

# Nowhere and Everywhere: The Causal Origin of Voluntary Action

Aaron Schurger · Sebo Uithol

© Springer Science+Business Media Dordrecht 2015

**Abstract** The idea that intentions make the difference between voluntary and non-voluntary behaviors is simple and intuitive. At the same time, we lack an understanding of how voluntary actions actually come about, and the unquestioned appeal to intentions as discrete causes of actions offers little if anything in the way of an answer. We cite evidence suggesting that the origin of actions varies depending on context and effector, and argue that actions emerge from a causal web in the brain, rather than a central origin of intentional action. We argue that this causal web need not be confined to the central nervous system, and that proprioceptive feedback might play a counter-intuitive role in the decision process. Finally we argue that the complex and dynamic origins of voluntary action and their interplay with the brain's propensity to predict the immediate future are better studied using a dynamical systems approach.

## 1 Introduction

More than 20 years ago, Dennett and Kinsbourne (1992) convincingly argued that there is no single place in the brain “where it all comes together” for processing by mechanisms of conscious perception. We argue that the same is also true for action initiation: There is no region or network in the brain “where it all comes from”—from which all decisions-to-act originate. Although most philosophers and scientists might readily agree, the idea of a “seat of volition” stems from an insidious line of thinking that is much harder to abandon. This reasoning is based on the intuition that voluntary

---

Aaron Schurger and Sebo Uithol contributed equally to this work.

A. Schurger (✉)

Laboratory of Cognitive Neuroscience, Brain-Mind Institute, Department of Life Sciences, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland  
e-mail: aaron.schurger@gmail.com

A. Schurger

Center for Neuroprosthetics, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

S. Uithol (✉)

Department of Neuroscience, University of Parma, Parma, Italy  
e-mail: sebo.uithol@gmail.com

actions are based on intentions, and that intentions are therefore somewhere represented in our brains.

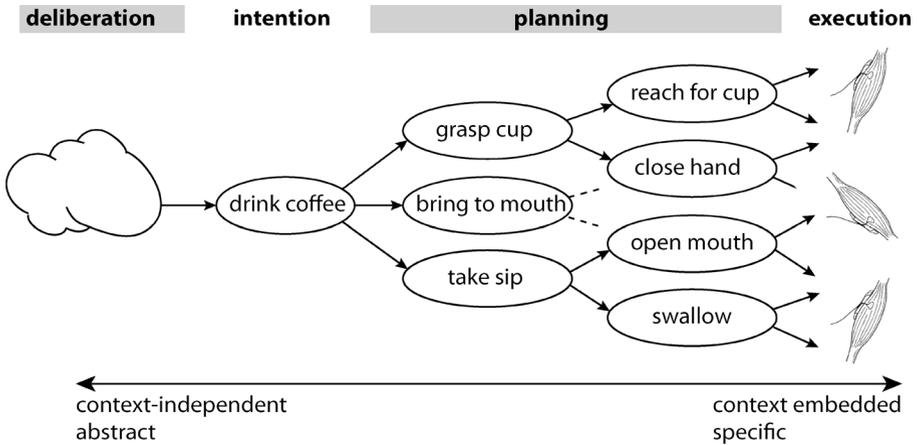
It is important to distinguish between the folk-psychological notion of intentions as the causes of future actions, and the theoretical construct of intentions as the causes of immediate actions. The latter is prevalent in philosophy of mind (Bratman 1987; Pacherie 2008; Searle 1983), and—either explicitly or implicitly—in cognitive psychology and neuroscience (Desmurget et al. 2009; Lafargue and Duffau 2008; Lau et al. 2004; Quiñ Quiroga et al. 2006; Sirigu et al. 2004; Vinding et al. 2014). Here we will challenge this latter notion by arguing that the neural mechanisms of action generation and control that are currently known are not compatible with the idea of discrete localizable states as the causes of action (see also Uithol et al. 2014). We will discuss how actions could be initiated if not stemming from a specialized “seat” in the brain (be it a single region or a network of regions), and we will argue that the cause of any given action is distributed across multiple processes and structures. We suggest that the complex interaction between dynamic processes in the brain and proprioceptive feedback from the body involved in the initiation of movement is better approached from a dynamical systems perspective.

## 2 The Search for Intentions in Brain Activity

If intentions are to be found anywhere, it seems reasonable to look for them in the brain. But what, precisely, would we be looking for? How would we recognize an intention, in the form of neuronal activity in the brain, if we were to find one? The neuroscientific search for “intentions in the brain” is arguably already under way (Brass and Haggard 2008; Carota et al. 2010; Desmurget and Sirigu 2009; Haggard 2008; Kriehoff et al. 2011; Lau et al. 2004), even though these foundational issues have yet to be squarely confronted. We begin with a rough sketch of the problem, with a focus on *immediate* (or *proximal*) intentions—the intention to act *now*.

The received picture of intentional action is as follows: After deliberation, we form an intention to act that is relatively abstract in the sense that the details of the execution are not part of the intention (Bratman 1981; Pacherie 2008; Pacherie and Haggard 2010). Next, this abstract idea is translated into more concrete sub-actions that are needed to bring the intention about. These sub-actions take the specific context into account, a process Pacherie (2008) calls “situational anchoring”. The sub-actions are translated into even more detailed sub-actions, and finally into motor programs that execute the movements. This model of action generation structured around a top-down causal hierarchy (Fig. 1) is widely accepted in cognitive (neuro)science (Grafton and Hamilton 2007; Haggard 2008; Hamilton and Grafton 2007; Kilner 2011; see Uithol et al. 2012 for a critical discussion), as a description of the brain’s solution to making the problem of action planning computationally tractable (Haggard 2005).

With this framework in mind, and starting with the onset of an intentional action we should, at least in theory, be able to trace the causal chain in reverse and converge upon the neural implementation of the intention. Hence the following thought experiment: Imagine that we have a database of very detailed neural recordings acquired while grasping actions were being performed. These recordings capture the entire sequence of neural events leading up to the initiation of a movement, from deliberation to muscle



**Fig. 1** The common view on action generation. Actions are planned in a hierarchical fashion, in which abstract intentions trigger a cascade of increasingly more concrete and context-sensitive processes (towards the right in this figure)

contraction, across a range of spatial scales. We can play the recordings back and forth, fast and slow, and zoom in and out as much as we want. Using these recordings, our aim is to find the neural event that corresponds to the “intention” to move. In order to keep things simple we will specify only two requirements: (1) the neural event should *cause* the action, and (2) the very same neural event should always be present when we attribute the same intention to the actor (i.e. the neural event should covary reliably with the attributed intention, see Section 8 below). Of course, if we were to converge on an event that causes the subsequent action, we could always look one step further back in time and find candidates for the cause of the cause, and so on. But note that criterion (2) ensures us that we will not go back in time endlessly because the dynamics leading to the intention need not be the same for different instances of the same intention. Note also that we do not require that the intention be conscious, but instead focus on the neural origins of the movement and remain agnostic as to their relation to consciousness.

Starting at the neuromuscular junction and following action potentials backwards in time, we might expect to arrive in primary motor cortex, from which we can step back in time yet another one or two synapses expecting that we should soon narrow in on the intention. However, this will not be the case: the primary motor cortex (M1), premotor cortex (PMC), supplementary motor area (SMA), and even primary somatosensory cortex (S1) all send projections directly to spinal motor neurons (Dum and Strick 1991; Galea and Darian-Smith 1994; He et al. 1993; Maier et al. 2013), so depending on what axon we choose to follow we could arrive in any one of these areas. In addition to receiving inputs from PMC and SMA, M1 also receives inputs, directly or indirectly, from primary somatosensory cortex (S1), multisensory area 5 of parietal cortex, basal ganglia, and cerebellum (Krakauer et al. 2000). And both PMC and SMA receive inputs from multiple other areas including prefrontal cortex and sensory areas, and all of these areas receive inputs from many other areas, as well as to and from each other (Krakauer et al. 2000; Rizzolatti and Luppino 2001). The above is simplified and almost certainly incomplete, but the point is that if we trace the causal path backwards in time from the neuromuscular junction we find a radical *divergence* of potential

causal inputs and no sign of convergence onto a localizable intention, as suggested in the classical hierarchy (Fig. 1).

