15

Action-Oriented Understanding of Consciousness and the Structure of Experience

Anil K. Seth, Paul F. M. J. Verschure, Olaf Blanke, Martin V. Butz, Judith M. Ford, Chris D. Frith, Pierre Jacob, Miriam Kyselo, Marek McGann, Richard Menary, Ezequiel Morsella, and J. Kevin O'Regan

Abstract

The action-oriented approach in cognitive science emphasizes the role of action in shaping, or constituting, perception, cognition, and consciousness. This chapter summarizes a week-long discussion on how the action-oriented approach changes our understanding of consciousness and the structure of experience, combining the viewpoints of philosophers, neuroscientists, psychologists, and clinicians. This is exciting territory, since much of the resurgent activity in consciousness science has so far focused on the neural, cognitive, and behavioral correlates of perception, independent of action. Our wide-ranging discussions included questions such as how actions shape consciousness, and what determines consciousness of actions. The specific context of self-experience, from its bodily aspects to its social expression were considered. The discussions were related to specific theoretical frameworks, which emphasize the role of action in cognition, and identified an emerging empirical agenda including action-based experiments in both normal subjects and clinical populations. An intensive consideration of action is likely to have a lasting impact on how we conceive of the phenomenology and mechanisms of consciousness, as well as on the ways in which consciousness science will unfold in the years ahead.

Introduction

There is a renewed emphasis within cognitive science on the role of *action*. Set in contrast to classical paradigms which emphasize computation involving mental representations, action-oriented perspectives emphasize the enactive,

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science," Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18, series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

embodied, and embedded nature of cognitive systems. While there are many variations of the action-oriented approach (for summaries, see Thompson and Varela 2001; Wilson 2002; Engel et al. 2013; Dominey et al., this volume), they have in common the important idea that actions are not just the outputs of a cognitive system; rather, cognitive processes are shaped and may even be partly constituted by the actions they subserve. Accordingly, cognition is "for" action, not for the generation of abstract world models subserving planning and problem solving. In this chapter, we examine how this pragmatic turn in cognitive science (Engel et al. 2013) impacts our understanding of consciousness and the structure of experience. This is challenging and exciting territory, especially when set in contrast to prevailing approaches in consciousness science, which generally studies the neurocognitive correlates of perceptual scenes (e.g., Dehaene and Changeux 2011) independently of action.

There is a trivial sense in which action impacts consciousness through the selection of sensory samples. The action-oriented approach suggests much deeper influences, which are explored in this chapter. We start by offering some working definitions of consciousness and action, noting that not only can action shape consciousness, but also that we are (sometimes) conscious of actions (our own, and those of others). Thereafter we outline candidate theoretical frameworks, which turn out to be useful in organizing a discussion of action and consciousness: the Bayesian brain (Friston 2009; Clark 2013b; Seth 2014), the sensorimotor contingency (SMC) theory (O'Regan and Noë 2001; O'Regan 2011), the distributed adaptive control (DAC) theory (Verschure et al. 2003), and enactive autonomy approaches (Varela et al. 1992; Di Paolo et al. 2010). Two key questions are addressed: How does action shape consciousness, and what determines our conscious awareness of action?[1] Notions of *goal-directedness* and *hierarchical organization* (of functional architectures) turn out to be critical in these discussions.

We next examine possible neuronal substrates and functional roles relating action and consciousness, again capitalizing on the four candidate theoretical frameworks. This leads to an analysis of action and consciousness in the specific setting of the *self*, where "selfhood" can be understood to operate at multiple levels: from physiological homeostasis and interoception, to social and cultural structures and norms. We consider the notion of *joint action* (in particular, the mother-infant dyad) to be an especially illuminating example of how also social influences structure conscious experience through action.

Considerations of selfhood provide a useful context to identify specific empirical challenges for an action-oriented view on consciousness. We discuss whether such a view could shed new light on pathologies involving disordered conscious experience (e.g., schizophrenia) and whether pathologies of motor control (e.g., like locked-in syndrome, and amyotrophic lateral sclerosis) might illuminate the action-consciousness relation. These considerations highlight

[1] "Consciousness" and "awareness" are used synonymously here.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

the experimental opportunities and limitations in studying consciousness from an action-oriented perspective, and we close with a brief survey of future and emerging directions.

Consciousness and Action: What Are They?

What Is Consciousness?

Discussions of consciousness always confront issues of definition. Watertight definitions are not needed in advance of scientific progress; they advance in lockstep and eventually emerge from a mature understanding of mechanisms. Here we offer instead some basic "identifications," recognizing that these cannot do justice to the diversity of views on what consciousness "is."

A first distinction can be made between *creature consciousness* (whether X is conscious at all), which can be titrated into distinct conscious *levels* (e.g., from dreamless sleep to vivid conscious wakefulness), and *state* or *content* consciousness (what X is conscious of; e.g., the components of a conscious visual scene at a given time) (Rosenthal 2005; Seth et al. 2008). Importantly, conscious contents are, at least for humans, remarkably diverse. Very broadly they include experiences of the world (including other selves), of one's own body (from the inside as well as from the outside), of action, of emotion, and even of abstract cognitive operations (cognitive phenomenology). Having conscious contents about one's own mental states is sometimes referred to as "higher-order" or "reflexive" consciousness (Rosenthal 2005). In what follows, we focus primarily on conscious contents, as opposed to creature consciousness.

A second and more controversial distinction is between *phenomenal consciousness* and *access consciousness* (Block 2005). Roughly, the former means "what it is like" to have a particular experience, whereas the latter refers to the information in conscious experiences that is accessible or available to consumer cognitive mechanisms, which include those mechanisms that can supply explicit behavioral report. "Report" here means an action (verbal or otherwise) that conveys what one is conscious of to an experimenter. Importantly, since report itself is an action, it may shape or even constitute conscious contents.

