



HAL
open science

Time to act: The dynamics of agentic experiences

Elisabeth Pacherie

► **To cite this version:**

Elisabeth Pacherie. Time to act: The dynamics of agentic experiences. Patrick Haggard; Baruch Eitam. The Sense of Agency: Social Cognition and Social Neuroscience, Oxford University Press, 2015, 9780190267278. hal-01110286

HAL Id: hal-01110286

<https://hal.science/hal-01110286>

Submitted on 5 Mar 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

published in Baruch Eitam & Patrick Haggard (eds), The sense of Agency: Social cognition and social neuroscience. Oxford: Oxford University Press, 2015

Time to act: The dynamics of agentic experiences

Elisabeth Pacherie

Institut Jean Nicod, Ecole Normale Supérieure, Paris

This paper is dedicated to the memory of Marc Jeannerod.

Abstract and Keywords

Actions unfold in time, and so do experiences of agency. Yet, despite the recent surge of interest in the sense of agency among both philosophers and cognitive scientists, the import of the fact that agentic experiences unfold in time remains to this day largely under-appreciated. I argue that we should think of agentic experiences as continuants, whose contents evolve as actions unfold. I try to characterize these content-shifts, distinguishing two main dimensions of change, changes in scale or fine-grainedness and changes in tense, as well as the main action control and action specification processes that underlie them. I further argue that taking into account this temporal dynamics of agentic experiences can also help us better appreciate in what sense some of the apparently conflicting empirical models of the sense of agency proposed in recent years can be seen as complementary rather than as rival and help us refine integrative models.

Keywords: sense of agency, agentic contents, action hierarchy, temporal dynamics; content shifts.

1. Introduction

Even the briefest of actions, such as pressing a button, flexing the wrist or raising an arm, takes time. Before one even acts, one has to decide what to do and when. When the time to act has come – which may be immediately following the decision, as when one acts on the spur of the moment –, and provided one recognizes that it has come, the action starts and then progresses until it either reaches successful completion, or, as the case may be, is aborted or ends in failure. Obviously then, actions unfold in time, but so do experiences of agency or, as I shall call them, agentive experiences.

Despite the recent surge of interest in the sense of agency among both philosophers and cognitive scientists, the import of the fact that agentive experiences unfold in time remains to this day largely under-appreciated. This is not to say that the temporal characteristics of agentive experiences have been completely neglected. Indeed, one of the main findings of experimental investigations into the sense of agency is the phenomenon of intentional binding, first reported by Patrick Haggard and his colleagues (Haggard, Clark & Kalogeras, 2002), whereby a voluntary action and its external sensory consequences are compressed together in subjective time. The effect of voluntary action on subjective time perception isn't however what I'll be primarily interested in. Rather my main concern will be with how the contents and the strength of agentive experiences evolve with time. While I will suggest a number of ways in which the contents of agentive experiences can shift as an action unfolds and a number of factors that can influence these shifts, I offer these as starting points for future research. My hope is to convince cognitive scientists that these issues are important, that many questions are yet unanswered and that they are worth investigating empirically.

I start with a preliminary characterization of agentive experiences in section 2. I then present a brief survey of recent empirical work on the sense of agency (section 3). In section 4, I discuss the contents of agentive experiences and try to characterize the two main

dimensions along which these contents shift as the action unfold, namely changes in scale or level of grain and changes in tense, as well as the factors that influence these shifts. In section 5, I argue that taking into account the time dynamics of agentic models can also help us better appreciate in what sense some of the apparently conflicting empirical models of the sense of agency proposed in recent years can be seen as complementary rather than as rival.¹

2. Agentic Experiences

States of agentic self-awareness comprise both agentic beliefs and agentic experiences (Bayne & Pacherie, 2007). *Agentic experience*—that is, our moment-by-moment sense of ourselves as performing various actions—constitutes our basic form of agentic awareness. Our agentic beliefs, while often based on our agentic experiences, can also be based on more indirect evidence. For instance, if I find my car keys in the fridge, I may form the belief that I was the one who put them there, despite have no recollection of having done so. In this case my agentic belief will be based on other doxastic states, such as my belief that I was alone at home and perhaps my further belief that I have a pronounced tendency towards absent-mindedness, rather on any specific agentic experience or memory thereof.

Here, my focus will be on agentic experiences. The burgeoning literature on the content and sources of agentic experiences highlights their many facets. A non-exhaustive list of proposed distinctions includes experience of deliberation, experience of decision, awareness of a goal, awareness of an intention to act, awareness of initiation of action, awareness of movement, sense of activity, sense of mental effort, sense of physical effort, sense of control, experience of authorship, experience of intentionality, experience of purposiveness, experience of freedom, and experience of mental causation. To avoid becoming overwhelmed by this terminological profusion, it may be useful, following Bayne (2010), to distinguish between core and non-core elements of agentic experiences. Core elements are elements that

must be possessed by any agentive experiences whatsoever and that are essential to agentive experiences, and non-core elements are elements that may but need not be present within experiences of agency.

Among non-core elements of agentive experiences, one may include, among other things, experience of effort, experience of deliberation or experience of decision. Arguably, not all actions are effortful or felt as such. Indeed, one finds reports in the clinical literature of people who after brain damage have lost the capacity to experience actions as effortful but still enjoy agentive experiences (Naccache et al., 2005). It is also not uncommon for one to act on the spur of the moment, or for actions, especially routine actions, to be triggered more or less automatically by cues in the environment. We can have agentive experiences for these actions, but these won't include experiences of deliberation or of decision. This is often the case for routine actions that can be triggered more or less automatically. For instance, I found myself a few minutes ago cleaning my glasses. As soon as I became aware of what I was doing, I had an agentive experience for the action I was performing but it did not include among its elements an experience of deliberating whether or not to clean my glasses or even deciding to do so.

Other elements, however, are essential to agentive experiences and in what follows I will concentrate on these core elements. At the very minimum, to have an agentive experience is to be aware of oneself as acting, that is to experience a sense of agency, narrowly conceived. Note that in saying that the sense of agency narrowly conceived is a core component of agentive experiences, I am not implying that this core component is itself necessarily atomic. Indeed, as we will see in the next section, it is often proposed that the sense of agency is itself decomposable into at least a sense of initiation and a sense of control.

