

Conscious and Unconscious Control of Spatial Action

B Hommel, Leiden University, Leiden, the Netherlands

© 2009 Elsevier Inc. All rights reserved.

Glossary

Action effect – Sensory consequences of actions (reafference), assumed to be integrated with the actions they accompany and to serve as their ‘mental’ retrieval cues.

Dual-pathway models of sensorimotor processing – Models assuming that sensory codes are translated into motor activity along two processing pathways, one subserving direct online sensorimotor transformation and another allowing for perceptual elaboration and off-line action planning.

Executive ignorance – The phenomenon that voluntary agents have conscious access to their action goals only but no insight into how these are translated into action.

Feedback – Feedback control considers action-produced, reafferent stimulation to fine-tune or stop the action, that is, preliminary results of an ongoing process can influence earlier stages of the process in a continuous loop.

Feedforward – Feedforward control consists of the complete prespecification of an action, which then runs off in a ballistic, context-insensitive fashion.

Motor program – Originally a concept to refer to lists of muscle instructions, later relaxed to allow for more abstract representations of only the invariant properties of actions.

Somatic marker – A representation of the affective consequences of an action (action effects), which can be used as a shortcut to the action in selection processes.

we actively navigate through our environment or for manual actions, such as grasping a cup of coffee or playing piano. But even actions that serve entirely nonspatial purposes, such as speaking to someone or singing, are controlled at lower motoric levels by targeting particular spatial end configurations of one’s jaws, lips, and tongue. Not to forget the actions we carry out to generate perceptual information, such as scanning a visual scene with one’s eyes, turning one’s head and ears to sound sources, and systematically exploring the texture of surfaces with one’s hand. Hence, we move our body in space all day long, and thereby change the configuration of our body parts and their relationship with the environment. At the same time, however, there is very little we know about all these changes and the ways we achieve them. Consider, for instance, you would be asked how you tie your shoes or how you ride a bike (assuming that you master these skills). Presumably, there is very little of interest you could say, simply because you know very little about the details. Most of your description will not differ much from what any other observer could see just as well (that you pick the right shoe lace with the index finger and thumb of your right hand, and so forth), which implies that you have no privileged knowledge about how you move your body in space. How is this possible, that you can move almost any limb of yours in almost any physically possible way, and yet have no idea how? This phenomenon of executive ignorance, as it is sometimes called, has been studied for more than 150 years by now and, fortunately, some progress has been made.

Action Control

Human action is characterized by its voluntary nature, that is, with very few exceptions we do not respond to stimuli in an invariant, reflexive manner but carry out movements to reach a particular goal. One may ask whether this is possibly

Introduction

Any overt action carried out with the muscles moves the body in space. This is particularly obvious if

an illusion and actions are only interpreted as being goal-directed after the fact. However, the final goal actions aim at often shapes the movement elements needed to reach it, so that actions reflect their goal right from the start. For instance, grasping an object entails a number of systematic movement features, with one being the aperture of the hand adjusting to the size of the object long before the contact is made. Another example is the way people pick up an object they want to move – such as when putting a book back into a shelf – which is commonly chosen to guarantee a comfortable position of the hand at the end of the action (the end state comfort principle). These and other observations have led to the idea that voluntary actions are mediated and driven by motor programs, which are conceived of as stored representations of actions that can be retrieved and used for controlling an action at will, even in the absence of feedback. Numerous findings have supported this idea of a central engram of action. For instance, patients with a complete loss of kinesthetic feedback are still able to carry out goal-directed actions with the afflicted limb and experimentally deafferented monkeys can still grasp, walk, and jump. Other interesting observations stem from analyses of action errors. Errors sometimes anticipate later action elements and thus are actually correct action components at the wrong time – famous examples are the expressions delivered from William Archibald Spooner, such as ‘the queer old dean’ (where he actually wanted to refer to ‘the dear old queen’). Less entertaining examples are everyday lapses in the order of sequential manual actions, such as when making tea or coffee, or typing errors, such as correctly doubling the wrong ‘letter.’ Further evidence is provided by experimental studies of actions varying in complexity, such as the number of action steps or their accuracy demands, that show that the time needed to execute an action increases with complexity. Hence, there is strong evidence that actions are commonly guided by internal representations, including the actions’ goal.