Other definitions of consciousness are more closely tied to particular theoretical frameworks. For instance, in SMC theory, consciousness is defined as a set of abilities to interact with the world (O'Regan and Noë 2001), whereby we are (phenomenally) conscious of X when X is a quality defined by a SMC and when we are poised to use X for flexible behavior. SMC theory is an example of an approach according to which consciousness constitutively rests on action.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

What Is Action?

Action is simpler to define, though there are still interesting boundary conditions. A straightforward approach is to say that *an action is any goal-oriented manipulation of an external or internal situation* (Jeannerod 2006). Not all movements are actions, and not all actions are movements. "Covert" actions that do not involve movements include top-down attention switching, manipulation of mental states, and autonomic control including, for example, glands and smooth muscle control. Covert actions could also include planned but unexecuted bodily movements. In addition, the same movement can be an action in some situations but not in others (contrast the patellar reflex with an attempt to kick a ball), and the same movement can participate in different actions depending on the goal (waving a hand to scare a fly or to say goodbye). Thus, the key feature of actions is the association with a *goal* or *intention* (Dretske 1988).

Actions can also be classified as *instrumental*, *epistemic*, or *communicative*, according to the goal (Kirsh and Maglio 1994; Gergely and Jacob 2013; Seth 2015). The goal of an instrumental action is to effect a change of a particular kind. The goal of an epistemic action is to generate new information. The goal of a communicative action is to effect a change in somebody else. Note that communicative actions can be either epistemic (to indicate to somebody else a property of the world or of a mental state) or instrumental (to get somebody to do something).

From a neurobiological perspective we can distinguish multiple levels of the control of the skeletal-muscle system starting with the spinal cord reflex circuits, which directly control the skeletal-muscle system defining movement primitives (Mussa-Ivaldi and Bizzi 2000), and brainstem-dependent discrete behaviors such as eye blinks, grasp, and posture associated with reticular formation and the red nucleus regions, and stereotyped patterns such as those involved in feeding, defense, and reproduction regulated by the periaqueductal gray (Panksepp 2005). These latter systems interface with learning systems that can render discrete experience-dependent action such as those observed in classical conditioning, which depend on the amygdala and the cerebellum. Lastly, forebrain structures (e.g., cerebral cortex, basal ganglia) are centrally involved in goal-directed voluntary actions which can comprise complex sequences of movements. These different layers of movement, behavior, and action are all tightly coupled in the interaction between the organism and the environment.

Frameworks for Cognition and Action

Considering the importance of action within consciousness can be helped by declaring (though not necessarily endorsing) *frameworks* or *architectures* which express particular theories or provide structures by which action and

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

experience can be related. Here we consider four candidate frameworks which put specific emphasis on action (see Figure 15.1):

1. Bayesian brain (Friston 2009), when equipped with concepts of active inference (Friston, Samothrakis et al. 2012; Seth 2014; Friston, this volume)
2. SMC theory (O'Regan 2011; O'Regan and Noë 2001)
3. Distributed adaptive control (Verschure et al. 2003; Verschure, this volume)
4. Enactive autonomy and autopoiesis (Varela et al. 1992; Di Paolo et al. 2010)

The Bayesian Brain

Within the Bayesian brain approach, perception is understood as a process of inference on the (hidden) causes of sensory signals. Although its origin in the work of Helmholtz emphasizes that the mechanisms and processes of inference can be (and usually are) *un*conscious, the outcomes from this process may shape or constitute conscious contents. An emerging consensus suggests that conscious phenomenology is shaped more by (Bayesian) priors or top-down expectations than by (bottom-up) prediction errors (Melloni et al. 2011; Hohwy 2013; Chang et al. 2015; Mathews et al. 2015), a position which fits nicely with evidence for the importance of top-down signal flow for consciousness (Lamme and Roelfsema 2000; Pascual-Leone and Walsh 2001; for a review, see Lamme 2010).

Three aspects of the Bayesian brain approach deserve emphasis in the context of the current discussion. First, the Bayesian brain is hierarchical, so that posteriors at one level can form priors in the level below, instantiating a process of "empirical Bayes." This means that high-level goals or intentions can percolate throughout the hierarchy to shape priors at levels descending all the way to the sensory epithelia or spinal cord.

Second is the concept of *active inference*, which says that prediction errors can be minimized not only by updating prior predictions but also by performing actions to change sensory samples. Accordingly, the active inference view underlies both perception and action: actions are generated through the minimization of proprioceptive prediction errors through engagement of classical reflex arcs (Friston, Samothrakis et al. 2012). Importantly, active inference emphasizes the deployment of predictive models for control rather than representation, calling on parallels with theories of predictive homeostasis in cybernetics (Conant and Ashby 1970; Seth 2015).

Third, priors, predictions, and prediction errors are always associated with *precisions* (inverse variances), which determine how strongly they affect inference. Attention corresponds to optimization of precision weighting, which corresponds to modulating the gain of prediction errors at specific hierarchical