It is difficult, however, to conceive of a sense of agency or a feeling of doing as existing independently of any experience of what it is one is doing, as a feeling simply floating in the air so to speak. It seems that an agentic experience also minimally involves some sense of what it is one is doing, whether it be an experience of moving one's body in a certain way (e.g., moving a finger), of performing an action of particular kind (e.g., pressing a key) or of acting with a certain goal (e.g., starting a new paragraph) Having an agentic experience does not require one to identify the action the experience is about at a particular level of specification, but it demands that the action be specified in some way, however vague the specification is. So I take it that agentic experiences include two core elements: a sense of agency and some specification of the action the sense of agency is for.²

It may be objected that in routine actions one typically has at least a minimal awareness of oneself as acting, but not an experience of what one is doing.³ In answering this objection, it is first important to note that the category of routine actions is a mixed bag. Actions such as playing with a hair lock or tapping one's foot while attending a lecture are sometimes classified as routine actions, but so is driving (at least for experienced drivers) or coffee making if you do it every morning. Routine actions of the first kind are typically involuntary and also unconscious in the sense that we lack even a minimal experience of ourselves as acting. Consider, for instance, the following dialogue:

- Could you please stop that?
- Stop what?
- This constant tapping with your pen on the table
- Oh, I'm sorry. I had no idea I was doing that.

One should also note that lacking an experience of acting is not tantamount to having an experience of passivity. An alternative to both is the lack of any experience. Thus, the

objection does not seem to apply to this first kind of routine actions, if, as I suggest, we are simply unaware that we are acting when performing them, and thus lack both a sense of agency *stricto sensu* and a sense of what it is we are doing.

Now, does the objection hold for routine actions of the second kind, actions that are voluntary but that one can perform without having to attend to what one is doing? The key word here is attention. There is a complex ongoing debate on the relationship between attention and consciousness, with some holding there is no consciousness without attention (O'Regan & Noe, 2001; Posner, 1994) and others holding that the two processes can be dissociated and thus that there can be conscious awareness without attention (Lamme, 2003; Koch & Tsuchiya, 2007, Wyart & Tallon-Baudry, 2008).⁴ I am not taking a position on this debate here, but what I want to do is sketch two stories about what might be going on in routine actions that preserve the two core-component view of agentic awareness.

The first story is meant to be compatible with the view that there can be some conscious awareness without attention. This conscious awareness would be marginal however and lack depth. For instance, when I am in the subway reading the newspaper, I am marginally aware of people around me, but my awareness of them is not such that I could recognize them if I was to meet them later elsewhere. Similarly, it may be that when I am performing a routine action I have some marginal awareness of myself as acting and also some marginal awareness of what I am doing, but that I consciously access only some general features of the action representations that drive my action and have therefore an agentic experience of what I am doing that is largely underspecified.

The second story is more radical and is meant to be compatible with the no-consciousness-without-attention view. Here's how it goes. When one is performing a routine action without any attentional resources being engaged, one has no agentic experience for

that action, not even a minimal sense of oneself as acting. However, (some of) the action representations that drive the action are 'poised' for consciousness. As soon as, for whatever reason, attention becomes engaged, one starts enjoying some awareness of oneself as acting and becomes aware of at least some aspects of what one is doing. On this story, the idea that during the course of routine actions we always enjoy some sense of agency is an illusion. It is, so to speak, an instance of the fridge light illusion. In the same way that we may come to mistakenly believe that the light is always on in the fridge because it is always on when we open the fridge, we may come to mistakenly believe that we always enjoy an experience of acting for routine action, because each time we attend to the action, we have the experience. This illusion would be reinforced by the fact that it is rarely the case that we deploy no attentional resources whatsoever while performing a routine action. Even though making coffee is something I do every morning, I still have to pay attention to how much water I pour in the coffee maker, how I position the cup, and so on. According to this second story, the question to ask is not at how many (and at which) of the stages of action an ongoing agentive experience must include an experience of what one is doing, but rather at how many (and at which) of the stages of action there must be an agentive experience.

Which of these two stories, if any, is correct remains at present an open and a difficult empirical question. My point here is simply that both stories can accommodate the claim that an agentive experience does not reduce to a pure, or naked, sense of agency but also involves some awareness of what one is doing, however partial and unspecific the content of this conscious representation. Both stories suggest that when you interrupt a routine action of mine and ask me what I am doing, I may give an answer by reporting the contents of the action representations that were active just before you interrupted me and that remain active and thus can be attended to precisely because you interrupted me and thus prevented their being erased by corresponding reafferences. In addition, the consciousness-without-attention

view allows for some marginal sense of agency and some marginal awareness of what one is doing even when no attentional resources are engaged; on this view, the deployment of attentional resources simply transforms and enriches our agentic experience, making it more vivid and allowing us to consciously access richer and more detailed representations of what we are doing. In contrast, on the no-consciousness-without-attention view, unless attentional resources are deployed, we simply have no agentic experience whatsoever: we experience no sense of agency and are not aware of what we are doing. However, when attentional resources are deployed we enjoy both a sense of agency and a sense of what we are doing and the more attentional resources are deployed the richer and more detailed the contents of our conscious experiences become.

3. The Sources of Agentic Experiences

Empirical research on agency has explored a variety of potential cues to agency, and a number of different models of how agentic experiences are generated have been proposed. Typically, these models focus more on core-elements of agentic experiences than on non-core elements, and, within these core elements, also tend to concentrate more on the sense of agency narrowly conceived than on how actions are represented in agentic experiences. These models all rely to a greater or lesser extent on a congruence principle: the sense of agency is produced when there is a match between cues x and y . The points on which they differ concern the nature of the cues being compared, the nature of the processes involved in the production of the sense of agency, and how closely these processes are related to action production and control processes.

Two theoretical positions define the two ends of the spectrum of possibilities: the motor prediction view and the cognitive reconstruction view. On the motor prediction view, the sense of agency is generated by processes dedicated to action control. On the cognitive

reconstruction view, the sense of agency is generated by general-purpose processes of retrospective causal inference.