In the next two sections we will discuss further evidence that the processes that control our behavior 1) cannot be localized or confined unambiguously, and 2) cannot be said to start at a certain point in time prior to the onset of the action itself. We argue that the brain and the skeletomotor system constitute a complex web of causal interactions. Consequently, the origins of actions cannot be confined unambiguously either in space or in time.

### 3 More Than One Road to Behavior

In a recent study, Vogelstein and colleagues (2014) used optogenetic methods combined with unsupervised machine learning techniques in order to produce a neuron-to-behavior map of the *entire* nervous system of the fruit fly larva (1054 genetic neuron lines that account for most of the roughly 10,000 neurons in the entire organism;  $N=37,780$  individual animals tested). The result is an atlas that maps all neurons and/or neuronal circuits that are causally related to each behavior in the animal's behavioral repertoire. One of the most noteworthy outcomes of this systematic study was that optogenetically (and very specifically) activating homologous neurons in different individual animals "did not always evoke the same behavior; rather it biased the probability toward a few possible behaviotypes." (Vogelstein et al. 2014, p.391). In addition, and most importantly, "even repeated activation of the same neurons in the same individual did not always evoke the same behaviotypes," although "the responses of individual animals were significantly more similar to each other than to the responses of distinct individuals" (Vogelstein et al. 2014, p.391).

This study raises two important points that are relevant to the present treatment. One is that the causal mapping from brain to behavior may be quite different between individuals, and may also vary considerably within the same individual over time. That is to say, very precise activation of a specific neuronal circuit does not fix the behavioral outcome, but rather depends on the state of other neurons in the system,<sup>1</sup> consistent with recent evidence of a "butterfly effect" in the brain (London et al. 2010).

Consequently, it may not be possible to trace a given behavioral outcome back to a specific neuronal cause in a way that generalizes robustly across different instances of the same behavior. This is consistent with the dynamical systems view of motor control advocated by Shenoy et al. (2013, see Section 6 below).

Other recent data reveals the emergence of stable patterns of neural activity in rodent M1 for cued movements (Peters et al. 2014), which might at first seem inconsistent with the argument above. But bear in mind that these patterns were observed in M1, where activity-pattern-to-behavior pairings are likely to be much more constrained than in premotor or other areas of cortex. Even more importantly, note that these were very well-learned and cued movements: As a result of training the effect of the cue may be to

<sup>1</sup> Although we have to keep in mind that activating a specific genetically-isolated sub-population of neurons may or may not correspond to something that ever happens in the life of the animal. Activating neurons through normal sensory pathways may yield more reliable neuro-behavioral mappings. However, other evidence of dynamic functional maps comes from optical imaging studies of spatial attention and optic flow processing in parietal cortex of behaving monkeys (Raffi 2005; Raffi and Siegel 2007).

constrain the state of the system (see Section 6) such that a specific activity-pattern-to-behavior mapping is selected, from among many that are possible in other contexts. Indeed Peters found that during learning the system “explored various activity patterns even during similar movements” (Peters et al. 2014, p. 263). Thus, even if highly reproducible patterns may emerge in primary motor cortex for well-learned cued movements, it seems clear from these two recent studies (Peters et al. 2014; Vogelstein et al. 2014) that it is possible for very different patterns of neural activity to lead to the very same movement, even at very late stages of motor control (as in M1).

In addition to being dynamic, the processes underlying action generation and control are highly distributed. To illustrate this we consider the well-studied grasping action. So-called *canonical neurons* are known to be involved in grasping actions in monkeys (Murata et al. 1997). Due to their response to both perceptual and motor features, they provide a fast and direct coupling between the perception of graspable objects and the execution of appropriate actions. Naturally, not every object in sight is immediately grasped, so these fast couplings need to be regulated by other control mechanisms. For example the proximity of the object (Rizzolatti et al. 1990), or the presence of a barrier (Bonini et al. 2014) influences their response. So even a simple grasping action involves several interdependent processes including processes that are dynamically coupled to context features, such as the availability of graspable objects.

The tight connectivity between motor processes and multimodal selection processes in action control suggests that action selection is not carried out in an abstract “reasoning space”, isolated from sensorimotor processes. Instead, there is mounting evidence that these decisions are made partly within the modality of the subsequent action. For example, the intention to speak is marked by activity in Broca’s area (Carota et al. 2010), the decision to move the hand or arm in pre-SMA (Lau et al. 2004), and the decision to either eat a piece of food or place it in a container is reported to be found in (monkey) IPL (Fogassi et al. 2005).

In line with this pattern of evidence, a recent theoretical account argues that the ‘what’ and the ‘how’ of an action are processed simultaneously (Cisek 2007; Cisek and Kalaska 2010). In his “Affordance Competition Hypothesis” Cisek (2007) explains that potential actions compete against each other for further processing. The dorsal visual system specifies actions, while a variety of biasing influences are provided by frontoparietal regions and the basal ganglia. A recent study provides support for this hypothesis by demonstrating that a supposedly “cognitive” competition between alternative options can be read out from the skeletomuscular system by perturbing the effector during perceptual decision making (Selen et al. 2012).

In addition to happening simultaneously and in parallel, the processes of action selection and action specification might be largely overlapping. For instance, the responses of canonical neurons discussed above are known to be highly specific for certain actions and certain objects (Murata et al. 1997). So the firing of a canonical neuron can herald the initiation of a grasping action and at the same time provide a specification of the action (i.e. how to grasp the object).<sup>2</sup> This is in line with the findings

<sup>2</sup> One might object by saying that the firing of canonical neurons is itself a consequence of a cognitive action plan, or an intention. However, the fact that canonical neurons respond very fast—100 ms after stimulus onset, (Murata et al. 1997)—and receive input relatively directly from visual areas—from IT through parietal area AIP (Borra et al. 2008) or from STP through PFG (Rozzi et al. 2006)—makes this suggestion highly unlikely.

of coordinated motor programs represented in monkey primary and premotor cortex (Graziano and Aflalo 2007; Graziano et al. 2002).<sup>3</sup>

The close interaction between sensorimotor and control processes is difficult to accommodate within the classical action-intention hierarchy in which deliberation, intention formation, and action planning are sequential stages in a causal chain. The data discussed above suggest that in many cases deliberation, action selection, and action specification occur in parallel as parts of a continuous and modal process. This is much more in line with the enactivist approach to cognition (Hutto and Myin 2013; Thompson and Varela 2001; Varela et al. 1991; see also Uithol and Maranesi 2014). We are not suggesting that we abandon the idea of an action hierarchy altogether. However, a more thorough rethinking of hierarchical organization is needed in order to capture this kind of complex interaction.

A fruitful alternative might be found in a hierarchy of dynamically interacting control processes (Kiebel et al. 2008; Koechlin and Summerfield 2007; Uithol et al. 2012). In this kind of control hierarchy, higher-level control processes are able to modulate or overrule lower-level processes, but the lower levels are themselves also capable of generating behavior.<sup>4</sup> The difference between this kind of control structure and the classical action hierarchy is that the latter is based on top-down causation, while the former is based on dynamic interaction. The processes “on top” of the control structure are different from the ones at lower levels in that they integrate information from multiple sources and over longer periods of time. Thus, in addition to being dynamic and distributed in space, the causes of our behavior may be spread out in time as well.