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

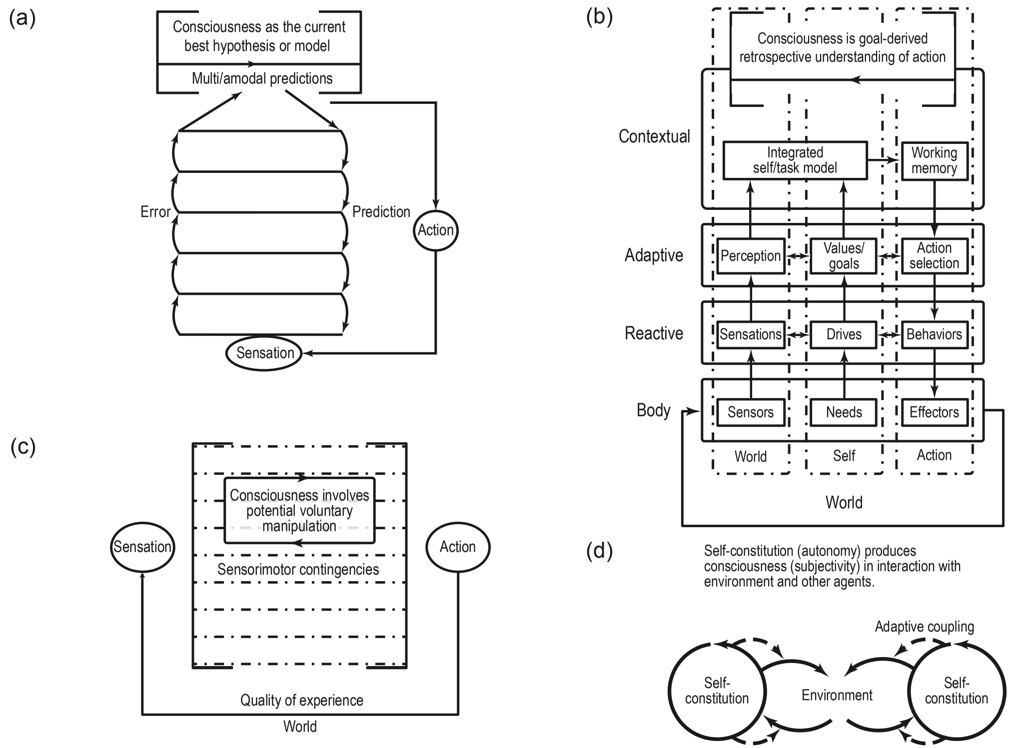

Figure 15.1 Schematics of four theoretical frameworks underlying discussions of action-consciousness relations: (a) Bayesian brain or predictive processing, (b) distributed adaptive control (DAC); (c) sensorimotor contingency (SMC) theory, and (d) enactive autonomy. Reprinted with kind permission from Ezequiel Di Paolo.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

levels (Feldman and Friston 2010). This modulation impacts the updating of the forward (generative) models underlying inference.

Although the "standard" Bayesian brain approach does not offer an explicit account of how goals are generated, broader formulations under the rubric of the free energy principle link goal states to a fundamental imperative to the maintenance of homeostatic or allostatic physiological integrity (Friston 2010; Seth 2015) through active inference. Explicit mechanistic models of goal-emergence in this framework remain, however, to be elaborated (see, however, Friston et al. 2014).

Sensorimotor Contingency Theory

SMC theory holds that perception is an activity involving skillful engagement with the world. Inheriting from the Gibsonian notion of "affordance" (Gibson 1979), conscious perceptual content is given by mastery of the SMCs, which are regularities governing how sensory stimulation depends on the activity of the perceiver. For instance, the conscious perceptual quality of redness is given by an implicit knowledge (or mastery) of the way red things behave given specific actions. Thus, in SMC theory, action plays a constitutive role in conscious phenomenology.

SMC theory also gives primacy to intentional actions, which effect manipulations of the objects of perception. Accordingly, conscious perception requires that the actions underpinning SMCs be exercised intentionally. Provocatively, there can be no conscious phenomenology without (potential) voluntary action.

Distributed Adaptive Control

DAC theory explains the mind in terms of the embodied brain interacting with its environment (Verschure et al. 2003; Verschure, this volume) and postulates that brains evolved to generate action to maintain the agent by means of a multilayered control structure (see Figure 14.2 in Verschure, this volume) involving somatic, reactive, adaptive, and contextual layers. Across these layers, three columns address states of the *world*, the *self* and *action*. DAC thus proposes a highly specific architecture that has been matched to specific brain regions and realized in real-world systems. It assigns explicit roles for goal representations in the structuring of action, as well as for Bayesian perception as inference. DAC also proposes that internal representations of SMCs lie at the heart of goal-oriented action, providing a link to SMC theory (although DAC, but not SMC theory, endorses a representationalist functional architecture). DAC also expresses central features of the enactive autonomy view (see below), by taking coupled mind, brain, body, and environment systems as its explanatory target.

The DAC architecture proposes a specific role for consciousness: the retroactive reconstruction of the sequentialized course of action(s) that led to goal

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

achievement, to enable the derivation of norms and values for the optimization of the unconscious parallel control systems that generated the course of action in real time. Put simply, consciousness serializes and interprets the present to define future goal-oriented action (Verschure, this volume).

Enactive Autonomy

This framework emphasizes the importance of autonomy and self-organization in cognition (Varela et al. 1992; Di Paolo et al. 2010) and shares with SMC theory an emphasis on skillful engagement between the agent and its environment. It is distinguished, however, by an emphasis on how this engagement supports the autonomous identity of the agent (or, more generally, system) (Thompson 2007). Accordingly, an autonomous system is one where component processes that continually self-produce the system's identity have a mutual dependence, creating an implicit organizational identity that the system maintains. This ongoing need for self-maintenance forms the basis of inherent values and goals for the system, and where such an agent can adaptively modulate the way it is coupled to the environment, these values and goals form the basis of cognitive activity (i.e., how it makes sense of, or adaptively copes with, the world around it). This emphasis echoes both the Bayesian brain approach, when seen in the wider context of the free energy principle, which grounds inference on the maintenance of organismic homeostasis (Friston 2010, 2013) and the DAC theory, which is predicated on the interaction between self-regulatory allostatic-embodied agents and their environment (Verschure et al. 2003). In contrast to the Bayesian brain and DAC approaches, proponents of enactive autonomy tend to place a very strong emphasis on phenomenology, but say less about the specifics of underlying mechanisms.