Actions and Effects

Combining the phenomenon of ‘executive ignorance’ (i.e., the observation that people do not have

much insight into the ‘how’ of their voluntary actions) with the assumption that voluntary actions are guided by something like motor programs poses an important question: If we have no conscious access to the contents of motor programs, how are we able to select them? One possible, philosophically very interesting answer will be considered in the next section: consciousness may not have anything to do with the selection of actions but, rather, may merely serve to make sense of or justify these actions after the fact. But there is another possibility. As ideomotor theorists in the nineteenth century such as Hermann Lotze and William James have suggested, conscious experience may well be associated with (and perhaps even causally effective in) the selection of the action goal, which then in one way or another takes care of the further motor programming itself. The developing infant and the novice facing a novel motor task may start by carrying out all sorts of involuntary movements, a process sometimes called ‘motor babbling.’ The motor patterns generating these movements may be set up entirely by chance or follow genetically specified reflexes, but they should systematically produce particular sensory feedback. An automatic integration mechanism may associate the motor patterns and the perceived action effects in a bidirectional fashion. If so, the motor pattern could later be intentionally retrieved by endogenously activating an aimed-at action effect – thinking of the goal triggers the movements necessary to reach it without any conscious insight into their inner workings.

Numerous studies have in fact revealed that people pick up novel effects of their movements spontaneously and continuously, thereby steadily increasing their action repertoire and the number of goals they can realize. Spontaneous acquisition has been observed in adults, children, and infants, and experiments have demonstrated that agents consider the novel action effects in the selection of voluntary actions. Brain imaging studies have shown that episodic action-effect associations are stored in the hippocampus and that reactivating them primes motor representations in the supplementary motor area, which again is known to play a crucial role in voluntary action planning.

Many studies have focused on the perhaps most obvious sensory effects, such as visual and auditory

consequences. However, recent theorizing has emphasized that the sensory codes of affective consequences may be of particular importance for action planning. According to Antonio Damasio and colleagues, we integrate motor patterns not only with codes of the reafferent feedback they produce but also with codes of the ways they make us feel. Carrying out an action can have more or less positive or negative consequences, which lead to corresponding affective states. Increasing evidence suggests that representations of these states (so-called somatic markers) are associated with the actions they accompany, so that reactivating an action tends to reevoke the associated somatic marker. If so, somatic markers can be used to select among possible actions and guide one to the action alternative that ‘feels best’ – a selection mode that one may call intuition or deciding ‘by gut.’ Interestingly, realistic studies of decision-making in complex everyday situations have revealed that good deciders are not at all following the rules of logic (which do allow good decisions on simple problems), but decide in ways they are often unable to explain themselves. Moreover, simulation studies have shown that logic-based decision-making is likely to be way too slow to work under realistic circumstances, in which decisions often have to be fast. In these cases, somatic markers may provide a kind of shortcut to appropriate decisions.

Will and Consciousness

According to ideomotor reasoning, consciously representing a goal (i.e., the intended action effect) is sufficient to retrieve and execute the necessary motor patterns. Conscious representations play thus a restricted but still crucial role in action control. Some researchers have questioned even this role however. If conscious representations would really be causal in bringing about intentional action, so they argue, the causal experience should temporally precede the action it causes. That this may not be the case is suggested by a classic experiment of Benjamin Libet. He asked subjects to carry out voluntary movements and measured the time at which they showed a readiness potential (an EEG component preceding voluntary actions) and when they began consciously intending the movement. Surprisingly, the readiness potential

was measured much earlier than the conscious intention. This suggests that the ‘brain’ had made the decision to act long before this decision was consciously experienced, which again undermines the idea that it is the experience that drives the action. Some authors have argued that the conscious experience of intentions may be the product rather than the cause of the neural processes leading to action, which means that consciousness does not play any causal role in action control. Other researchers have argued that this holds true only for the initiation of actions while a conscious ‘veto’ may still be conceivable. In other words, conscious experience may be functional not in producing but in preventing actions.

On the one hand, the available evidence clearly shows that the relationship between conscious intentions and the actions they refer to is more complicated than the common sense model of action (perception → conscious decision-making → action, see the section ‘[Sensorimotor processing](#)’) suggests. Clearly, voluntary actions can be prepared and carried out without tight conscious monitoring, which is also evident from many everyday observations: We walk without thinking of every single step and drive home with very little cognitive involvement. And yet, there is no evidence that voluntary actions are possible without a conscious representation of the goal. We cannot exclude that, in Libet’s study, consciously constructing the overarching task goal and preplanning the possible actions once in the beginning of the experiment was sufficient to drive remaining activities more or less automatically, just as consciously intending to drive home is enough to have the remaining action run in ‘autopilot’ mode.