#### 4 Evidence for a Late Rather Than Early Neural Decision to Move

More than 30 years ago Libet et al. (1983) reported that the conscious intention to move follows the onset of the cortical “readiness potential” (RP; Kornhuber and Deecke 1965) by 300 ms or more. Libet and colleagues asked subjects to repeatedly perform a single, spontaneous finger flexion at a moment of their choosing—without deciding in advance when to move. Subjects also monitored a fast-moving clock dial and retrospectively estimated the time on the clock when they had first been aware of their capricious decision to move. While the RP stretches back in time for up to 500 ms or more prior to movement onset, the average subjective decision time was approximately 200 ms prior to movement onset (Haggard and Eimer 1999; Libet et al. 1983).

If some form of decision process is supposed to cause both the RP and the ensuing action, then how is it possible that the action is prepared before the decision becomes

<sup>3</sup> Graziano and colleagues find that although these chunks of behavior were complex—they consisted, for instance in an action sequence of reaching towards an object, grasping the object, bringing it to the mouth, and opening the mouth—they also needed adaptation to fit the current context. While this is a form of action specification, it is a form that is incompatible with the traditional hierarchy in which details are only added at the lowest levels. In this case details are already present, and are adapted to fit the specific context.

<sup>4</sup> In robotics, Brooks (1986, 1991) suggested a similar control structure, based on several autonomous control layers. Higher layers were able to modulate the activity of lower layers, but these lower layers were by themselves also capable of producing behavior. This alternative structure proved to cope with challenging environments much better than its contemporaries based on a top-down structure (Nilsson 1984).

conscious? The standard answer (and the one proposed by Libet) is that the decision is not conscious. We refer to this view of action initiation, where a decision-to-move (conscious or not) precedes the RP which in turn leads to an action, as the ‘early decision’ account.

Libet’s (1983) experiment has been replicated numerous times (Haggard and Eimer 1999; Keller and Heckhausen 1990; Miller et al. 2011; Schurger et al. 2012; Trevena and Miller 2002, 2010), including a recent replication using single-neuron recordings in humans that confirms a gradual increase in firing rate of individual neurons in advance of reported conscious-decision time (Fried et al. 2011).

How best to interpret the apparent buildup of neuronal activity preceding conscious decisions has been debated extensively (Gomes 1999). The general view is that the ‘early-decision’ account is not compatible with the idea of conscious intentions causing our actions, and many have taken this to be evidence that unconscious process causes both motor preparation and conscious intentions (Desmurget and Sirigu 2009; Haggard 2011; Wegner 2003). We suspect that this view remains tied to the traditional framework to the extent that it assumes a single moment or a single location within which the unconscious decision to act emerges (Dennett 2003). Here we will argue that there is another, less problematic interpretation for these findings, which we will call the ‘late-decision’ account.

We have suggested above that action control is instantiated by various interacting processes. Some of these processes are fast, and provide a relatively direct coupling between perception and action, while other processes are less direct, and capable of regulating or biasing these faster processes. None of these interacting processes can be said to be *the* initiator of the action, as none of these processes can, in and of itself, account for the subsequent behavior; it is through the interaction of these processes that action control emerges.

How, on this account, can the RP *precede* the conscious decision to move? Schurger et al. (2012) modeled the RP by averaging together the output of a leaky stochastic accumulator time-locked to moments when the output of the accumulator crossed a certain threshold (the “stochastic decision model”). The distribution of first-crossing times and the average accumulator output preceding threshold crossings fit the distribution of waiting times and average readiness potential (respectively) of subjects performing Libet’s task. This simple model thus reproduced the results of prior studies (Fried et al. 2011; Haggard and Eimer 1999; Libet et al. 1983), but with a completely different interpretation: the early pre-movement buildup reflects ongoing random fluctuations in cortical activity that bias the precise moment of movement onset.

While intrinsic fluctuations in brain activity may play a weakly causal role in self-initiated movement, they do not reflect a purposeful or intentional preparation for movement—their causal role is better described as incidental or accidental. Thus, according to this account, the slow early rise in the readiness potential in the Libet task is not caused by a preconscious “decision” to move, as most have assumed. In order to further test their account, Schurger et al. (2012) performed an adapted task in which subjects were sometimes given a surprise cue to perform the movement immediately, as quickly as possible. They found that relatively fast responses to surprise interruptions were preceded by a slow negative-going buildup in the EEG signal, just as the model predicted. This buildup even preceded the randomly-delivered surprise cues,

which does not make sense if the buildup is supposed to be preparatory (the brain cannot “prepare” for a movement that it does not know it is going to make<sup>5</sup>).

The rationale behind the model is that when there is no specific external imperative dictating *when* to move, the precise time of movement onset is largely determined by internal factors, which include ongoing spontaneous fluctuations in brain activity. On average, the initiation of movement will tend to coincide with crests in these ongoing fluctuations, because they (partly) determine when the motor threshold is crossed which triggers the movement. Since the relationship between the time of movement onset and the time of crests in spontaneous fluctuations is not random, the movement-locked average will in part reflect the temporal profile of spontaneous fluctuations (dominated by a slow exponential character). Thus, the buildup of neuronal activity prior to initiating movement is not caused by an unconscious decision; it is part of the process leading up to the decision: the crossing of the motor threshold. This is the ‘late decision’ account.

Noise plays an important role in this explanation of the RP, with the implication that not every rise towards the threshold will actually result in a crossing. The underlying neural activity could drop back down just before crossing the threshold, whereas the decision to press the button is only made when the threshold is actually crossed. Schurger et al. (2012) suggest reserving the term “decision” for a true neuronal commitment (Gold and Shadlen 2007) to “move now”, and suggest, based on prior evidence (R. Chen et al. 1998; Haggard and Eimer 1999), that this event occurs approximately 150 ms prior to movement onset – indeed this is approximately when subjects themselves estimate that the decision was made (Libet et al. 1983). According to the ‘late-decision’ account the initiation of an “intention” to act cannot be established unambiguously – although the commitment to “act now” can.

In the short time since the stochastic decision model was first introduced (Schurger et al. 2012), other studies have emerged providing evidence that ongoing stochastic variability might play a role in the choice of action (Bengson et al. 2014) and might account for the shape of the readiness potential (Jo et al. 2013). The model offers a parsimonious explanation of the mechanisms underlying ‘when’ a pre-specified action is initiated. Similar mechanisms could also underlie the decision of ‘what’ action to perform, and recent evidence confirms this speculation (Bengson et al. 2014). As mentioned above, Cisek (2007, Cisek and Kalaska 2010) developed a model in which multiple action possibilities are considered simultaneously. At some point the system settles on one of the options, and the other options are suppressed (Caruana et al. 2014). However, the increase in magnitude of one of the alternatives need not be caused by an earlier decision– it is part of the decision processes itself, which may be biased by spontaneous internal fluctuations, especially when the external imperative to move is weak or absent.

When intentional action is interpreted using the ‘late-decision account’, a number of puzzling findings can be made a lot less puzzling. For example, attempts to predict or detect the onset of movement in real time using the readiness potential (e.g. to drive a brain-computer interface) have met with only limited success (Bai et al. 2011; Lew

<sup>5</sup> Trials on which subjects were interrupted just as they were about to move were excluded from the analysis, and the distribution of slow and fast responses were the same early or late in the trial, thus ruling out the two most obvious alternative accounts (Schurger et al. 2012).

et al. 2012; Mason and Birch 2000). The stochastic decision model predicts that the RP should be modestly predictive of movement onset (better than chance), but not highly predictive. According to the theory, this will happen if ongoing fluctuations often approach the motor threshold without necessarily crossing it.