Summary

These frameworks are just four among a range of possibilities, chosen for their broad representation within the field of cognitive science and their explicit consideration of either action or consciousness (or both). The frameworks are not necessarily exclusive. SMC theory and enactive autonomy both underline the importance of skillful engagement with the world, though neither makes claims about mechanistic implementation. It is possible to provide SMC theory with a mechanistic foundation in Bayesian terms by conceiving of SMCs as (active and counterfactually informed) inferences about sensorimotor regularities (Seth 2014) or via the principles of DAC theory (Verschure et al. 2003). Free-energy principle readings of the Bayesian brain also connect with the enactive emphasis on organismal homeostasis and autopoeisis (Friston 2013). Finally, the DAC architecture is more concrete than SMC theory (and enactive autonomy) or Bayesian brain while arguably encompassing concepts from both.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

Action and the Structure of Experience

Equipped with these working definitions and candidate frameworks, we turn to the core questions regarding the relationships between consciousness and action, broadly grouped into two categories: the influence of action on consciousness and awareness of action.

Influence of Action on Consciousness

There is a trivial sense in which action determines conscious contents, by specifying the configuration of sensory receptors in relation to the environment (e.g., through eye movements). In our discussions of the action-oriented approach, we asked whether there are deeper, nontrivial influences that arise from the tight coupling of perception and action.

Prefiguring this question, Sperry (1952) noted that the outputs of a system are often more informative about its working than are its inputs. For instance, one can learn much more about the inner workings of a blender by observing its output (e.g., juices) than its input (e.g., fruits). Sperry thus concluded—consistent with Gibson—that conscious percepts are more isomorphic (i.e., similar in form) with potential action plans than with the proximal patterns of sensory inputs.

A global perspective on this issue comes from Merker, who argues that the organization of the entire phenomenal or conscious field is defined to support flexible action. For adaptive influence over the motor domain, most conscious contents should appear *as if* they transpired from a first-person perspective (Merker 2013) (perhaps with the exception of contents underlying communicative actions). In addition, accurate spatial representation is essential for adaptive action selection. It would not be adaptive for a nearby object to be represented as if positioned far away. Insofar as the action selection process in the motor domain must take into account *spatial distance from the organism* as one of its factors in the selection process, then all perceptual contents about the external world (including the body) must have a common, egocentric reference (Merker 2013). An interesting corollary of this idea is that different action possibilities (e.g., from different body morphologies) would necessarily give rise to different organizations of the global conscious scene.

A second relatively global aspect reflecting the influence of action on consciousness concerns the influence of goals. Several issues are involved: The existence of a goal distinguishes an action from a (mere) movement (a kick from a patellar reflex). In addition, goals distinguish between different types of action (instrumental, epistemic, communicative). These distinctions may operate at the level of the unfolding movements (internal or external) or at the level of how resulting changes in sensory samples are used to update conscious perceptions. For example, it seems plausible (although evidence is lacking) that a movement, if deployed as an epistemic action, may preferentially modulate conscious contents related to a (target) object of the world, but when deployed

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

as a communicative action may modulate specific conscious contents related to others' mental states. Another way to say this, again using vision as an example, is that we never simply "see"; rather, we "see socially" or "see walkingly" (Gibson 1979).

There are several examples where action and action preparation have been shown to impact directly on conscious contents. In one study, distance perception was found to be altered when participants were holding a baton for use in reaching movements (Witt et al. 2005). Vishton et al. (2007) extended this finding in the context of the Ebbinghaus illusion, showing that planning a reaching movement affected the perceived size of a visual object. Interestingly, Fleming et al. (2015) recently demonstrated action-specific disruption of perceptual metacognition. In a simple visual discrimination task, they used transcranial magnetic stimulation to disrupt motor responses underlying the response not chosen (targeting premotor cortex). Although objective discrimination performance was unaffected, confidence in correct responses was reduced. This selective reduction in metacognitive capacity implies an effect of motor representations of unexecuted actions on conscious perception.

In a different context, Hayhoe et al. (2003) used eye-tracking data to show how visual memory is affected by motor planning in natural tasks (e.g., making a sandwich). Their results suggest that object- and hand-relative spatial structures as well as object identities are determined and constructed on-the-fly. Thus, current working memory content appears to be modified by current intentional actions, and this is relevant to conscious content inasmuch as current conscious contents are closely associated with working memory (Bor and Seth 2012; see also Hagura et al. 2012, who show that motor planning of ballistic reaching movements induces subjective time dilation).

Other examples of modulation of conscious content by action include experiments which show that directional instructions or directional intentions can prime—or bias—the perception of bistable, ambiguous visual motion displays in the direction of the intended action (Wohlschlager 2000). Similarly, Butz et al. (2010) have shown that rotating tactile stimuli can bias the conscious perception of similar, bistable ambiguous visual motion displays. In this case, tactile bias depended on the orientation of the hand in space relative to a head-centered frame of reference (see also Salomon et al. 2013, who show that proprioceptive signals can bias conscious access, in a masking paradigm).

More dramatic effects of action on consciousness can be found in the domain of self- and body-related consciousness (discussed later) and in the phenomenon of sensory suppression. It is well known that the performance of actions causes the *attenuation of sensory experiences*, which explains why self-tickling is largely ineffective (Blakemore, Wolpert et al. 2000). Recent interesting evidence suggests that actions may suppress *auditory* sensory experience across a broad range of actions, mediated by direct ipsilateral projections from secondary motor cortex to auditory cortex (Schneider et al. 2014).

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

Behavioral report (i.e., actions which convey to an experimenter information about current conscious contents) provides an intriguing context in which to consider the influence of action on consciousness. Methodologically, it is (on most accounts) imperative for investigating consciousness, as explicit reports (verbal or otherwise) are the primary means for obtaining information about the conscious contents of another person. Yet report, by definition, involves a communicative action. The question therefore arises as to whether report actions change or shape the reported experience. As yet, evidence is thin. Recent work has shown that frontal brain activations, often associated with (access) consciousness, are absent when report is not required (Frassle et al. 2014), but this does not demonstrate directly any change in the experience as a consequence of report. This question also pertains to early work by Marcel (1993) which demonstrated that different perceptual judgments were made when different reporting channels were employed (e.g., button presses vs. eye blinks).