The motor prediction view is inspired by computational theories of motor control. According to these theories, when the motor system generates a motor command, an efference copy of this command is sent to forward models whose role is to generate predictions about its sensory consequences in advance of actual execution. Error signals arising from the comparison of desired, predicted, and actual states (as estimated from sensory reafferences) are used to make corrections and adjustments. The motor prediction view holds that the signals used for motor control also provide cues to agency (Frith *et al.* 2000). In particular, it holds (1) that awareness of initiating an action is based on a representation of the predicted consequences of making that action, rather than its actual consequences, and on the congruence between the predicted state and the desired state, and (2) that for this experience of agency to continue, the predicted consequences would also have to remain congruent with the sensory reafferences when they become available. These two claims are associated with the idea that the sense of agency has two main components: the sense of initiation linked to the congruence between predicted and desired state and the sense of control linked to the degree of congruence between predicted and actual state. Indeed, more recent empirical work within the motor prediction framework suggests a more complex picture of the determinants of the sense of agency, highlighting in particular the contribution of action selection fluency, and thus potentially of metacognitive processes, to the sense of agency (Chambon *et al.*, 2014; Chambon & Haggard, 2012; Sidarus *et al.*, 2013; Wenke *et al.*, 2010).

Claim (1)—and therefore the possibility that the sense of agency can emerge in advance of actual sensory effect and be based on premotor processes alone—is supported by evidence that awareness of initiating a movement in healthy subjects is reported by the agent between 80-200 milliseconds before the movement actually occurs (Libet *et al.* 1983; Libet 1985;

Haggard & Eimer 1999) or, in the case of patients with anosognosia for hemiplegia, can be reported despite no movement being actually produced (Berti et al., 2007). Evidence for claim (2)—that the sense of agency also depends on the congruence between predictions and sensory reafferences—comes from studies where these reafferences are artificially manipulated by introducing temporal delays and spatial distortions of feedback. These studies demonstrate that the sense of agency is gradually reduced as these discrepancies increase (Fournieret & Jeannerod 1998; Knoblich & Kircher 2004; Sato & Yasuda 2005).

In contrast, the cognitive reconstruction view downplays the contribution of the motor system to the sense of agency and proposes that it is inferred retrospectively from the existence of a match between a prior thought and an observed action. Thus, on Wegner's 'theory of apparent mental causation' (Wegner 2002), a general-purpose causal inference process is at play. If an action is consistent with a prior thought of the agent and other potential causes of the action are not present or salient, a sense of agency for the action will be induced.

There is also empirical evidence that high-level inferential processes play a role in determining the sense of agency for an action. Studies of Wegner and colleagues have demonstrated that cognitive cues can alter the sense of agency for an action independently of changes in sensorimotor and perceptual cues. For instance, in their 'I-Spy' study (Wegner & Wheatley 1999), a participant and a confederate of the experimenter had joint control of a computer mouse that could be moved over any one of a number of pictures on a screen. When participants had been primed with the name of an item on which the mouse landed, they expressed a stronger sense of agency for the action of stopping on that object (when in fact the stop had been forced by the confederate). Further studies also suggest that subliminally priming an outcome just before the outcome is produced can enhance the sense of agency for that outcome (Aarts *et al.* 2005) and that priming an outcome relatively far in advance can

augment self-agency as well, but only if the outcome is attached to positive affect (Aarts *et al.* 2009).

There is now a growing consensus that the motor prediction view and the cognitive reconstruction view are not mutually exclusive but complementary and that intrinsic cues (cues provided by the motor system) and extrinsic cues (such as cognitive primes) both contribute to the sense of agency (Pacherie 2008, 2010; Sato 2009; Synofzik *et al.* 2008; Moore *et al.* 2009; Moore & Fletcher 2012). Researchers are now trying to develop integrative frameworks and to get a better understanding of how all these agency cues interact.

One way to try and combine the motor prediction view and the cognitive reconstruction view is to appeal to the distinction between pre-reflective agentic experiences and reflective agentic beliefs or judgments (Bayne & Pacherie 2007; Gallagher 2007; Haggard & Tsakiris 2009) and to argue that while motor processes contribute mainly to agentic experiences, interpretive and general inferential processes contribute mainly to agentic beliefs or judgments. Thus, the example given at the beginning of section 2 of my finding my car keys in the fridge and judging that I must have put them there myself may constitute an instance where an agentic judgement is not based at all on an agentic experience but is the result of inferential processes exploiting general causal knowledge about the world. This, I know that car keys don't move on their own and that no one else was around, leading me to conclude that I must have done it. Typically though, we may expect judgments to be based at least in part on agentic experiences when such experiences are available.

This conceptual distinction between agentic experiences and judgements is echoed methodologically in the ways agency is measured in experimental studies. While some studies (Farrer *et al.* 2003; Metcalfe & Greene, 2007; Sato & Yasuda 2005) investigate agency by

asking participants to explicitly judge whether they caused a particular sensory event, other studies use implicit agency measures such as intentional binding and sensory suppression. Intentional binding is a phenomenon, first reported by Haggard and his colleagues (Haggard *et al.* 2002), whereby an action and its external sensory consequences are compressed together in subjective time. As intentional binding occurs only for voluntary actions (Tsakiris & Haggard 2003) and is furthermore modulated by the statistical relation between events (Moore & Haggard 2008), it is considered to provide an implicit measure of agency. Sensory attenuation of self-produced action effects has also been used as an implicit measure of agency. When the internally generated motor predictions about the sensory consequences of one's on-going actions and their actual sensory consequences are congruent, the sensory percept is attenuated, thereby enabling a differentiation between self-generated and externally generated sensory events (Blakemore *et al.* 2000). However, recent studies showing that prior authorships beliefs can modulate both sensory attenuation and intentional binding (Desantis *et al.* 2011; Desantis *et al.* 2012) suggest that drawing a sharp distinction between agentic experiences supported by motor processes and agentic beliefs supported by interpretive processes may be over-simplistic.

A promising approach is to appeal to a Bayesian integrative framework involving a hierarchy of prediction and model building. Thus, Moore & Fletcher (2012) propose that the sense of agency is determined by a Bayesian process of cue integration, where the predictions generated at higher levels of the hierarchy provide the priors for the lower levels; i.e. constrain the interpretation of cues available at lower levels. In this model, cue integration is itself the product of both the strength of the priors and the weights attached to the available cues as a function of their reliability. When priors are weak—as, for example, when one is quite unsure what the effects of pressing a button will be—one may still have a strong sense of agency for the ensuing consequence, provided that perceptual reafferences carrying

information about it are very reliable. Conversely, if my priors are very robust, I may have a strong sense that I produced a certain effect in the world, even though the feedback I get is weak or ambiguous. When both priors and reafferent cues are weak, my sense of agency may be correspondingly weakened. While this Bayesian approach does not allow for a sharp distinction between agentic experiences and agentic judgments, it can accommodate the idea that high-level priors exert more influence on agentic judgments than on agentic experiences. As we will see in section 5, another important parameter we must take into account when trying to understand how cue integration works is time and, in particular, the different times at which different cues become available.