As another example, Kagaya and Takahata (2011) recently discovered that the crayfish exhibits a readiness discharge in advance of spontaneously initiated bouts of foraging, with a time course similar to that of the human RP. The classical account of the RP, which is grounded in cognitive and cortical phenomena, can do little more than cast a blank stare at these data - the crayfish is an invertebrate, and as such has no cortex and very little in common with the mammalian motor system. According to the stochastic decision model, on the other hand, all that is required for both vertebrates and invertebrates to exhibit a similar buildup in neuronal activity preceding spontaneous behavioral events is for the spectral properties of ongoing “background” neuronal activity to be similar in both vertebrate and invertebrate phyla, and this appears to be the case (Garcia-Perez et al. 2007; Mazzone et al. 2007).

As a final example, using functional magnetic resonance imaging (fMRI), Soon et al. (2008) were able to predict whether subjects would choose to press a left or a right button with above-chance accuracy up to 8 s before the choice was consciously made. Thus, the likely outcome of the decision (using left or right hand) appears to manifest well in advance of the conscious decision to act. The surprise with which these results were received, both within and outside of the scientific community, betrays an implicit dualist stance. Consider what it would mean if there existed *no prior information in the brain* about upcoming decisions: it would mean that decisions can literally appear spontaneously in the brain, as if from a vacuum. In our view Soon et al. (2008, 2013) made explicit what we already implicitly presumed to be true—decisions do not emerge from a vacuum, they emerge from the causal milieu of ongoing brain activity, although they may not be fully determined until they are actually made. This research has made an important contribution, not by identifying the locus from which decisions originate, but by exposing the time course over which brain states influence future decisions, and regions where the influence is apparent. Nevertheless, the successful application of fMRI and multivariate decoding strategies in this area of research has been widely considered as evidence that intentions *can* be localized in the brain. In the next section we critically evaluate the role of fMRI in the study of immediate intentions.

## 5 fMRI in the Study of Immediate Intentions

In a controlled fMRI experiment (and in neuroscientific experiments in general) as many variables as possible—ideally all-but-one—are held equal for two different conditions. Any extraneous context or stimulus feature that could influence the response is carefully banished from the experimental setting. This is of course an important prerequisite for any neurocognitive investigation. However, if you succeed in controlling all but one factor, you also ensure that the dynamical system you are studying will appear to function like a representational system: you will find differences in the activation of brain loci that are correlated with the experimental variables *under these specific experimental constraints*. It is tempting to conclude that activity in these brain loci *represents* the cognitive operation at hand (e.g. conclusions of the type: “the

intention to do A is represented in area X"). Yet, such an interpretation ignores the dynamic interaction between the process at hand and other factors that are equal in the two conditions (such as the task context), or at least highly similar in terms of gross activation.

While fMRI has traditionally focused on identifying "regions of interest" (ROIs), a more recent trend has been to apply multivariate pattern classification methods to the very-high-dimensional space of voxels (Haynes 2011). This approach can in theory identify widely-distributed patterns of brain activity that distinguish between different cognitive or perceptual states (Haxby et al. 2001), rather than trying to isolate a single ROI or group of ROIs. In addition, statistical methods for fMRI analysis have advanced appreciably in recent years, allowing us to make inferences about functional connectivity between regions.

The machine-learning and functional-connectivity approaches are indeed highly promising and have already contributed greatly to moving neuroimaging research away from an overly-modular view of cortical function. However, the very low temporal resolution of fMRI relative to the time scale of neuronal events remains a limiting factor in the study of immediate intentions. Consider that it takes on the order of 1.5 or 2 s to acquire a single whole-brain sample of brain activity using fMRI, and that a given causal pattern of brain activity might appear and disappear (and contribute whatever effect it has on behavior) faster than the blink of an eye. From the point of view of fMRI, a distributed pattern with nodes that were activated in series, in parallel, or in a more complex interleaved fashion will all appear identical. Also, using fMRI, areas with sustained activity are more readily detected than areas with brief patterns of activity, even if the latter happen to play a more direct causal role.

Another caveat regarding fMRI pattern classification concerns patterns of activity that are only modestly informative as to an upcoming decision (i.e. better-than-chance, but far from ceiling; (e.g. Soon et al. 2008, 2013)). Such effects may reflect the causal milieu rather than the intention itself. Here "causal milieu" refers to ongoing brain activity not specifically related to or directed at producing the action, but that plays a secondary role in biasing the decision. As long as such effects are reliable, they will be picked up by a pattern classifier, even if the proximal causal patterns were too fleeting or spatially variable to be detected.

Although multi-voxel pattern classification is almost always performed within-subject (cf. Hasson 2004), inferences about the location of informative voxels are normally made based on averaging the brain maps across many subjects. fMRI might therefore tend to give the impression of focal and localizable effects due to the nature of group-level averages, which capitalize on the spatial smearing inherent in the fMRI BOLD signal (compounded by the spatial smoothing that is commonly applied during pre-processing). Group-level averages discount variability between subjects in the discriminative patterns and emphasize local regions where informative voxels tend to overlap across subjects (Schurger et al. 2010), giving the impression that a causal arrow can be drawn from a given bright spot in the brain to a given decision.

Thus even multivariate pattern classification is not immune to the ROI-centric tendency in fMRI research: as soon as one has access to the pattern of informative voxels that the classifier used, one is drawn towards trying to find a cluster somewhere (Bode et al. 2011; Schurger et al. 2010; Soon et al. 2008). After all, what researcher would bother to publish voxel maps that looked like television noise, and on top of that

were different for each subject? Maybe this is because we tend to infuse brain maps with the assumption that there is a “where” in the first place—that is what a *map* is supposed to be for! There may indeed be a tendency for at least some of the informative voxels to commonly appear in roughly the same region(s) across subjects. But the fact that a statistical map has one or more isolated regions in common does not in any way imply that the cognitive process with which they are correlated somehow “localizes” causally to these regions. Thus, while multivariate approaches to fMRI might reveal patterns of activity associated with decisions that play out on the order of several seconds, the group-level analysis of fMRI data is not ideally-suited to the study of proximal neural decisions to act *now*.

So far we have discussed what we perceive to be some of the limitations of the contemporary framework and tools used in the empirical study of intentions. One point that we have tried to make is that the utility of powerful techniques is limited by the conceptual framework within which those techniques are applied. We have argued that the classical hierarchical control framework fails to capture the dynamic and distributed nature of causal interactions in the brain. In the next session we discuss the dynamical systems approach, which we believe is a more appropriate framework for the empirical study of intentions.