Awareness of Actions

The relationship between consciousness and action is perhaps most directly expressed in *when* and *how* we experience actions themselves. To underline this point, awareness of actions should be understood as awareness of the actions themselves (at different levels of abstraction), not as awareness of the consequences of actions. A general way to think of action awareness is that a focus on action execution and on the required motor (or mental) control (i.e., a focus on the control of the unfolding SMCs) might constitute the continuously experienced phenomenology of action.

An examination of action awareness first requires, however, a phenomenology of action. Compared to perceptual phenomenology, the terrain is only dimly lit (Pacherie 2008). The phenomenology of action is “thin and evasive” (Metzinger 2006). It is useful to characterize two distinct dimensions along which action phenomenology can be organized. The first reflects a scale from high-level goals and intentions to low-level individuated muscle contractions prescribing movements (or their equivalents for autonomic control or mental acts). The second distinguishes experiences of intention from experiences of agency or authorship.

An extensive line of work, originating with Benjamin Libet, has studied experiences of intention and agency (for a helpful review, see Haggard 2008). Put simply, the experience of intention has to do with awareness of a goal and is sometimes associated with the phenomenology of an “urge,” whereas the experience of agency depends on action-outcome association, often reflecting goal fulfillment. A now classic finding in this area is that of *intentional binding*: if an event is experienced as the effect of an intended action, the time between the action and the event is perceived as being shorter in duration (Haggard et al. 2002). This not only provides an additional example of an effect of action

From “The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,”
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

on consciousness (in this case, of perceived time), it also shows that the conscious experience of agency depends on both the prior expectations about outcome of the action and the observed outcome of the action. This implies that a role for action in consciousness is to create the experience of agency; that is, I am in control of my actions, in the sense that it was I who chose the option (action selection) and caused the outcome, and that I could have chosen another option, giving rise to the possibility of regret (Frith 2014). It is interesting to note that voluntary actions appear to have a specific kinematic signature, as compared to stimulus-driven actions (Becchio et al. 2014).

Goals remain important in action awareness. Our awareness of our actions tends to follow the scale of the goal driving the action in question. Habitual and well-practiced actions are performed relatively automatically, with little deliberate awareness. Generally, as our expertise increases so does the abstractness of associated goals, and this changes our awareness of the situation (Speelman and Kirsner 2005). Expert musicians are frequently as focused on the emotional tone of their performance as on the playing of the music, much more so than the movements of their fingers required to produce notes. Similarly, your consciousness of your actions in a conversation has much more to do with the message you wish to convey, or the effect you wish your words to have, than on the physical movements involved in producing the sounds (and sometimes even on the choice of words themselves). The execution of practiced action is frequently robust to minor perturbations (as in the case of dynamic stability of the articulation system for jaw movement interruptions; Kelso et al. 1984). However, where an action sequence meets failure, when goals and reality sharply diverge, we become keenly aware of the more fine-grained resolution of the movements involved.

Architectures, Mechanisms, and Implementations

At this point it is useful to return to the four candidate frameworks (Bayesian brain, SMC, DAC, enactive autonomy) to try to connect action-consciousness relations to specific mechanistic properties. The four frameworks specify these relations to different degrees, as will be seen.

Under the Bayesian brain account, conscious content can be associated with the Bayesian model or "hypothesis"—spanning multiple hierarchical levels—that best suppresses prediction error (Hohwy 2013). When these hypotheses have to do with proprioceptive (and possibly vestibular[2]) predictions, they reflect awareness of actions. As discussed above, a key feature of the Bayesian brain approach is the deployment of precision weighting (attention) to optimize inference. Accordingly, the specific phenomenology of action may

[2] The "broken escalator" phenomenon is a wonderful example of vestibular mis-prediction. It is the sensation of losing balance reported by most people when they step onto an escalator that is not working. Strikingly, the illusion is not diminished by knowledge that the escalator is broken (Reynolds and Bronstein 2003).

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

depend on the hierarchical assignment of precision to proprioceptive predictions. Action awareness may shift from low-level movements to high-level goals during the acquisition of expertise. The converse shift may occur when an action is blocked, frustrated, or otherwise goes wrong: the resulting cascade of prediction errors will lead to reassignment of precision weighting targeting lower hierarchical levels. (If I try to pick up a glass and almost knock it over, I suddenly become aware of the specific movements of my fingers.) Sensory attenuation is also naturally accommodated within the Bayesian brain framework: the precision of proprioceptive prediction errors needs to be underweighted so that predictions are fulfilled rather than updated. Under the SMC theory account, actions (including "mental" actions) are constitutive of consciousness. SMC theory makes no claims about specific brain mechanisms and thus is compatible with both Bayesian brain and DAC accounts. Enactive autonomy approaches, however, tend to be incompatible with representational accounts (Hutto 2012).

With respect to goals, priors at one level of a hierarchy of control operate as goals for the level below. The character of our awareness of an action may in this view be governed by the level of either highest prediction error (indicating where more attention and control is most needed) or highest prediction confidence (where action is at the most abstract level of description consistent with the amount of practice or habituation). DAC theory makes a similar prediction, pointing out that consciousness involves a maintenance of coherence between an agent's predictions (including forward models of its own actions) and the interaction with the world (Verschure, this volume). Within the DAC perspective, consciousness serves to derive value from goal-oriented action and its outcomes through retrospective reconstruction, thus marking a distinction with generic Bayesian brain approaches in which conscious experiences are temporally aligned with the dynamics of perception and action. Both SMC and autonomy-based approaches assume that consciousness always involves action or potential action, so that our awareness of our own actions is complementary to our awareness of properties of the world. Insofar as we are frequently engaged with multiple goals simultaneously—or at least hierarchies of goals at different granularities—conscious experience of our actions and the world will be variable and textured by the various forms of goals and the associated skills being deployed (O'Regan and Noë 2001; Di Paolo et al. 2010; McGann 2010; O'Regan 2011).