4. Agentic Experience: Content Shifts

I proposed in section 2 that agentic experiences include two core elements: a sense of agency and some specification of the action the sense of agency is for. I pointed out, however, that having an agentic experience does not require one to identify the action the experience is about at a particular level of specification. Not imposing such a requirement makes sense if agentic experiences evolve in time and if their intentional contents change as actions unfold. My aim in this section is to characterize the kinds of changes the intentional contents of agentic experiences can undergo, to explore the ways in which these changes relate to the hierarchical structure of action specification and control processes, and to discuss additional factors that could also modulate which action representations enter the contents of agentic experiences.

Contents-shifts in agentic experiences can involve both changes in tense and changes in the level of grain at which the action is identified. With respect to tense, one can have an experience of being about to *A*, of initiating *A*, of being in the middle of *A*-ing, of nearing the completion of *A*-ing and of having just *A*-ed. We should expect that, as an action unfolds,

there should be tense-shifts in the contents of the agentive experience that accompanies it. Additionally, if one considers actions somewhat more complex than simply flexing a wrist or raising an arm, it becomes clear that multiple tenses can in principle co-exist within the same agentive experience. Suppose that I am preparing a chocolate mousse, my agentive experience at a given moment may be of having melted the chocolate and being about to whip the egg whites. A second form of content-shift concerns the level of grain at which the action is identified. As the chocolate mousse example can also serve to illustrate, I can experience my action as preparing a chocolate mousse, whipping the eggs white, rotating my wrist with great energy, and so on. More generally, it seems that we can zoom in and out as we perform an action, focusing at some points on specific details of action execution and at other points of more abstract aspects of our action plan.

All of this sounds commonsensical enough, or so I hope. But can we say more about what makes one identify an action at a given level of grain, or what makes one concentrate on the past, the present or the future aspects of the action being performed? As I remarked earlier, empirical work on agentive experiences tends to be more concerned with investigating the cues and the processes that generate the sense of agency than with studying how actions are represented in agentive experiences. Nevertheless, it is quite plausible that the various cues that contribute to the sense of agency also contribute to an agent's awareness of what he or she is doing. As the discussion in section 3 indicates, a variety of agency cues contribute to the sense of agency and, as hierarchical cue integration models suggest, different cues may pertain to different levels in a hierarchy of action representations and predictions. One early proponent of this idea was Marc Jeannerod (1995, 1997) who captured it in the form of a flow chart model, reproduced as Figure 1. In what follows, I propose to use this model as a guide to address the following questions: What are the action representations available at given stages of action preparation and execution? When should we expect an agentive experience to

represent the whole action? When should we expect it to represent a particular component of an action? When should we expect transitions from one component to another or from one level of representation to another (zooming in and zooming out) to take place?

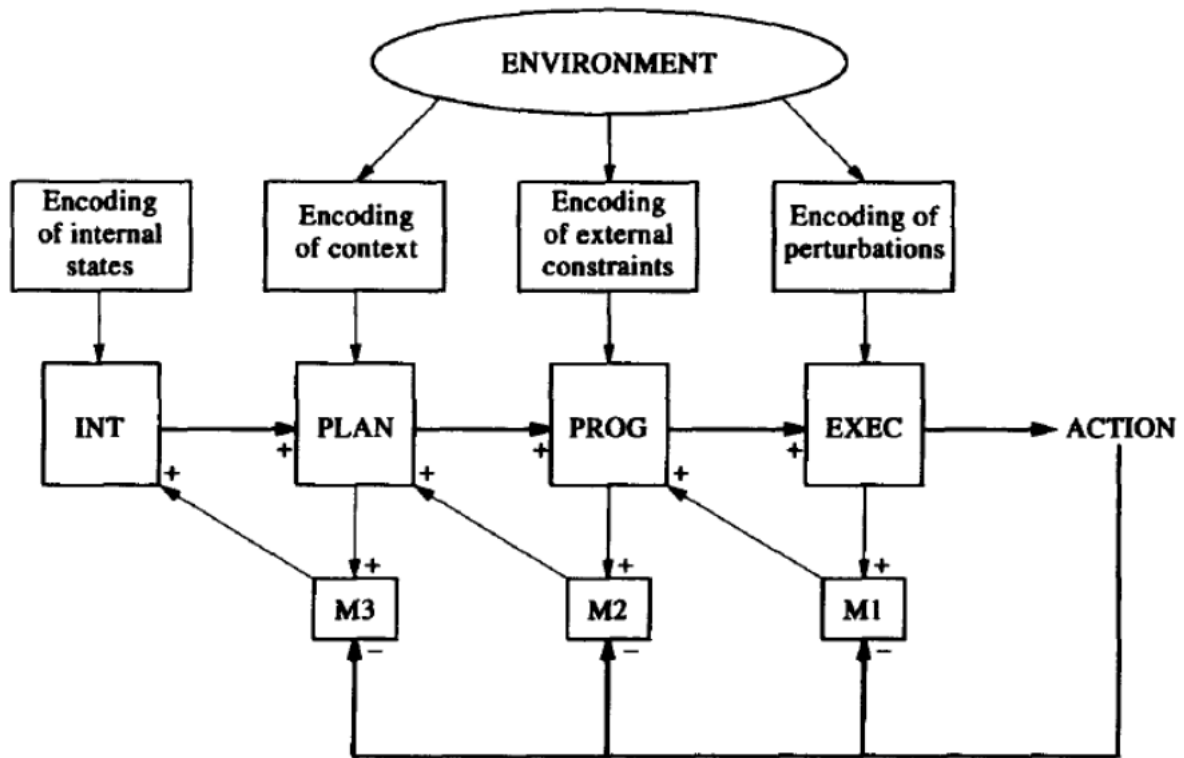


Figure 1. A flow-chart model of the ongoing regulation of action. Reproduced from Jeannerod, 1995.