## 6 A Dynamical-Systems Approach to the Neuroscience of Action

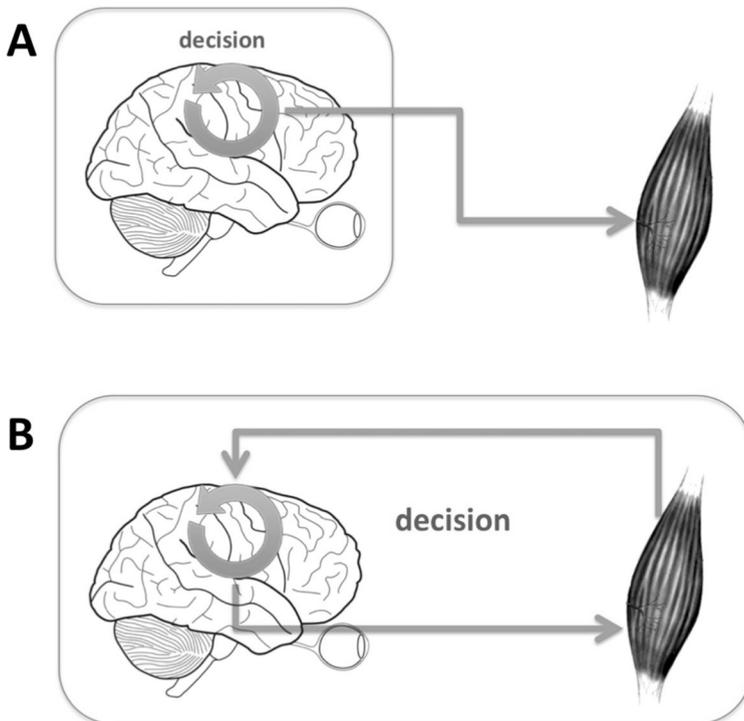
Although the idea of applying dynamical systems theory to the study of neural and cognitive mechanisms has been around for at least two decades (Beer 2000; Kelso 1995; Thelen and Smith 1994; van Gelder 1998), the dynamical systems approach has only recently begun to figure prominently in neuroscience (Albantakis and Deco 2011; Ames et al. 2014; Churchland 2006; Churchland et al. 2012, 2010; Deco and Romo 2008; Domnisoru et al. 2013; Jezek et al. 2011; Kaufman et al. 2014; Machens 2005; Wong and Wang 2006).

In line with the evidence discussed in Section 3, Shenoy et al. (2013) points out that knowledge of the motor output alone is unlikely to be sufficient to determine the pattern of neuronal activity that produced it. The tuning of neurons that participate in generating movement is constantly changing even when the functional mapping from stimulus to response remains fixed. In Shenoy’s words, the nervous system has to “generate a pattern of activity appropriate to drive the desired movement”, even if that pattern is different every time the same movement is made (because e.g. the starting conditions are different). If a given intention causes the pattern of activity needed to bring about a specific movement, and the pattern of activity needed to bring about the movement is variable, then the pattern of activity corresponding to the intention must also be variable—even when, functionally speaking, the intention remains the same. Thus, even with the ability to account for the activity of every single neuron in an organism, the very nature of a highly complex connectome drenched through and through with subtle stochastic fluctuations may limit the degree to which causal patterns generalize across individuals, or even across different episodes in the same individual. What we might find in brain-to-behavior mappings may be more indefinite and probabilistic in nature—statistical tendencies that approach (but never reach) a

definitive causal arrow as the context, behavioral parameters, and recent past are ever more tightly constrained.

Another intriguing consequence of this dynamic perspective for the study of movement initiation appears when we consider how *decisions* are made in dynamical systems and what this might mean for our conception of movement decisions. A seemingly obvious and common assumption about movement initiation is that movement decisions are made in the brain and then transmitted to the muscle via spinal motor neurons (Fig. 2a), which we refer to as the “central-decision” assumption. However, when the functional role of feedback pathways between the body and brain in decision-making is taken into account, then the blurring of cause and effect inherent in dynamical systems suggests a very different possibility.

One prevalent source of abundant and rapid motor-related feedback are the specialized proprioceptive organelles called muscle spindles. These relay continuous information about change in muscle stretch and are commonly encountered in textbooks because they are responsible for the well-known “knee-jerk” reflex. Muscle spindles are engaged in a continuous dynamic feedback interaction with both central and peripheral motor systems that not only receive information from the spindles, but also dynamically control the gain of muscle spindles. Feedback from muscle spindles is impressively fast, with latencies as short as 12 ms, and Dimitriou and Edin (2010) provide evidence that muscle spindles even *anticipate* upcoming kinematic states— a



**Fig. 2** **a** The classical picture of a decision being formed in the brain and subsequently propagated to the muscles. **b** The proposed alternative view in which a decision is formed using feedback from within the motor system as well as from the periphery

property that can only be accounted for through interaction with central motor systems. In addition, Selen et al. (2012) have shown that skeletal muscles carry information about the decision process leading up to a cortical motor command, suggesting that muscles are “plugged in” to the cortical dynamics within which the decision plays out—not just receiving the outcome of the decision.

These observations afford us a very different view of movement initiation (Fig. 2b) in which decisions-to-move *now*, on the immediate time-scale, are realized by the *brain + body*. This view suggests the intriguing possibility that not only the process leading to a decision (Selen et al. 2012), but also the fact of having decided in the first place may incorporate feedback from the body.<sup>6</sup> Thus if this decision model is accurate, then the seat of intentions may not only prove to be elusive *within* the brain, but the arena of intentions may even extend *outside* the brain (Beer 2003; Buhrmann et al. 2013).

## 7 Intentions as Perceptual Attributions

A possible reason why it may be difficult to find intentions in the brain is simply because intentions are not neural phenomena—they are perceptual attributions that we use to describe behavior, of ourselves and others (Uithol and Paulus 2013). We do not have access to the complex interaction of our motor control processes, let alone those of others, but we do have access to the result of these processes: behavior. When we see an ongoing stream of behavior, we extract and isolate the movements that we are interested in and call this the ‘action’, a process called the ‘teleological approach’ (McFarland 1989) to behavior. An intention, seen in this light, is an explanation that accounts for precisely that aspect of behavior that happens to be most relevant in the current context. Our assumption that behavior stems from a discrete cause might derive from our tendency to chunk behavior in discrete events (Uithol et al. 2014), and infer a discrete cause somewhere inside of the agent.

Evidence for the fact that the attributed intention is partly dependent on the observer can be found in the fact that the judgment of an action as intentional or unintentional seems to be dependent in part on the valence of the outcome of an action (Knobe 2006). Given this dependence, it is unlikely that any given intention will have a fixed one-to-one correspondence with a specific spatial pattern of brain activity, as this would require the neural state to be dependent on who is observing the behavior.

If intentions are not brain states, how is the distinction between intentional and non-intentional actions to be understood? It has been suggested that intentional actions are accompanied by (higher-level) predictions, whereas they are absent with passive movements (Blakemore et al. 2002). However, it is unlikely that the distinction can be phrased entirely in terms of the neural causes of action because of its contextual dependencies described above. Thus we should not always expect to find a consistent or even coherent mapping between intentions, as we attribute them to others and to ourselves, and a reliable, reproducible, and localizable antecedent brain state. We might

<sup>6</sup> Of course, even if our conjecture is correct (and it remains speculative), feedback from the body may exert only a small modulatory influence in most contexts: i.e. the “central-decision” assumption may still be valid to a first approximation, much like Newtonian dynamics in physics; relativity theory offers a fuller, subtler explanation, but very many phenomena can be accounted for without it.

find candidate brain states when the context is highly constrained and/or the data analyses involve spatiotemporal smoothing and averaging, but we must be very careful about how we interpret such data.

## 8 Concluding Remarks

Here we have argued that in general our behavior is based on many coupled processes that are each modulated by external events to a greater or lesser degree. Consequently, trying to designate a single neural cause as the origin of an intentional action will be a self-limiting approach, as it fails to capture the complex and dynamic interaction between decision-making and action control, with causation being distributed over various dynamically interacting brain processes.

We have focused on the processes underlying immediate actions. Are the processes underlying future actions any different? This might be the case, but it need not be. Indeed, part of our argument is based on the fact that the mechanisms for action control are context-sensitive, and for future actions the context for the action need not be the current context. So one could take this as evidence that future actions are caused by intentions. Yet, the extent to which actions in the future rely on states that that can rightfully be considered representations of intentions remains an open question. We consider it a matter of scientific soundness not to posit such discrete states as long as there might be a more parsimonious explanation available.