Multiple Levels of Self-Experience

Experiences of being a self, having a body, and perceiving the world from a first-person perspective are intimately tied to action, perhaps more so than experiences of the external world. Some aspects of self-experience have already been discussed (namely, experiences of intention and agency). Other

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

more basic experiences of selfhood have to do with *being* and *having* a body as well as the first-person perspective (Merker 2013). At the other extreme, self-experience—at least in humans—is co-constituted by social interactions and the elaboration of a self-narrative which evolves over time. Importantly, all aspects of the self involve actions, but the actions are different, ranging from real bodily actions to mental and ultimately social actions.

The Bodily Self

Considering the body brings into focus processes of *interoception* (sense of internal physiological condition) and *autonomic control* (Craig 2003), as well as multisensory integration (Blanke and Metzinger 2009; Tsakiris 2010). Importantly, autonomic control signals can be thought of as internalized actions: they are associated with goals and have discrete somatic effects, following a hierarchical structuring akin to that associated with external action (see above). At the foundation of this stand notions of homeostatic control and allostasis (i.e., maintenance of multiple homeostatic loops) linked to concepts of drive reduction (Sánchez-Fibla et al. 2010). A long tradition extending at least as far back in time as James and Lange has associated conscious emotional or feeling states with the perception of changes of physiological condition (James 1894). Recently, this tradition has been extended via the Bayesian brain approach to suggest that emotional states arise through a process of *active interoceptive inference* (Seth 2013, 2015).[3] Accordingly, autonomic control can be understood as implementing active interoceptive inference to suppress interoceptive prediction errors and as a consequence, maintain physiological homeostasis.[4] Thus, internal actions may be fundamental to emotional experience. In this light, it is fascinating to note recent evidence for direct projections from motor cortex to internal organs (Levinthal and Strick 2012).

Interoceptive signals also influence the experience of body ownership and first-person perspective, another key constituent of conscious selfhood. This is shown by the fact that individual interoceptive sensitivity predicts susceptibility to the "rubber hand illusion" (Tsakiris et al. 2011), and that cardiac feedback, presented visually on augmented-reality representations of body parts and avatars, can modulate experienced ownership and first-person perspective depending on the accuracy of this feedback (Aspell et al. 2013; Suzuki et al. 2013).

Action seems especially critical for shaping the experience of body ownership and first-person perspective. This might be because actions induce specific correlations between proprioceptive, visual, and perhaps other modalities,

3 SMC theory claims that only those internal actions that are voluntarily controlled and executed with an intent to sense an individual's emotional state can be experienced consciously.

4 Alternative views associate emotion with a prediction error between allostatic targets and states of the world (Verschure 2012b), rates-of-change of prediction error or free energy (Joffily and Coricelli 2013), or embodied predictions shaping action selection (Barrett 2012).

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

which can be used to infer a distinction between one's own body and the rest of the world. These actions, therefore, are epistemic actions and might be related to motor babbling during infancy (Bullock et al. 1993). Notably, proprioceptive motor tactile correlations, along with interoceptive signals, are available before birth and thus might play an extremely important role in the development of the experience of body ownership before vision comes into play (Gori et al. 2008; Rochat 2010).

The Stable Self

A Bayesian brain perspective may shed light on the phenomenology of self-stability: though the experienced contents of the world continually change, our experience of "being a self" is highly continuous (as pointed out by William James, long ago). The reasoning is as follows. As we move through the world (and continuously move our eyes), we are bombarded with an ever-changing sequence of sensory signals and experiences which, in their raw form, are extremely difficult to predict (by analogy with pre-Copernican attempts to predict the motion of the planets). The solution is to develop the prior that there is a self, moving through a relative stationary and stable three-dimensional world (Butz 2008). Given this prior, we can then predict the changing pattern of stimulation received by our senses. The same process applies in the social sphere. Predicting the various interactions we have with others is greatly simplified given the prior that there is population of individuals, each with their own (and this includes the self) relatively stable preferences and styles (Robalino and Robson 2012).

The Social Self and Joint Action

An important dimension of the experience of selfhood and the structure of the sense of self involves intersubjectivity and our social and cultural relations with the environment. Several, not necessarily compatible, perspectives in current embodied cognition help shed light on this issue. Here we broadly distinguish three: (a) the neural implementation of the social self, (b) the particularly rich set of evolved social abilities in humans, and (c) the intersubjectivity as constitutive for minimal sense of self and the dynamic construction of embodied identity.

Neural Implementation

One way to think about the (experience of the) social self is in terms of its neurological implementation and the question of how the social self can be functionally represented and processed in the brain. "Mirror" neurons and activations in biological motion areas when an individual performs and observes actions (Rizzolatti and Sinigaglia 2010) suggest that actions, intentions, and

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

goals of others are partially represented by means of self-representations. The resulting inference-making capabilities about others imply that we understand others somewhat in the way we understand and control ourselves. Importantly, once such representations and simulations of others are active and overlap with our self representation and self-control mechanisms, the representations of others need to be separated from those pertaining to self (Gallese et al. 2004; Butz 2008). Moreover, by simulating actions of others through self-grounded representations, perspective taking becomes possible in a social context (including self-reflection) by monitoring one's self from the perspective of another person (Frith and Frith 2006; Hassabis et al. 2014).