In his comments on this flowchart, Jeannerod (1995, 1997) makes several important points. First, while the main flow of information goes through the cognitive steps of action specification (e.g., intending, planning, programming and executing), Jeannerod notes that the impression given by this flowchart representation that these steps are organized serially must be corrected. In other words, while the initial component or segment of an action is prepared an executed, the preparation of the next components of the action is already engaged. As he puts it, 'the diagram, at least for what concerns its distal steps, should be considered as one of many parallel lines with a similar structure' (1997: 172: Second, the predictions made are

stored in as many short-term memories (represented as M1, M2 and M3 in Figure 2) as there are levels, and reafferences documenting the current state of the action are later fed into the memories. If the action has been completed, the reafferences and the content of the memories will coincide, and the latter will be erased. If the action is incomplete, the memories remain active, which reactivates the corresponding module and generates corrections. While Jeannerod himself doesn't use this terminology, the mechanism he proposes is akin to the comparison processes appealed to in the motor prediction view discussed in section 3.

The idea of a set of short-term memories storing predictions and remaining activated until matching reafferences erase their contents gives us a way of understanding which representations of an action are active at a given moment and can thus potentially enter the contents of agentic experiences. This model suggests certain constraints on which action representations can contribute to agentic experience at a given moment, and suggest more generally that our agentic experiences should tend to be present and future oriented, given that short-term memories for actions or action components already completed are promptly erased. Yet, it still implies that many representations of an action can be simultaneously active, first because, as Jeannerod puts it, in a given line, corresponding to the preparation of a given component or segment of an action, several memories storing representations of this segment at different levels of specification can be active, and second, because this line is only one of possibly many such lines where memories are also currently active. Yet, it seems unlikely that all the action representations active at a given time contribute, or at least contribute equally, to agentic experiences.

Suppose action representations are, as the model suggests, arranged in a hierarchy with the more abstract representation of the action at the top of the action tree (e.g., making coffee), representations of component actions (e.g., filling the water reservoir of the coffee machine, installing a new filter in the filter basket, putting coffee in the filter, turning the

machine on, etc.) at the next level, representations of subcomponents (e.g., fill coffee pot with required amount of water, open water reservoir, pour the water from the coffee pot to the reservoir, etc.) one further level down, then representations specifying the exact effectors used (e.g., open water reservoir with right hand, and pour water with left hand, or vice-versa), and so on until one reaches the level of basic motor actions. One prediction we can make is that while a number of different, more or less specific, action representations would be active at any point during the action, the most active at a given point would be those that are most directly relevant in guiding the transition from the action segment currently being performed to the next. Using tree terminology, the representations most relevant to guiding the transition from the current terminal node of the action tree to the next adjacent terminal node are those attached to the closest common ancestor node of these two terminal nodes. If these two nodes happen to belong to the same sub-branch of the tree, the representation at this common ancestor node should remain relatively specific, if the two nodes are on two different sub-branches, they will have as a common ancestor a more abstract action representation. In other words, one would predict that the abstract representation "preparing coffee" would be more readily available when one has put coffee in the filter and is about to turn the machine on, then when one is opening the water reservoir and about to pour water in it. One way to test such predictions, and thus to explore to what extent our awareness of what we are doing is a reflection of how we navigate a tree structure of action representations, would be to have participants perform such action scripts, interrupt them at various junctures and ask them to answer a yes/no question as quickly as possible (e.g., "are you making coffee?", "are you filling the water reservoir?", "are you raising your right arm?"). If we assume that they would be faster at answering a question when its content corresponds to how the content of their agentive experience at the time of the interruption, their reaction time would presumably give

us an indication of what the contents of agentic experience are and how they shift at various points during action execution.

Several other factors can further influence which action representations enter the contents of agentic experiences. First, while, as I just suggested, we normally attend to what we must do next and navigate our action representation tree accordingly, what we are attending to can also be modulated by endogenous and exogenous factors. On the one hand, we can voluntarily allocate attention to some aspects of an action more than to others. For instance, when I play a serve in tennis, I can choose to attend to distal aspects of the action, (e.g., where in the court the ball lands), or to more proximal aspects of the action (e.g., how I move my arm and shoulder). On the other hand, stimulus events can also involuntarily capture attention. The short-term memories discussed by Jeannerod are short-lived, since they are erased as soon as matching reafferences become available. The longer they last, the more likely it is that they will capture attention, since their longevity is a sign of trouble.

One further factor that may modulate the content of our agentic experiences is the agent's preferred level of action identification. According to Action Identification Theory (Vallacher & Wegner, 1985, 1989), any behavior can be represented at multiple levels but there are individual differences in agents' preferred level of action identification. Some people show a general tendency to (consciously) represent their behavior at a low level, i.e., in terms of the concrete, mechanistic aspects of their action (e.g. turning a doorknob), whereas other people show a general tendency to represent their behavior at a higher level, i.e. in terms of the goals or purposes of their action (e.g., opening the door or going out: While other factors can also influence the level at which an action is represented, still, *ceteris paribus*, agents whose preferred level of action identification is low will tend to experience their action in terms of producing sensorimotor consequences or outcomes, while agents whose preferred level of

action identification is high will tend to tend to experience their action in terms of the goal it serves.⁵

Yet another factor is the level of skill of the agent. When an action is overlearned and has become routinized, the transition between action segments or components can proceed more or less automatically. If, however, the agent is not performing a routine, he or she may have to consciously consider what comes next when a given segment is completed. However, it can be argued that what characterizes expert performance is not so much that it can proceed automatically, without the agent needing to pay attention to what he is doing or to transitions between action segments, as the fact that the agent can flexibly navigate between levels of action representations, attend to global aspects of his performance as well as to details of execution. On the one hand, there is evidence that highly skilled performers are able to attend to details of motor execution that remain inaccessible to less proficient performers (Knoblich & Repp, 2009). On the other hand, while less proficient performers are often unable to execute certain actions without consciously attending to what they are doing, expert performers may be able to perform them without mobilizing attentional resources, focusing their attention instead on more global or distal aspects of their performance. Thus, the expert piano player may be able to concentrate her attention on the phrasing of a passage, while all the attention of the novice player is concentrated on how to play the next chord. Expertise is thus characterized both by the range of action representations the expert can consciously access, from low-level sensorimotor representations to much more global and abstract representations of his action, and by the flexibility with which he or she can navigate between representations at various levels. While the less proficient agent may be forced to attend to representations at a certain level, the expert can choose what to attend to.

