993. Executions, motivations and accomplishments. *The Philosophical Review* 102: 515–540.
- Jacob, P., and M. Jeannerod. 2003. *Ways of Seeing, the Scope and Limits of Visual Cognition*. Oxford: Oxford University Press.
- Jeannerod, M. 1997. *The cognitive neuroscience of action*. Oxford, UK: Blackwell Publishers, Inc..
- . 2006. *Motor Cognition: What actions tell the self*. New York, NY: Oxford University Press.
- Kiefer, M., and F. Pulvermüller. 2012. Conceptual representations in mind and brain: theoretical developments, current evidence and future directions. *Cortex* 48: 805–825.
- Levine, J. 2010. Demonstrative thought. *Mind & Language* 25(2): 169–195.
- Lucas, C.G., S. Bridgers, T.L. Griffiths, and A. Gopnik. 2014. When children are better (or at least more open-minded) learners than adults: developmental differences in learning the forms of causal relationships. *Cognition* 131(2): 284–299.
- Lycan, W.G. 1996. *Consciousness and experience*. Cambridge, MA: Bradford Books/MIT Press.
- Maloney, L.T., and P. Mamassian. 2009. Bayesian decision theory as a model of human visual perception: testing Bayesian transfer. *Visual Neuroscience* 26: 147–155.
- Mele, A. 1992. *Springs of action*. New York: Oxford University Press.
- Milner, A.D., and M.A. Goodale. 1995. *The visual brain in action*. Oxford: Oxford University Press.
- Nanay, B. 2013. *Between perception and action*. Oxford: Oxford University Press.
- Ochipa, C., S.Z. Rapcsak, L.M. Maher, L.J.G. Rothi, D. Bowers, and K.M. Heilman. 1997. Selective deficit of praxic imagery in ideomotor apraxia. *Neurology* 49: 474–480.
- Orbán, G., and D.M. Wolpert. 2011. Representations of uncertainty in sensorimotor control. *Current Opinion in Neurobiology* 21(4): 629–635.
- Orbán, G., J. Fiser, R.N. Aslin, and M. Lengyel. 2008. Bayesian learning of visual chunks by human observers. *Proceedings of the National Academy of Sciences* 105(7): 2745–2750.
- Pacherie, E. 2006. Towards a dynamic theory of intentions. In *Does Consciousness Cause Behavior? An Investigation of the Nature of Volition*, eds. S. Pockett, W.P. Banks, and S. Gallagher, 145–167. Cambridge, MA: MIT Press.
- Pacherie, E. 2008. The phenomenology of action: A conceptual framework. *Cognition* 107: 179–217.
- Pacherie, E. 2011. Nonconceptual representations for action and the limits of intentional control. *Social Psychology* 42(1): 67–73.
- Perfors, A., J.B. Tenenbaum, T.L. Griffiths, and F. Xu. 2011. A tutorial introduction to Bayesian models of cognitive development. *Cognition* 120(3): 302–321.
- Pulvermüller, F., O. Hauk, V.V. Nikulin, and R.J. Ilmoniemi. 2005. Functional links between motor language systems. *European Journal of Neuroscience* 21: 793–797.
- Reason, J. 1990. *Human error*. Cambridge, MA: Cambridge University Press.
- Schmidt, R.A. 1975. A schema theory of discrete motor skill learning. *Psychological Review* 82(4): 225.
- . 2003. Motor schema theory after 27 years: reflections and implications for a new theory. *Research Quarterly for Exercise and Sport* 74(4): 366–375.
- Searle, J.R. 1983. *Intentionality: An essay in the philosophy of mind*. Cambridge, MA: Cambridge University Press.

- Shepherd, J. 2014. The contours of control. *Philosophical Studies* 170(3): 395–411.
- Tenenbaum, J.B., and T.L. Griffiths. 2001. Structure learning in human causal induction. In *Advances in neural information processing systems*, vol 13, eds. T. Leen, T. Dietterich, and V. Tresp, 59–65. Cambridge, MA: MIT Press.
- Tenenbaum, J.B., C. Kemp, T.L. Griffiths, and N.D. Goodman. 2011. How to grow a mind: statistics, structure, and abstraction. *Science* 331(6022): 1279–1285.
- Ungerleider, L.G., and M. Mishkin. 1982. Two cortical visual systems. In *Analysis of visual behavior*, eds. D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield, 549–586. Cambridge, MA: MIT Press.
- Willems, R.M., L. Labruna, M. D’Esposito, R. Ivry, and D. Casasanto. 2011. A functional rôle for the motor system in language understanding: evidence from theta-burst transcranial magnetic stimulation. *Psychological Science* 22: 849–854.
- Wolpert, D.M., Z. Ghahramani, and M.I. Jordan. 1995. An internal model for sensorimotor integration. *Science* 269(5232): 1880–1882.
- Wolpert, D.M., R.C. Miall, and M. Kawato. 1998. Internal models in the cerebellum. *Trends in Cognitive Sciences* 2(9): 338–347.
- Wolpert, D.M., J. Diedrichsen, and J.R. Flanagan. 2011. Principles of sensorimotor learning. *Nature Reviews Neuroscience* 12(12): 739–751.
- Wu, W. (2015). Experts and deviants: the story of agentive control. *Philosophy and Phenomenological Research*, 90(3), online first.