Social Abilities and Joint Action

Distinguishing self-representations from representations of others is particularly important in the performance of cooperative and communicative actions, awareness of communicative goals, and awareness of others' psychological states. Indeed, much comparative evidence suggests that compared to nonhuman primates, humans are uniquely cooperative (Tomasello 2014) in at least three fundamental ways: through *trust*, a basic ability for *perspective taking*, and a basic receptivity to *cultural learning*, each of which contributes to shaping the human social self. First, cooperative breeding, which is uniquely human among great apes, requires an unusually high level of trust in one's in-group members (Hrdy 2009). Second, humans engage in joint actions in an unprecedented way compared to nonhuman primates. In joint action, two (or more) agents must represent a common goal and fulfill two (or more) complementary roles, and switch roles according to context. Arguably, role taking and role switching require some perspective taking (see above). Third, developmental evidence suggests that human infants are uniquely tuned to the ostensive signals that reflect assumed intentions to convey relevant information. These ostensive signals include speech in "motherese," contingent responses, and direct gaze, which is facilitated by the fact that, compared to the eye of nonhuman primates, the human eye involves a uniquely dark iris on a white background (Kobayashi and Kohshima 1997). In addition to this three-tiered basic social self, the social self is further enriched by membership to linguistic, religious, nationalistic, and other groups of various sorts.

Autonomous Dynamics in the Social Domain

An important question emphasized by the enactive autonomy perspective is whether social interactions are constitutively necessary for individual agency and self-identity. Social contexts can be incorporated into Bayesian brain, SMC, and DAC through specification of hierarchically high-level predictions that transcend the individual. The enactive autonomy account (at least in line with most readings) goes further in this regard: the social is not merely the

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

context in which cognitive systems are embedded, but figures in the constitution of agency and self-identity. In other words, an agent's mechanism of self-organization is not fully determined in terms of *individual* sensorimotor activity but is continuously *co*-enacted with others. Arguably, this entails a sense of openness and connectivity associated with individual agency and selfhood (Kyselo 2014).

Empirical Challenges

The issues discussed so far in this chapter point to a number of important empirical challenges, which we outline below.

Disorders of Motor Control

If action is important in shaping, or even constitutive in, conscious contents, then disorders of motor control (where the capacity for movement is severely diminished or even lost altogether, as in amyotrophic lateral sclerosis or locked-in syndrome) should have measurable effects on consciousness. In addition to investigating basic levels of consciousness in these patients, it would be interesting on a more fine-grained level to use functional magnetic resonance imaging to examine a range of action-related brain responses. These include whether the mirror system still responds to the perception of goal-directed actions (e.g., reaching and grasping) executed by others, whether the superior temporal sulcus still responds to the perception of biological motion, and whether one finds responses in motor systems, as in healthy subjects, during perception of action sentences or action verbs (e.g., "kick," "lick," "pick"). Finally, one might examine whether their responses are in accordance with the somatotopic representation of the respective effectors with which the action described by the verb is normally executed, as has been found by Pulvermüller and colleagues in healthy subjects (see Pulvermüller 2005; Pulvermüller, Hauk et al. 2005).

Another interesting context arises in neuroprosthetics, for instance in cases where a robotic arm is controlled directly by brain signals obtained from motor cortex (see Dominey et al., this volume). Here, in some cases it is possible to modulate explicitly the degree to which the robotic arm is controlled by the brain signals, or by "helping" signals provided by an external computer (given an identifiable goal, like reaching toward a target) (e.g., Miele et al. 2011). This raises the intriguing possibility of calculating a psychometric curve to relate the experience of agency to the degree of control the subject's brain has over the robotic limb.[5]

5 In a recent study with stroke patients, in which the movements of the paretic limb were amplified beyond its physical characteristics, the point of subjective equality of the upper extremities could be shifted without loss of sense of agency (Ballester et al. 2015).

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

Regret

The experience of regret provides intriguing possibilities for the empirical study of the importance of action in consciousness (Frith and Metzinger, this volume). Like disappointment, regret is a decision-related emotion, closely tied to action (or inaction). Disappointment is elicited if the outcome of a chosen action is not as good as expected on the basis of previous experience. In contrast, regret occurs when we realize that an action not chosen (including doing nothing) would have achieved a better outcome. Thus, regret involves counterfactual cognition. Regret is an unpleasant experience and thus behavior is altered by counterfactual thinking about possible future regret. For example, humans can take steps to avoid finding out about the outcomes associated with alternative actions (Reb and Connolly 2009) or we can adjust our choices to minimize anticipated regret (Filiz-Ozbay and Ozbay 2007). It has also been argued that rats experience regret (Steiner and Redish 2014).

Several open questions concerning regret require further study: How does the phenomenology of regret relate to other aspects of awareness of action such as agency, responsibility, and selfhood? How does the phenomenology of regret relate to that of other emotions, and are there specific interoceptive and autonomic signals associated with regret? How does the conscious monitoring of regret have its impact on planning and future action? What is the relation of regret to social emotions such as shame and guilt?

Empirical Possibilities in the Further Study of Action and Consciousness

One reason why action-consciousness relations have not been studied extensively, as compared to perceptual consciousness, is that it can be difficult to design experimental protocols involving rich movement. Most psychophysical and especially neuroimaging environments impose severe restrictions on movement, such that action repertoires are often restricted to simple eye movements and button presses.

Emerging technologies like virtual and augmented reality may ameliorate some of these limitations, by providing "virtual" analogs of rich action repertoires in physically restrictive situations. Even these technologies, though, face important limitations, since virtual actions are necessarily not grounded in the subject's own physiology. Nonetheless, a great many interesting experimental opportunities can be identified. Many of these have to do with experiences of body ownership, where virtual reality can induce systematic and highly controllable manipulations. Prospectively, these techniques could be used to examine social impacts on body experience and misperception of self-consciousness as somebody else's consciousness; for instance, the clinical condition of feeling "a presence" (i.e., the feeling of being in the presence of another person, although there is nobody there) may be of interest for study (see Blanke et al. 2014).