It is a controversial issue whether low-level sensorimotor representations can contribute anything to agentive experiences, given that detailed sensorimotor specifications of actions

and sensorimotor reafferences may both remain below the threshold of awareness. This is to a large degree an empirical question and I therefore limit myself to some tentative remarks. For a start, the question itself is in need of clarification. Let me throw in three distinctions. First, we can distinguish between contributions to the core components of agentic experiences and contributions to their non-core components. Second, we can also distinguish between contributions to one or the other of the two core components of agentic experiences: sense of agency and sense of what it is one is doing. Third, we can distinguish between direct and indirect contributions to the contents of agentic experiences, where by direct contributions I mean that an action representation contributes (some of) its contents to the intentional content of agentic experiences.

Presumably, the version of the question that is most controversial is whether low-level sensorimotor representations can make a direct contribution to the intentional content of an agentic experience. For instance, there is experimental evidence that when performing a target reaching task subjects can make rapid corrections of their hand movements in response to biased feedback on their hand trajectory without being aware that they are making these corrections (Fournieret & Jeannerod, 2008). These results suggest that the motor system constructs precise sensorimotor representations, since small discrepancies between predicted and actual sensorimotor reafferences are detected and used to make appropriate corrections of the hand movement, but that these sensorimotor representations are not consciously accessed, since subjects are not aware of these discrepancies and not aware that they are making corrections. However, evidence from expert performance suggests that highly skilled performers acquire very specific sensorimotor mappings and may be able to attend to details of execution and to access sensorimotor representations that remain inaccessible to less skilled agents (Knoblich & Repp, 2009). Whether or not certain sensorimotor representations directly contribute to our awareness of what we are doing may thus depend on the agent's

level of skill or expertise. This is not to say however, that experts have unlimited access to their sensorimotor representations. Distinguishing between high-level cognitive representations, mid-level perceptual representations and low-level sensorimotor representations is probably a simplification as each of these three layers of action representations is likely to contain many sub-layers. It may thus remain true that even for experts, very low-level sensorimotor representation remain below the threshold of awareness and thus make no contributions to our awareness of what we are doing.

With regard to the possible indirect contribution of sensorimotor representations to the contents of agentic experiences, it seems plausible that prediction errors arising from mismatches between sensorimotor predictions and sensorimotor reafferences could contribute to attentional amplification at the next level and thus to an increased awareness of what is going on as it is represented at the next level. In such a scenario, low-level sensorimotor predictions and reafferences would have an effect on how rich the contents of an agentic experiences would be not by contributing their own contents but because their comparison generates prediction errors that trigger the deployment of attention resources one level upwards. Finally, there are two further contributions these sensorimotor representations may make to agentic experiences. First, it may also be that the number and strength of the error signals generated at the sensorimotor level modulate to some extent the sense of agency for an action. Agents may not have conscious access to what went wrong at the sensorimotor level, yet have a sense that something was not quite right. Second, and perhaps relatedly, the absence or near absence of error signals at that level may also contribute to the sense of flow that can accompany flawlessly executed actions.

To recap, hierarchical models of action organization suggest that action is guided and controlled by multiple action representations, representing different components of the actions and representing them at various levels of grain and that many such representations can be

activated in parallel. Following Jeannerod, I suggested that the action representations that contribute to the contents of agentic experiences should be predominantly forward-looking, that is, be mostly representations of what the agent is currently doing or about to do. I also suggested that the action representations most active at a given point should be those that are most directly relevant in guiding the transition from the action segment currently being performed to the next and that their degree of abstractness should depend on whether these two segments belong or not to the same sub-branch of the action representation tree. In other words, the changing contents of our agentic experiences should reflect the way in which we navigate an action representation tree when performing an action, zooming in when performing a particular component and zooming out at transition points between components. This idea is in principle empirically testable and I mentioned one way it could be tested. However, other factors may also modulate the contents of agentic awareness, including endogenous and exogenous factors affecting what the agent is attending to, the agent's preferred level of action identification, and her degree of expertise.

5. Sense of Agency: Variations in sources and strength

In section 2, I distinguished two core components of agentic experience: awareness of oneself as acting – i.e., the sense of agency narrowly conceived – and awareness of what one is doing – i.e., the intentional content of agentic experience. In the previous section, I discussed shifts in the intentional contents of agentic experiences. In what follows, I come back to the sense of agency.

First, as I mentioned at the outset of section 3, empirical research on agentic experiences has concentrated more on the sense of agency *stricto sensu* than on the content of agentic experiences. This is not to say that it has completely neglected content. For instance, Haggard & Eimer (1999) asked participants to choose voluntarily between left and right key-press

actions in a modified version of Libet's task and showed that awareness of initiating an action correlated not with the general readiness potential but covaried with the lateralized readiness potential (the later phase of preparation, in which brain activity contralateral to the selected hand exceeds ipsilateral activity), suggesting that the awareness of initiating an action is tied to the preparation of a specific body movement rather than to a general preparation to move. Experimental studies have also examined the effects of action identity manipulations on intentional binding and sensory attenuation (for a review, see Hughes et al., 2013). However, existing experimental studies typically manipulate action representations in order to assess the effects of these manipulations on the sense of agency narrowly conceived rather than to investigate their impact on the contents of agentic experiences.

Second, as the discussion in section 3 also indicated, one important area of disagreement in empirical research concerns the nature of the processes underpinning the sense of agency: Are these processes essentially predictive or postdictive? Are the cues or representations they exploit rather low-level or high-level? While a consensus seems to be emerging that both kinds of processes may be at play and that a variety of lower and higher level agency cues contribute to the sense of agency, the integrative frameworks that have been proposed so far still lack sufficient specificity. In particular, they do not sufficiently take into account the temporal dimension of the phenomenon at stake and thus fail to avail themselves of one important source of constraint. For instance, while Bayesian hierarchical frameworks propose that the sense of agency is the product of the interaction between sensory or perceptual (bottom-up) evidence, weighted according to its reliability, and prior (top-down) expectations weighted according to their strength, they do not explicitly take into account the temporal properties of these different sources of information and therefore the way the sense of agency and its sources may evolve as the action unfolds.