We suggest that the dynamical systems approach offers a potent, albeit less intuitive explanatory framework. Neuroscience can take on the project of finding the neural origins of voluntary action; not (just) by isolating brain regions that are supposed to represent certain types of intentions, but by studying the dynamic interaction between modal decision processes, motor specification processes, and motor predictions.

This new approach brings with it a number of new research questions, both philosophical and empirical. If actions do not stem from intentions, then what makes certain behaviors intentional, and others non-intentional? How can we define intentions so that searching for them in the brain is both conceptually sound and empirically tractable? What empirical techniques can we use to identify non-trivial causal relationships between brain activity and behavior? To what extent do predictions in the skeletomotor system play a role? What is the most appropriate shared definition of ‘causality’ (Kistler 2006) for both philosophical and empirical approaches to intentional behavior? We should take account of the ways in which the brain interacts dynamically with the body and the world, effectively bringing these seemingly “external” factors into that causal arena. Future research that embraces the dynamic interactions within the brain and between the brain, body, and sensory context is needed to understand the causal origins of voluntary action, and dynamical systems theory offers an approach that is suited to this endeavor.

**Acknowledgments** Aaron Schurger was supported by a grant from the Association Robert Debre Pour la Recherche Medicale. Sebo Uithol was supported by the EU grant ‘Towards an Embodied Science of Intersubjectivity’ (TESIS, FP7-PEOPLE-2010-ITN, 264828).

## References

- Albantakis, L., and G. Deco. 2011. Changes of mind in an attractor network of decision-making. *PLoS Computational Biology* 7(6): e1002086. doi:10.1371/journal.pcbi.1002086.
- Ames, K.C., S.I. Ryu, and K.V. Shenoy. 2014. Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron* 81(2): 438–451. doi:10.1016/j.neuron.2013.11.003.
- Bai, O., V. Rathi, P. Lin, D. Huang, H. Battapady, D.-Y. Fei, et al. 2011. Prediction of human voluntary movement before it occurs. *Clinical Neurophysiology* 122(2): 364–372. doi:10.1016/j.clinph.2010.07.010.
- Beer, R. 2000. Dynamical approaches to cognitive science. *Trends in Cognitive Sciences* 4(3): 91–99.
- Beer, R. 2003. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* 11(4): 209.
- Bengson, J.J., T.A. Kelley, X. Zhang, J.-L. Wang, and G.R. Mangun. 2014. Spontaneous neural fluctuations predict decisions to attend. *Journal of Cognitive Neuroscience* 47: 1–7. doi:10.2307/2531248.
- Blakemore, S.-J., D.M. Wolpert, and C.D. Frith. 2002. Abnormalities in the awareness of action. *Trends in Cognitive Sciences* 6(6): 237–242. doi:10.1016/S1364-6613(02)01907-1.
- Bode, S., A.H. He, C.S. Soon, R. Trampel, R. Turner, and J.-D. Haynes. 2011. Tracking the unconscious generation of free decisions using ultra-high field fMRI. *PLoS One* 6(6): e21612. doi:10.1371/journal.pone.0021612.
- Bonini, L., M. Maranesi, A. Livi, L. Fogassi, and G. Rizzolatti. 2014. Space-dependent representation of objects and other's action in monkey ventral premotor grasping neurons. *Journal of Neuroscience* 34(11): 4108–4119. doi:10.1523/JNEUROSCI.4187-13.2014.
- Borra, E., A. Belmalih, R. Calzavara, M. Gerbella, A. Murata, S. Rozzi, and G. Luppino. 2008. Cortical connections of the macaque Anterior Intraparietal (AIP) area. *Cerebral Cortex* 18(5): 1094–1111.
- Brass, M., and P. Haggard. 2008. The what, when, whether model of intentional action. *The Neuroscientist* 14(4): 319–325. doi:10.1177/1073858408317417.
- Bratman, M.E. 1981. Intention and means-end reasoning. *The Philosophical Review* 90(2): 252–265.
- Bratman, M.E. 1987. *Intention, plans, and practical reason*. Cambridge: Harvard University Press.
- Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* 2(1): 14–23.
- Brooks, R. 1991. Intelligence without representation. *Artificial Intelligence* 47: 139–159.
- Buhrmann, T., E.A. Di Paolo, and X. Barandiaran. 2013. A dynamical systems account of sensorimotor contingencies. *Frontiers in Psychology* 4: 285. doi:10.3389/fpsyg.2013.00285.
- Carota, F., A. Posada, S. Harquel, C. Delpuech, O. Bertrand, and A. Sirigu. 2010. Neural dynamics of the intention to speak. *Cerebral Cortex* 20(8): 1891–1897. doi:10.1093/cercor/bhp255.
- Caruana, F., S. Uithol, G. Cantalupo, I. Sartori, Russo, G. Lo, and P. Avanzini. 2014. How action selection can be embodied: Intracranial gamma band recording shows response competition during the Eriksen flankers test. *Frontiers in Human Neuroscience* 8(668): 1–9. doi:10.3389/fnhum.2014.00668.
- Chen, R., Z. Yaseen, L.G. Cohen, and M. Hallett. 1998. Time course of corticospinal excitability in reaction time and self-paced movements. *Annals of Neurology* 44(3): 317–325. doi:10.1002/ana.410440306.
- Churchland, M.M. 2006. Neural variability in premotor cortex provides a signature of motor preparation. *Journal of Neuroscience* 26(14): 3697–3712. doi:10.1523/JNEUROSCI.3762-05.2006.
- Churchland, M.M., J.P. Cunningham, M.T. Kaufman, S.I. Ryu, and K.V. Shenoy. 2010. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron* 68(3): 387–400. doi:10.1016/j.neuron.2010.09.015.
- Churchland, M.M., J.P. Cunningham, M.T. Kaufman, J.D. Foster, P. Nuyujukian, S.I. Ryu, and K.V. Shenoy. 2012. Neural population dynamics during reaching. *Nature*. doi:10.1038/nature11129.
- Cisek, P. 2007. Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B-Biological Sciences* 362(1485): 1585–1599. doi:10.1146/annurev.neuro.20.1.25.
- Cisek, P., and J.F. Kalaska. 2010. Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience* 33(1): 269–298. doi:10.1146/annurev.neuro.051508.135409.
- Deco, G., and R. Romo. 2008. The role of fluctuations in perception. *Trends in Neurosciences* 31(11): 591–598. doi:10.1016/j.tins.2008.08.007.
- Dennett, D.C. 2003. *Freedom evolves*. New York: Viking.
- Dennett, D.C., and M. Kinsbourne. 1992. Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences* 15: 183–247.
- Desmurget, M., and A. Sirigu. 2009. A parietal-premotor network for movement intention and motor awareness. *Trends in Cognitive Sciences* 13(10): 411–419.