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

Experimental opportunities can also be identified with clinical populations, other than those with extreme deficits in motor control (the locked-in syndrome and amyotrophic lateral sclerosis patients mentioned earlier). These are predicated on the notion that inaccurate signaling related to actions may cause a variety of symptoms. In schizophrenia, inaccurate signaling may underpin delusions of control (the belief that one's actions are being controlled by another agent) (C. D. Frith 2012). The normal attenuation of sensations resulting from our own movement depends on relating the intention to move (or speak) with the anticipated sensory changes (the forward model) (Frith et al. 2000a; Ford et al. 2007). This explains why we cannot tickle ourselves. Schizophrenic patients with delusions of control *can* tickle themselves, presumably because something has gone wrong with this action-perception loop (Blakemore, Smith et al. 2000). Abnormalities in neural activity that precede simple, self-paced, noncontingent button presses have been related to a lack of motivation and general apathy (Ford et al. 2008). Perhaps the mismatch between predicted sensations and the resulting experiences ultimately diminishes the motivation for action. In terms of the Bayesian brain framework, failure to suppress proprioceptive feedback would create problems for active inference and thus for action. This might contribute to general motor awkwardness and neurological soft signs in schizophrenia (Bachmann et al. 2005) as well as the motor symptoms that often precede the onset of the illness (Walker and Lewine 1990; Cannon et al. 1999). The implications of this system for the social self and its dysfunction in schizophrenia are also evident. Patients with schizophrenia have reduced activity in the posterior superior temporal sulcus during imitation of actions and during action observation, possibly reflecting a breakdown in internalized mimicking, one route to understanding the minds of others and increasing social facility (Thakkar et al. 2014).

Open Issues

The topic of action and consciousness is unusually broad and any single review is necessarily incomplete. Several important areas have not been discussed in any detail, including the *role of language* (speech acts are canonical examples of actions which, like any other, can be instrumental, epistemic, or communicative). Another neglected area has been the role of *synthetic modeling*. Synthetic models are extremely valuable in connecting theoretical frameworks to specific empirical predictions. Here, concrete architectures like DAC might play an important role (Verschure, this volume). Another aspect of experience emphasized by an enactive autonomy approach is its *temporally extended* nature. Here, relations to different timescales in action need further attention. Finally, the scope of this treatment has not extended very far into an examination of the brain regions and circuits involved in action awareness and consciousness. This, too, demands close attention.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

There is also a need for more work on action phenomenology, following Pacherie (2008). It is worth noting that there is also phenomenology to *reflex* movements (which are not actions) or to the movement of my hand when you move my hand, but I don't, and which is therefore *your* action, but *not mine*.

Concluding Remarks

Bringing an action-oriented perspective to the study of consciousness clarifies several important issues. A first clear outcome is that *action actively shapes and structures conscious experiences* in ways that extend beyond the trivial case of selecting sensory samples. Action emphasizes the *openness of consciousness* to extrapersonal influences. More controversial is the suggestion, emerging in particular from SMC theory and enactive autonomy approaches, that actions (possibly social actions) are constitutive of (some) conscious experiences. Although this seems highly plausible for nonhallucinatory experiences of action, it is less clear that all conscious experiences constitutively rest on action. Interestingly, actions may shape conscious experiences in specific ways depending on whether they are instrumental, epistemic, or communicative. Data on this question would be very valuable.

A second outcome is that actions both shape conscious scenes and engage a specific phenomenology of action awareness, which depends on goals and emerges at multiple hierarchical levels and spatiotemporal grains. Action awareness, in the specific guises of experiences of intention and agency, is also a key determinant of experiences of selfhood.

A third outcome is that it makes sense to speak about action and awareness at multiple scales, spatial and temporal. Actions can take place within the body (autonomic control), at the body-world interface, and within larger social spheres. Timescales can range from fine finger movements and rapid autonomic contractions to long-term intentions to, for example, become an academic and write book chapters on action and consciousness. Actions may shape experiences at each of these scales and may give rise to specific action-related phenomenology at each of these scales. Social actions, including joint actions as an illuminating example, have emerged as especially important domains in which action determines self-experience, carrying important implications for psychological well-being.

A fourth outcome is a general (though not universal) appreciation that a Bayesian framework seems valuable when considering the effects of action and embodiment on awareness. Importantly, this framework does not exclude (and may benefit from) insights gained from alternative frameworks like SMC theory and enactive autonomy as well as specific operational architectures like DAC. Indeed, suggesting specific mechanisms brings to light interesting new questions, such as whether action-related experiences are constructed prospectively (as suggested by Bayesian brain) or retrospectively (as suggested

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

by DAC). These points emphasize the importance of systems-level embodied modeling.

A final cautionary remark is that the degree to which action affects or constitutes consciousness remains unclear and thus it is not possible to conclude, at this stage, whether the "pragmatic turn" will constitute a revolution in our understanding of consciousness. Action is important, but it may not be everything, when it comes to experience.

Acknowledgments

The authors gratefully acknowledge the following. AKS: the Dr. Mortimer and Theresa Sackler Foundation, which supports the work of the Sackler Centre for Consciousness Science. MVB: Emmy Noether grant BU1335/3-1 from the German Research Foundation (DFG). JF: Department of Veterans Affairs (I01CX000497) and National Institute of Mental Health (MH-58262). CDF: Wellcome Trust and Aarhus University. MK: Marie-Curie Network TESIS: FP7-PEOPLE-2010-ITN, 264828. MM: Convergent Science Network of Biomimetics and Neurotechnology. JKO'R: ERC Advanced Grant "FEEL" No. 323674. All authors are grateful for illuminating input from other participants at this Ernst Strüngmann Forum.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.