Motor Programs and Action Plans

The original concept of a motor program that controls the execution of a movement was derived from the domain of computer programming. The idea was that sequences of muscle movements could be stored and rerun in a purely feedforward fashion whenever needed. Historically, this approach was a counterreaction to the strong emphasis on stimulus-driven behavior associated with American behaviorism, which dominated the psychological scene in the first half of the twentieth century. Behaviorist

approaches tried to reconstruct actions as responses, that is, as logical and empirical consequences of stimuli impinging on the sense organs. These approaches were surprisingly successful and did a good job in disenchanting the concept of voluntary action quite a bit. But in the 1950s and 1960s it became increasingly clear that there are too many indications that actions can be generated in a purely endogenous fashion and in the absence of any sensory and evaluative feedback to further rely on behaviorist concepts of action control. In sight of the upcoming computer revolution and the increasing temptation to apply computer metaphors to human cognition, it seemed only logical to consider action being controlled by programs that have a structure similar to computer software. If such software could produce 'overt' behavior on a computer monitor, why should mental software not generate muscle activity?

However, even if people would be able to maintain muscle-specific information associated with a particular movement, the idea is unlikely to solve a number of problems. One main problem relates to 'storage.' Consider all the possible pointing actions you could perform, all the possible locations in space that your index finger may occupy, and all the possible configurations of the limbs involved this would imply. If a motor program would really be a literal record of muscle movements, each single combination of all the factors would require a program on its own. Worse, which muscles need to move in which way to reach a particular movement goal is dependent on context and starting conditions, such as the current positions of the body and the limbs involved. A separate program would need to be constructed for each possible context. Considering that pointing is just one simple movement out of the many simple and complex movements you can perform, this would imply enormous amounts of memory to store all those programs, not speaking of the search time needed to relocate a required program in that memory. Even though the human brain entertains smart routines to handle large amounts of data, the implied memory demands seem too excessive. A second problem relates to 'novelty.' Having learned to point to one location is enough to generate any other pointing movement – even if different muscles are involved, and other movements generalize just as

well. It is difficult to see, however, how generalization from muscle-specific programs would work. Hence, very detailed representations of movements, as implied by the original motor program concept, would be counterproductive.

The major theoretical move to fix these and other problems was to assume that the motor program proper consists of invariant information only (a kind of motor schema or plan), whereas parameter slots were defined for configuring a program for a particular task and purpose. The idea was that action control consists of two phases, one in which the proper program is selected and another in which the (commonly spatial and/or timing-related) parameters of this program are specified to tailor it to the situation at hand. Pointing would be a good example: Selecting a generalized pointing program would specify which limbs are involved and how they move in relation to each other, whereas a location parameter would specify the spatial target of the movement and a timing parameter would specify how fast the action is carried out. This approach successfully overcomes both the storage problem, as very few programs are actually stored, and the novelty problem, as using different parameters would be a natural way to generalize.

Sensorimotor Processing

The main motivation for the original motor program approach were findings that voluntary actions can be carried out even in the absence of any sensory feedback. Accordingly, the intention was to conceive of a control structure that can run in a completely feedforward fashion without any need for sensory information. However, even though it is possible to carry out at least some actions without any feedback, it is clear that the accuracy of actions is often enhanced if feedback is available. Moreover, sensory contributions to action programming do not need to be restricted to feedback, that is, to information that is produced by the action, but the sensory information available at the onset or during the execution of an action may be just as important. Indeed, once programs are thought to accept parameters, it makes sense to consider that parameter values are delivered by the environment. This

raises the question of how sensory information affects and shapes action control.

Single-Pathway Models

One way or another, all organisms are capable of responding to stimuli from their environment, be it flowers stretching toward the sun or bacteria fleeing areas of high phenol concentration. In lower organisms the links between sensors and effectors are rather short and direct, so that the action of a given effector is easy to understand from tracking the transmission of signals picked up by the sensors. In the beginning of the systematic investigation of human sensorimotor processing, the same logic was applied: Motor activity was thus conceived of a more or less direct extension of sensory processing, with the task being to trace stimulus signals through the processing system until it makes contact with the muscles. A famous example of this line of reasoning is the empirical and theoretical work of René Descartes, who suggested that the pineal gland transforms the electric signals it receives from the eye into hydraulic energy driving the muscles. Later researchers like Franciscus Cornelis Donders further developed this approach and suggested that human cognition can be understood as a (rather extended) sequence of processing stages beginning with the sensory registration of a stimulus and resulting in the movements of the muscles. The interesting implication of Donders' approach was that it allows for the separate measurement of the time demands of each single processing stage by systematically manipulating the processes necessary for performing a given task.