- Desmurget, M., K.T. Reilly, N. Richard, A. Szathmari, C. Mottolese, and A. Sirigu. 2009. Movement intention after parietal cortex stimulation in humans. *Science* 324(5928): 811–813. doi:10.1126/science.1169896.
- Dimitriou, M., and B.B. Edin. 2010. Human muscle spindles act as forward sensory models. *Current Biology* 20(19): 1763–1767. doi:10.1016/j.cub.2010.08.049.
- Domnisoru, C., A.A. Kinkhabwala, and D.W. Tank. 2013. Membrane potential dynamics of grid cells. *Nature* 495(7440): 199–204. doi:10.1038/nature11973.
- Dum, R.P., and P.L. Strick. 1991. The origin of corticospinal projections from the premotor areas in the frontal lobe. *The Journal of Neuroscience* 11(3): 667–689.
- Fogassi, L., P.F. Ferrari, B. Gesierich, S. Rozzi, F. Chersi, and G. Rizzolatti. 2005. Parietal lobe: From action organization to intention understanding. *Science* 308(5722): 662–666.
- Fried, I., R. Mukamel, and G. Kreiman. 2011. Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron* 69(3): 548–562. doi:10.1016/j.neuron.2010.11.045.
- Galea, M.P., and I. Darian-Smith. 1994. Multiple corticospinal neuron populations in the macaque monkey are specified by their unique cortical origins, spinal terminations, and connections. *Cerebral Cortex* 4(2): 166–194. doi:10.1093/cercor/4.2.166.
- Garcia-Perez, E., A. Mazzoni, and V. Torre. 2007. Spontaneous electrical activity and behavior in the leech *Hirudo medicinalis*. *Frontiers in Integrative Neuroscience* 1(8): 1–9.
- Gold, J.I., and M.N. Shadlen. 2007. The neural basis of decision making. *Annual Review of Neuroscience* 30(1): 535–574. doi:10.1146/annurev.neuro.29.051605.113038.
- Gomes, G. 1999. Volition and the readiness potential. *Journal of Consciousness Studies* 25(2): 157–181.
- Grafton, S.T., and A.F.C. Hamilton. 2007. Evidence for a distributed hierarchy of action representation in the brain. *Human Movement Science* 26(4): 590–616.
- Graziano, M.S.A., and T.N.S. Aflalo. 2007. Mapping behavioral repertoire onto the cortex. *Neuron* 56(2): 239–251.
- Graziano, M.S.A., C. Taylor, and T. Moore. 2002. Complex movements evoked by microstimulation of precentral cortex. *Neuron* 34(5): 841–851.
- Haggard, P. 2005. Conscious intention and motor cognition. *Trends in Cognitive Sciences* 9(6): 290–295.
- Haggard, P. 2008. Human volition: Towards a neuroscience of will. *Nature Reviews. Neuroscience* 9(12): 934–946. doi:10.1038/nrn2497.
- Haggard, P. 2011. Decision time for free will. *Neuron* 69(3): 404–406. doi:10.1016/j.neuron.2011.01.028.
- Haggard, P., and M. Eimer. 1999. On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research* 126(1): 128–133.
- Hamilton, A.F.C., and S.T. Grafton. 2007. The motor hierarchy: From kinematics to goals and intentions. In *Attention & performance 22. Sensorimotor foundations of higher cognition attention and performance*, ed. P. Haggard, Y. Rossetti, and M. Kawato, 381–408. Oxford: Oxford University Press.
- Hasson, U. 2004. Intersubject synchronization of cortical activity during natural vision. *Science* 303(5664): 1634–1640. doi:10.1126/science.1089506.
- Haxby, J., M. Gobbini, M. Furey, A. Ishai, J. Schouten, and P. Pietrini. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293(5539): 2425–2429.
- Haynes, J.-D. 2011. Decoding and predicting intentions. *Annals of the New York Academy of Sciences* 1224(1): 9–21. doi:10.1111/j.1749-6632.2011.05994.x.
- He, S.Q., R.P. Dum, and P.L. Strick. 1993. Topographic organization of corticospinal projections from the frontal lobe: motor areas on the lateral surface of the hemisphere. *The Journal of Neuroscience* 13(3): 952–980.
- Hutto, D.D., and E. Myin. 2013. *Radicalizing enactivism: Basic minds without content*. Cambridge: MIT Press.
- Jezek, K., E.J. Henriksen, A. Treves, E.I. Moser, and M.-B. Moser. 2011. Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478(7368): 246–249. doi:10.1038/nature10439.
- Jo, H.-G., T. Hinterberger, M. Wittmann, T.L. Borghardt, and S. Schmidt. 2013. Spontaneous EEG fluctuations determine the readiness potential: Is preconscious brain activation a preparation process to move? *Experimental Brain Research* 231(4): 495–500. doi:10.1007/s00221-013-3713-z.
- Kagaya, K., and M. Takahata. 2011. Readiness discharge for spontaneous initiation of walking in crayfish. *The Journal of Neuroscience* 30(4): 1348–1362.
- Kaufman, M.T., M.M. Churchland, S.I. Ryu, and K.V. Shenoy. 2014. Cortical activity in the null space: Permitting preparation without movement. *Nature Neuroscience*. doi:10.1038/nn.3643.
- Keller, I., and H. Heckhausen. 1990. Readiness potentials preceding spontaneous motor acts: Voluntary vs. involuntary control. *Electroencephalography and Clinical Neurophysiology* 76(4): 351–361. doi:10.1016/0013-4694(90)90036-J.
- Kelso, J.A.S. 1995. *Dynamic patterns: The self-organization of brain and behavior*. Cambridge: MIT Press.
- Kiebel, S.J., J. Daunizeau, and K.J. Friston. 2008. A hierarchy of time-scales and the brain. *PLoS Computational Biology* 4(11): e1000209. doi:10.1371/journal.pcbi.1000209.

- Kilner, J.M. 2011. More than one pathway to action understanding. *Trends in Cognitive Sciences* 15(8): 352–357. doi:10.1016/j.tics.2011.06.005.
- Kistler, M. 2006. La causalité comme transfert et dépendance nomique. *Philosophie* 89(1): 53–77. doi:10.3917/philo.089.0053.
- Knobe, J. 2006. The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies* 130(2): 203–231.
- Koechlin, E., and C. Summerfield. 2007. An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences* 11(6): 229–235.
- Kornhuber, H.H., and L.D. Deecke. 1965. Hirnpotentialänderungen bei Willkürbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. *Pflügers Archiv Für Die Gesamte Physiologie Des Menschen Und Der Tiere* 284(1): 1–17. doi:10.1007/BF00412364.
- Krakauer, J.W., Z.M. Pine, M.-F. Ghilardi, and C. Ghez. 2000. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *The Journal of Neuroscience* 20(23): 8916–8924.
- Kriehoff, V., F. Waszak, W. Prinz, and M. Brass. 2011. Neural and behavioral correlates of intentional actions. *Neuropsychologia* 49(5): 767–776. doi:10.1016/j.neuropsychologia.2011.01.025.
- Lafargue, G., and H. Duffau. 2008. Awareness of intending to act following parietal cortex resection. *Neuropsychologia* 46(11): 2662–2667. doi:10.1016/j.neuropsychologia.2008.04.019.
- Lau, H.C., R.D.D. Rogers, P. Haggard, and R.E. Passingham. 2004. Attention to intention. *Science* 303(5661): 1208–1210. doi:10.1126/science.1090973.
- Lew, E., Chavarriaga, R., Silvoni, S., and J.D.R. Millán. 2012. Detection of self-paced reaching movement intention from EEG signals. *Frontiers in Neuroengineering* 5(13). doi:10.3389/fneng.2012.00013
- Libet, B., C.A. Gleason, E.W. Wright, and D.K. Pearl. 1983. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential): The unconscious initiation of a freely voluntary act. *Brain* 106(3): 623–642. doi:10.1093/brain/106.3.623.
- London, M., et al. 2010. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* 466(7302): 123–127.
- Machens, C.K. 2005. Flexible control of mutual inhibition: A neural model of two-interval discrimination. *Science* 307(5712): 1121–1124. doi:10.1126/science.1104171.
- Maier, M.A., P.A. Kirkwood, T. Brochier, and R.N. Lemon. 2013. Responses of single corticospinal neurons to intracortical stimulation of primary motor and premotor cortex in the anesthetized macaque monkey. *Journal of Neurophysiology* 09(12): 2982–2998.
- Mason, S.G., and G.E. Birch. 2000. A brain-controlled switch for asynchronous control applications. *IEEE Transactions on Bio-Medical Engineering* 47(10): 1297–1307. doi:10.1109/10.871402.
- Mazzoni, A., et al. (2007). On the Dynamics of the Spontaneous Activity in Neuronal Networks. *PLoS ONE* 2(5): e439.
- McFarland, D. 1989. Goals, no-goals and own goals. In *Goals, no-goals and own goals: A debate on goal-directed and international behaviour*, ed. A. Montefiore and D. Noble, 39–57. London: Unwin Hyman.
- Miller, J., P. Shepherdson, and J.A. Trevena. 2011. Effects of clock monitoring on electroencephalographic activity: Is unconscious movement initiation an artifact of the clock? *Psychological Science* 22(1): 103–109. doi:10.1177/0956797610391100.
- Murata, A., L. Fadiga, L. Fogassi, V. Gallese, V. Raos, and G. Rizzolatti. 1997. Object representation in the ventral premotor cortex (Area F5) of the monkey. *Journal of Neurophysiology* 78(4): 2226–2230.
- Nilsson, N. 1984. *Shakey The Robot, Technical Note 323*
- Pacherie, E. 2008. The phenomenology of action: A conceptual framework. *Cognition* 107(1): 179–217.
- Pacherie, E., and P. Haggard. 2010. What are intentions? In *Conscious will and responsibility. A tribute to benjamin libet*, ed. W. Sinnott-Armstrong and L. Nadel, 70–84. Oxford: Oxford University Press.
- Peters, A.J., S.X. Chen, and T. Komiyama. 2014. Emergence of reproducible spatiotemporal activity during motor learning. *Nature*. doi:10.1038/nature13235.
- Quian Quiroga, R., L.H. Snyder, A.P. Batista, H. Cui, and R.A. Andersen. 2006. Movement intention is better predicted than attention in the posterior parietal cortex. *Journal of Neuroscience* 26(13): 3615–3620. doi:10.1523/JNEUROSCI.3468-05.2006.
- Raffi, M. 2005. Functional architecture of spatial attention in the parietal cortex of the behaving monkey. *Journal of Neuroscience* 25(21): 5171–5186. doi:10.1523/JNEUROSCI.5201-04.2005.
- Raffi, M., and R.M. Siegel. 2007. A functional architecture of optic flow in the inferior parietal lobule of the behaving monkey. *PLoS ONE* 2(2): e200.
- Rizzolatti, G., and G. Luppino. 2001. The cortical motor system. *Neuron* 31(6): 889–901. doi:10.1016/S0896-6273(01)00423-8.