My main point here is very simple: when thinking of the sense of agency as determined by a Bayesian process of cue integration, one shouldn't lose sight of the fact that these cues are not all available at the same time. Both higher and lower level priors and predictions based on them are available from the start while feedback and signals arising from the comparison of feedback with predictions only become available at later stages. We should therefore expect our sense of agency to be mostly based on predictive processes and prospective cues in the early stages of an action with a progressive shift towards postdictive processes and retrospective cues as the action unfolds. We should also expect our sense of agency to fluctuate as the action unfolds, according to the congruence or lack thereof between prospective and retrospective cues and to their respective strength. For instance, our sense of agency may initially be based on weak priors and be correspondingly weak but increase as reafferences that match our expectations become available. Conversely, our sense of agency may initially be based on strong priors and thus be strong but decrease when we receive reliable incongruent feedback. As proposed by Moore and Haggard (2008), retrospective inferential processes may kick in when predictions are invalidated by reafferences, in order to ensure that our agentic experience is a correct reflection of the actual relations between our actions and their consequences. What these retrospective inferential processes do, in such a case, is tie efferent signals to their actual effects in the world. In so doing they help us regain a sense of agency over our action by changing the contents of our agentic experience and improving their accuracy.

Moore and Haggard's proposal also highlights one way in which the two core components of agentic experience, their contents and their strength, can interact. This interaction could also be manifested in other ways. I argued in section 4 that the intentional content of an agentic experience changes as the action unfolds and that such changes may reflect the way in which we navigate an action representation tree when performing an action. The changing

intentional focus of our agentic experience may also have an effect on our moment-to-moment sense of agency by modulating the weights attached to the various available agency cues. For instance, if, as I suggested, we zoom in and out as we perform an action, being centrally aware at some points of very specific aspects of action execution and at other points of more abstract aspects of our action plan, we may expect lower-level cues to be given more weight when zooming in and higher level cues to be given more weight when zooming out.

Conclusion

In a nutshell, I argued here that when attempting to characterize the core components of agentic experiences we need to pay closer attention to their contents and temporal dynamics. I tried to emphasize two dimensions along which the contents of an agentic experience can vary as the corresponding action unfolds in time. The first has to do with the level of grain or scale at which the ongoing action is consciously represented. Many representations at various levels of the action representation hierarchy could in principle contribute to the intentional contents of agentic experience. It is important that we gain a better understanding of what determines which of these action representations contribute to the contents of agentic experience at a given moment and how these contents may shift as the action unfolds. I considered several factors that could in principle influence what enters the contents of agentic awareness and suggested some ways in which these ideas could be empirically tested. I proposed that the changing contents of our agentic experiences should reflect the way in which we navigate an action representation hierarchy when performing an action, zooming in when performing a particular component and zooming out at transition points between components. The second dimension of variation in the content of an agentic experience concerns its tense, that is whether its content is forward looking ("I am about to jump"), backward looking ("I just jumped") or present-centered ("I am jumping"). Whether these two dimensions of variation are correlated or orthogonal is yet another issue awaiting

empirical scrutiny. My personal hunch is that there is some degree of correlation. If we take a simple action, we should probably expect that, on the scale dimension, we will initially zoom in and then zoom out and that, on the tense dimension, experiential contents should initially be forward looking but later backward looking. However, the picture is likely to become more complicated, when we consider complex, extended action sequences.

Regarding the other core component of agentic experiences, the sense of agency *stricto sensu*, we still lack a clear understanding of how the many agency cues that can in principle contribute to the sense of agency interact. I suggested that in order to improve our understanding of how these cues are integrated, we should take into account their temporal properties and their effects on the evolution of the sense of agency in the course of the action. Finally, I pointed out that sense of agency and intentional content are not two completely independent components of agentic experience. Rather, they interact and changes in one component may induce changes in the other.

Acknowledgments

The preparation of this paper was supported by ANR grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL*. Some of the ideas discussed here were presented at a Special Session on Action Consciousness at the Fourth Online Consciousness Conference, February 17th-March 2nd 2012 (<http://consciousnessonline.com/2012/02/17/special-session-on-action-consciousness/>) I wish to express my heartfelt thanks to Myrto Mylopoulos, the organizer of this special session, to my two commentators, John Michael and Markus Schlosser, and to all the participants in this session for their many insightful comments and suggestions.

Endnotes

1. Note that my concern is with agentic experiences for physical actions. Whether and how the account proposed here could be extended to cover mental actions as well are difficult issues on which I shall not pronounce myself here.
2. Bayne (2010) leaves it open whether the agentic core should take the richer form I suggest or can reduce to a bare experience of oneself as acting.
3. Thanks to Markus Schlosser for raising this objection.
4. See also Kouider et al. (2010) and Cohen & Dennett (2011) for discussions of the methodological difficulties involved in attempting to empirically confirm or falsify the claim that there can be consciousness without attention.
5. See for instance Van der Weiden et al. (2010) for a study of how the effects of priming on agentic experiences crucially depend on action identification level and Belayachi & van Linden (2009) for a discussion of the role of low-level action identification in checking behavior in subjects with obsessive-compulsive disorder.

References

- Aarts, H., Custers, R., & Marien, H. (2009). Priming and authorship ascription: When nonconscious goals turn into conscious experiences of self-agency. *Journal of Personality and Social Psychology*, 96, 967-979.
- Aarts, H., Custers, R., & Wegner, D. (2005). On the inference of personal authorship: Enhancing experienced agency by priming effect information. *Consciousness and Cognition*, 14, 3, 439-458.
- Bayne, T. (2008). The phenomenology of agency. *Philosophy Compass*, 3, 1, 182-202.