Not all researchers working with this single-pathway conception of human information processing were interested in the role of consciousness, but those who were located conscious experience right in the middle between sensory coding and muscle movement. That is, it was not assumed that the whole chain of processes would be accessible for conscious experience, but rather, conscious experience and decision-making was thought to be associated with the higher ends of perception (such as the conscious experience of meaningful objects and events) and the decision which action this perceptual experience calls for (hence, the

conscious experience of will). In other words, the role of consciousness, or of the processes associated with conscious experience, was to separate perception from action, so that actions would be no longer directly driven by stimuli but could be planned ahead and triggered in the absence of any external stimulation. The theoretical idea shared by single-pathway models is sketched in **Figure 1**: In lower organisms sensory coding more or less directly leads to motor output, whereas higher organisms (and humans in particular) are assumed to have acquired means to decouple stimulus processing from motor programming – with consciousness being either functional in, or at least associated with, this decoupling.

Multiple-Pathway Models

Even though it seems clear that humans are no longer purely reflexive organisms, which certainly implies some additional processing capabilities, one can ask whether this necessarily implies the loss of more direct linkages between sensory and motor processing. Indeed, increasing evidence suggests that the two pathways sketched in **Figure 1** do not represent 'alternative' processing routes but, rather, a tandem of concurrent processing streams that distribute labor in a particularly efficient way. Numerous observations have suggested that (1) there must be more than one pathway from sensors to effectors and that (2) not all pathways are monitored by consciousness and accessible for conscious experience.

Early evidence for the presence of multiple pathways was provided by prism studies, in which human subjects wore goggles that systematically




Figure 1 Dual-pathway models of sensorimotor processing distinguish between a low-level pathway translating sensory codes into motor activity and a high-level pathway subserving higher-order cognitive operations and action planning.

distorted their visual input, for example, by turning it upside down. From these studies, two observations are particularly impressive. First, the conscious experience of the visual world (or, perhaps better, the conscious visual experience of the world) turned to normal after some time, which means that the subjects have learned to ‘correct’ the visual input for the distortion. Second, however, the subjects regained their ability to carry out spatial actions (such as walking or riding a bike) much earlier, long before conscious experience was successfully readapted. This implies that actions have access to sensory information that is not available for, and certainly not controlled by phenomenal experience. Numerous dissociations between conscious perception and judgment on the one hand and (commonly) manual motor action on the other have been demonstrated since then – for example, people can correctly point at stimuli they consciously mislocate in space and correctly grasp objects they consciously misperceive in size. Likewise, even if people are unable to detect a displacement of an object they are in the process of grasping (by moving the object during an eye blink) their hand moves straight to the new location without any noticeable delay or hesitation.

Patient studies provided converging evidence. For instance, patients with lesions in their visual cortex are unable to consciously perceive stimuli falling into the visual field corresponding to the lesioned area; and yet they can correctly point to these stimuli (the so-called ‘blindsight’ phenomenon). Patients with bilateral damages in the posterior parietal cortex (the Balint–Holmes syndrome) are perfectly able to judge the orientation of a slit in a wooden board, but are unable to orient their hand so that it can pass through the slit. Reversely, patients suffering from visual form agnosia (due to lesions affecting visual cortex) are able to orient their hand accordingly but at the same time cannot judge the orientation of the slit. Hence, numerous findings converge on the conclusion that multiple pathways exist from sensory coding to motor processing and that not all pathways are consciously accessible.

Even though researchers agree that multiple pathways exist, they differ with respect to the question of how such pathways should be characterized. Early approaches assumed that one pathway is

mainly responsible for the processing of identity-related information (the ‘what’ pathway), such as shape and color, while another pathway focuses on spatial information (the ‘where’ pathway). A later approach has suggested that one pathway mainly serves to inform conscious perception (the ‘ventral’ pathway, named after its assumed anatomical location) while another feeds motor processes with action-related information (the ‘dorsal’ pathway). Other approaches have focused more on the distinction between off-line processing (be it for perception or action planning) and online processing (be it for sensory coding or motor programming). Even though the debate is not entirely settled, most differences are rather semantic and restricted to details. A general agreement is that there must be two or more different processing pathways that are likely to represent a combination of what previously was thought to be alternatives: A direct link between sensory and motor processes that we share with lower organisms and a much more cognitively penetrated and in part consciously accessible link, including higher-level perception and the planning of goal-directed actions (see [Figure 1](#)).