- Rizzolatti, G., Gentilucci, M., Camarda, R. M., Gallese, V., Luppino, G., Matelli, M., and L. Fogassi. 1990. Neurons related to reaching-grasping arm movements in the rostral part of area 6 (area 6a?). *Experimental Brain Research* 82(2). doi:10.1007/BF00231253
- Rozzi, S., R. Calzavara, A. Belmalih, E. Borra, G.G. Gregoriou, M. Matelli, and G. Luppino. 2006. Cortical connections of the inferior parietal cortical convexity of the macaque monkey. *Cerebral Cortex* 16(10): 1389–1417.
- Schurger, A., F. Pereira, A. Treisman, and J.D. Cohen. 2010. Reproducibility distinguishes conscious from nonconscious neural representations. *Science* 327(5961): 97–99. doi:10.1126/science.1180029.
- Schurger, A.A., J.D. Sitt, and S.S. Dehaene. 2012. An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences of the United States of America* 109(42): E2904–E2913. doi:10.1073/pnas.1210467109.
- Searle, J. 1983. *Intentionality, an essay in the philosophy of mind*. Cambridge: Cambridge University Press.
- Selen, L.P.J., M.N. Shadlen, and D.M. Wolpert. 2012. Deliberation in the motor system: Reflex gains track evolving evidence leading to a decision. *Journal of Neuroscience* 32(7): 2276–2286. doi:10.1523/JNEUROSCI.5273-11.2012.
- Shenoy, K. V., Sahani, M., and M.M. Churchland. 2013. Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annual Review of Neuroscience* 36(1). doi:10.1146/annurev-neuro-062111-150509
- Sirigu, A., E. Daprati, S. Ciancia, P. Giroux, N. Nighoghossian, A. Posada, and P. Haggard. 2004. Altered awareness of voluntary action after damage to the parietal cortex. *Nature Neuroscience* 7(1): 80–84. doi:10.1038/nn1160.
- Soon, C.S., M. Brass, H.-J. Heinze, and J.-D. Haynes. 2008. Unconscious determinants of free decisions in the human brain. *Nature Neuroscience* 11(5): 543–545. doi:10.1038/nn.2112.
- Soon, C. S., Hanxi He, A., Bode, S., and J.-D. Haynes. 2013. Predicting free choices for abstract intentions. *Proceedings of the National Academy of Sciences* 1–6. doi:10.1073/pnas.1212218110
- Thelen, E., and L. Smith. 1994. *A dynamic systems approach to the development of cognition and action*. Cambridge: MIT Press.
- Thompson, E., and F. Varela. 2001. Radical embodiment: Neural dynamics and consciousness. *Trends in Cognitive Sciences* 5(10): 418–425.
- Trevena, J.A., and J. Miller. 2002. Cortical movement preparation before and after a conscious decision to move. *Consciousness and Cognition* 11(2): 162–190. doi:10.1006/ccog.2002.0548.
- Trevena, J.A., and J. Miller. 2010. Brain preparation before a voluntary action: Evidence against unconscious movement initiation. *Consciousness and Cognition* 19(1): 447–456. doi:10.1016/j.concog.2009.08.006.
- Uithol, S., and M. Maranesi. 2014. No need to match: A comment on Bach, Nicholson and Hudson's 'affordance-matching hypothesis'. *Frontiers in Human Neuroscience* 8(710): 1–2. doi:10.3389/fnhum.2014.00710.
- Uithol, S., and M. Paulus. 2013. What do infants understand of others' action? A theoretical account of early social cognition. *Psychological Research* 78(5): 609–622. doi:10.1007/s00426-013-0519-3.
- Uithol, S., I. van Rooij, H. Bekkering, and W.F.G. Haselager. 2012. Hierarchies in action and motor control. *Journal of Cognitive Neuroscience* 24(5): 1077–1086. doi:10.1162/jocn\_a\_00204.
- Uithol, S., D. Burnston, and W.F.G. Haselager. 2014. Why we may not find intentions in the brain. *Neuropsychologia* 56: 129–139. doi:10.1016/j.neuropsychologia.2014.01.010.
- van Gelder, T. 1998. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21: 615–665.
- Varela, F.J., E. Thompson, and E. Rosch. 1991. *The embodied mind: Cognitive science and human experience*. Cambridge: MIT Press.
- Vinding, M.C., M. Jensen, and M. Overgaard. 2014. Distinct electrophysiological potentials for intention in action and prior intention for action. *Cortex* 50: 86–99. doi:10.1016/j.cortex.2013.09.001.
- Vogelstein, J.T., Y. Park, T. Ohyama, R.A. Kerr, J.W. Truman, C.E. Priebe, and M. Zlatić. 2014. Discovery of brainwide neural-behavioral maps via multiscale unsupervised structure learning. *Science* 344(6182): 386–392. doi:10.1126/science.1250298.
- Wegner, D.M. 2003. *The illusion of conscious will*. Cambridge: The MIT Press.
- Wong, K.-F., and X.-J. Wang. 2006. A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience* 26(4): 1314–1328. doi:10.1523/JNEUROSCI.3733-05.2006.