- Bayne, T. (2010). Agentive Experiences as Pushmi-Pullyu representations. In J. Aguilar, A. Buckareff and K. Frankish (eds) *New Waves in the Philosophy of Action*. Palgrave Macmillan, pp. 219-36.
- Bayne, T., & Pacherie, E. (2007). Narrators and comparators: The architecture of agentive self-awareness. *Synthese*, 159, 3, 475-491.
- Belayachi, S., & van der Linden, M. (2009). Level of agency in sub-clinical checking. *Consciousness and Cognition*, 18, 293–299.
- Berti, A., Spinazzola, L., Pia, L., Rabuffetti, M., Haggard, P., & Rossetti, Y. (2007). Motor awareness and motor intention in anosognosia for hemiplegia. *Sensorimotor foundations of higher cognition*, 163-181.
- Blakemore, S. J., Wolpert, D., & Frith, C. (2000). Why can't you tickle yourself? *Neuroreport*, 11, 11, R11-R16.
- Chambon, V., Filevich, E., & Haggard, P. (2014). What is the Human Sense of Agency, and is it Metacognitive?. In *The Cognitive Neuroscience of Metacognition* (pp. 321-342). Springer Berlin Heidelberg.
- Chambon, V., Haggard, P. (2012). Sense of control depends on fluency of action selection, not motor performance, *Cognition*, 125, 441-51.
- Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Science*, 15, 358-364.
- Desantis, A., Roussel, C., & Waszak, F. (2011). On the influence of causal beliefs on the feeling of agency. *Consciousness and Cognition*, 20(4): 1211-1220.
- Desantis, A., Weiss, C., Schutz-Bosbach, S., & Waszak, F. (2012). Believing and Perceiving: authorship belief modulates sensory attenuation. *PLoS ONE*, 7(5): e37959.

- Farrer, C., Franck, N., Georgieff, N., Frith, C.D., Decety, J., & Jeannerod, M. (2003). Modulating the experience of agency: A positron emission tomography study. *Neuroimage*, 18: 324–333.
- Fourneret, P., & Jeannerod, M. (1998). Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia*, 36, 11, 1133-1140.
- Frith, C., Blakemore, S., & Wolpert, D. (2000). Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 355, 1404, 1771-1788.
- Gallagher, S. (2007). The natural philosophy of agency. *Philosophy Compass*, 2, 2, 347-357.
- Haggard, P & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, 126: 128-133.
- Haggard, P., Clark, S., & Kalogeras, J. (2002) Voluntary action and conscious awareness. *Nature Neuroscience*, 5, 4, 382-385.
- Haggard, P. & Tsakiris, M. (2009). The experience of agency: feeling, judgment and responsibility. *Current Directions in Psychological Science*, 18, 4: 242-246.
- Hughes, G., Desantis, A., & Waszak, F. (2013). Mechanisms of intentional binding and sensory attenuation: The role of temporal prediction, temporal control, identity prediction, and motor prediction. *Psychological bulletin*, 139(1), 133.
- Jeannerod, M. (1995). Mental imagery in the motor context. *Neuropsychologia*, 33, 11: 1419-1432.
- Jeannerod, M. (1997). *The cognitive neuroscience of action*. Oxford: Blackwell Publishers.

- Knoblich, G., & Kircher, T. T. J. (2004). Deceiving oneself about being in control: Conscious detection of changes in visuomotor coupling. *Journal of Experimental Psychology- Human Perception and Performance*, 30, 4, 657-666.
- Knoblich, G., & Repp, B. H. (2009) Inferring agency from sound. *Cognition*, 111, 2, 248-262.
- Koch, C. and Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Science*, 11, 16–22.
- Kouider, S., de Gardelle, V. Sackur, J. & Dupoux, E. (2010). How rich is consciousness? *The partial awareness hypothesis*. *Trends in Cognitive Science*, 14, 301-307.
- Lamme, V.A. (2003). Why visual attention and awareness are different. *Trends in Cognitive Science*, 7, 12–18.
- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8: 529-566.
- Libet, B., Gleason, C. A., Wright, E. W. & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activities (readiness potential): the unconscious initiation of a freely voluntary act. *Brain*, 106: 623-642.
- Metcalf, J., & Greene, M. (2007). Metacognition of agency. *Journal of Experimental Psychology: General*, 136, 2, 184–199.
- Moore, J. W., & Fletcher, P. C. (2012). Sense of agency in health and disease: A review of cue integration approaches. *Consciousness and Cognition*, 21, 1, 68-59.
- Moore, J. W., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition*, 17, 136–144.
- Moore, J. W., Wegner, D. M., Haggard, P. (2009) Modulating the sense of agency with external cues. *Consciousness and Cognition*, 18, 1056-64.

- Naccache, L., Dehaene, S., Cohen, L., Habert, M.-O., Guichart-Gomez, E., Galanaud, D. & Willer, J.-C. (2005). Effortless control: executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia*, 43: 1318-1328.
- O'Regan, J.K. and Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–973.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107, 1, 179-217.
- Pacherie, E. (2010). Self-agency. *The Oxford handbook of the self*, 440-462, Oxford University Press.
- Posner, M.I. (1994). Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Science U. S. A.* 91, 7398–7403.
- Sato, A. (2009). Both motor prediction and conceptual congruency between preview and action-effect contribute to explicit judgment of agency. *Cognition*, 110, 1, 74-83.
- Sato, A., & Yasuda, A. (2005). Illusion of sense of self-agency: discrepancy between the predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership. *Cognition*, 94, 3, 241-255.
- Sidarus, N., Chambon, V., & Haggard, P. (2013). Priming of actions increases sense of control over unexpected outcomes. *Consciousness and cognition*, 22(4), 1403-1411.
- Synofzik, M., Vosgerau, G., & Newen, A. (2008). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, 17, 1, 219-239.
- Tsakiris, E., Haggard, P. (2003). Awareness of somatic events associated with a voluntary action. *Experimental Brain Research*, 149(4): 439-446.

Vallacher, R. R., & Wegner, D. M. (1985). *A theory of action identification*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Vallacher, R. R., & Wegner, D. M. (1989). Levels of personal agency: Individual variation in action identification. *Journal of Personality and Social Psychology*, 57, 660–671.

van der Weiden, A, Aarts, H. & Ruys, K. I. (2010). Reflecting on the action or its outcome: Behavior representation level modulates outcome-priming effects on self-agency experiences. *Consciousness and Cognition*, 19, 21-32.

Wegner, D. (2002). *The illusion of conscious will*: The MIT Press.

Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation - Sources of the experience of will. *American Psychologist*, 54, 7, 480-492.

Wenke, D., Fleming, S. M., & Haggard, P. (2010). Subliminal priming of actions influences sense of control over effects of action. *Cognition*, 115(1), 26–38.

Wyart, V., & Tallon-Baudry, C. (2008). Neural dissociation between visual awareness and spatial attention. *The Journal of Neuroscience*, 28(10), 2667-2679.