Establishing multiple concurrent pathways for what looks to be the same purpose (bridging the gap between sensors and effectors) seems redundant and creating needless confusion and communication problems. However, close consideration reveals that the purposes the multiple streams apparently serve are rather different and to some degree mutually incompatible. First, there are different demands in terms of processing ‘speed.’ If we consider direct sensorimotor transformation processes (the lower branch in [Figure 1](#)) to continuously feed motor actions with online information about the current environmental state of affairs, these processes need to be fast and specific enough to guide a hand toward a visible goal. In contrast, speed is not so much of an issue when recognizing an object and planning an action via the higher-level route.

Second, the ‘type of processes’ operating in the two pathways is likely to differ. Sensorimotor transformation consists in the application of acquired transformation rules (which need readjustment when wearing goggles with reversal prisms, for instance), whereas recognition and action planning

require the extensive use of memory retrieval, prediction and probability estimates, and reasoning – all processes unnecessary for sensorimotor functioning.

Third, a related point, the ‘amount of information’ to be considered differs. Sensorimotor transformation works best if the incoming sensory information is as pure as possible, whereas recognition and action planning benefit from extensive use of anticipation, elaboration, and retrieved memory codes – the informationally richer the resulting representation, the better.

Fourth, sensorimotor and the more cognitive pathways need rather different ‘spatial reference frames’ to code information efficiently. Motor control needs to consider the body of the agent and the configuration of the effectors involved. Action goals need to be translated into body- and effector-related reference frames, so that the required movements can be properly programmed and executed. Many of such so-called pragmatic body maps have been identified in the primate brain and it is interesting to consider that the way most of these maps represent action space is not reflected at all in our conscious experience. In contrast, conscious experience commonly conceives of space as egocentric (considering the body as a whole) and allocentric (i.e., independent of the body). Even though the mechanism underlying this ability is not yet well understood, we are somehow able to construct the latter from the former by integrating many different egocentric ‘snapshots’ of experiences into a coherent allocentric representation, such as of the city we live in by navigating through it by foot, bike, subway, and car. The emerging allocentric representation is particularly important for planning purposes, but it is at the same time way too coarse to control the movements of our feet. Allocentric representations are also important for the recognition of objects irrespective of their orientation, another ability that requires the integration of egocentric snapshots into an ego-independent representational format.

Interacting Levels of Action Control

So far, we have seen three converging lines of theoretical developments. First, modern approaches to

action control started with a strong emphasis on purely feedforward programming but became increasingly lenient with respect to the consideration of the available sensory information to adapt motor programs to current needs. Recent models assume that stored action-control structures are more like Swiss cheese: The general structure, a mere schema or action plan, is maintained after use but a number of slots are intentionally left open to be specified online whenever needed. Second, ideomotor approaches explain the executive ignorance of conscious experience with respect to the details of action selection and execution by referring to the automatic integration of motor patterns with representations of their sensory consequences (i.e., action effects). Accordingly, action-effect codes become retrieval cues for the associated motor patterns, so that the conscious representation of intended action effects is sufficient to launch the programs necessary to reach them. Third, modern approaches to sensorimotor processing assume that there are multiple streams from sensory processing to motor activity. Some of these streams link sensory codes more or less directly to motor execution, thus providing a fast and pure picture of the sensory state of affairs, whereas other streams are busy with the cognitive elaboration of the input and the setting up of action strategies.

Figure 2 shows how these assumptions fit together. The observation that conscious representations are restricted to the perceivable action consequences is consistent with the claim that conscious access is restricted to higher cognitive processing pathways, pathways that are specialized for the off-line elaboration of perceived events and anticipatory action plans. In some sense, the only difference between perceiving a current event and planning an action is the temporal reference: What we call perception is awareness of the sensory aspects of an event that happens to be right now, based on the activation of sensory codes describing this event, whereas what we call action planning is awareness of the sensory aspects of an event, again based on the activation of sensory codes describing it, that is still to occur – a sort of imagery. One might object that action planning has motor consequences while perception has not. However, apart from the fact that requiring information for perception almost




Figure 2 Action control in a dual-pathway system. Specifying a goal leads to the selection of codes representing the sensory consequences of actions suited to bring about the intended effects (action-effect codes or AE). Action-effect codes activate associated action plans (the off-line component of action control), which contain open parameter slots to be filled online by available sensory information (the online component of action control).

always involves motor activity, there is increasing evidence that perceiving an object primes the actions it affords and perceiving actions of other people activates compatible action tendencies. Moreover, if people judge and compare visual objects, their performance is affected by relations between the actions these objects imply, suggesting that action-related information is automatically activated in what looks like straightforward perceptual tasks. Hence, perception and action planning seem to bias motor processes in comparable ways, so that it makes sense that our consciousness treats them equally. Selecting a plan does not fully determine its motoric realization. Rather, it only makes sure that the intended sensory consequences are produced in one way or another (the equifinality principle). In which way they are is eventually determined by the present sensory situation, which fills in the open parameter slots. The whole filling-in operation and, presumably, the filled-in parameter values are not only commonly unconscious but presumably not accessible by conscious operations in principle.

Even though some authors have argued that the two main processing pathways are independent, this stands in contrast with a number of empirical findings, but it is also unlikely for logical reasons. If it is the case that low-level sensorimotor transformation procedures are entirely unaffected by higher-level perception, memory retrieval, decision-making, and action planning, how do they know which parameters to select and feed into which motor program? Consider a grasping movement. You are often facing numerous objects, so that without selecting one target object transformation procedures would not know which sensory information to use as input. Minimally, the perception-and-planning pathway would need to communicate to the sensorimotor transformation pathway which object or location in space is to be attended. Next, the sensorimotor pathway would need to know for which action the transformation is required. Grasping mainly calls for shape and location information while pointing is likely to rely more on contrast and/or color. This means that attention also needs to be directed to particular

stimulus dimensions, a job that relies on information about the planned action. Furthermore, analyses of sequential actions have revealed context effects and optimization strategies, which seem to imply rather massive interactions between pathways. For instance, the way one grasps a given object depends on what one is going to do with this object next – to move it somewhere else, to bring it to the mouth, or to throw it away. This means that rather concrete action parameters are biased by the next upcoming action elements, which means that parameter specification at lower levels must have some access to, or be affected by higher-level planning processes. Thus, even though the sensorimotor transformation pathway and the perception-and-planning pathway differ in many respects, and are better off doing so, good communication between them is essential for coherent goal-directed action.

Any distortion of this communication can lead to abnormal phenomena. As mentioned earlier, patients with optic ataxia (the Balint–Holmes syndrome) have no problem to consciously perceive objects, but at the same time they have difficulty to manipulate them manually, whereas patients with visual agnosia show the opposite pattern. In line with dual-pathway models, this seems to reflect selective impairments of sensorimotor transformation and perception-and-planning pathways, respectively. Damage of the frontal lobes, which are particularly important for action planning, has been reported to lead to so-called ‘utilization behavior’: Patients tend to carry out actions on objects that are ‘sensomotorically correct’ without making sense or following any goal. This seems to suggest that the sensorimotor transformation pathway is still functional but no longer under control and in the service of action planning. A possible explanation for delusion of control experiences reported by schizophrenics might be that action planning is still working but no longer represented in consciousness. Accordingly, the outcomes of actions are not consciously anticipated and therefore surprising. Daniel Wegner and others have suspected successful anticipation of action outcomes to be a major component of the attribution of agency, with anticipated outcomes indicating self-performed actions. If so, a loss of conscious anticipation is likely to prevent the attribution of self-agency.

Role and Function of Conscious Control

Now that we have roughly circumscribed the processes that are related to consciousness, it seems appropriate to ask what role consciousness may play in and for them. Why is it that we have conscious access to higher-order perceptual processes and action planning, but are widely ignorant with respect to the specifics of sensory and motor processing? Depending on one’s philosophical stance regarding the concept of consciousness, different types of answers to this question are appropriate.

From an evolutionary standpoint, it is certainly difficult to say why conscious experience exists anyway and, if it exists, why it does not reflect all existing processes. However, it is also clear that the phylogenetic development of processing pathways on top of and in addition to the basic, relatively direct linkage between sensory organs and motor systems is associated with the emergence of conscious experience. In fact, the more a species shows evidence of similar multiple-pathway structures, the more signs they show for aspects of consciousness. This does not account for the *qualia* of conscious experience, but it points to an interesting correlation between the existence of consciousness and a particular alternative way to transform sensory information into motor activity, which allows for the decoupling of perception and action, for prospective planning, and for stimulus-independent, endogenously generated action.

From a functionalist point of view, it seems interesting to consider why this correlation exists. That is, what are the characteristics of processes that accompany conscious experience (irrespective of the issue whether this correlation expresses a particular causality or not)? Most interestingly, there is no evidence of any conscious access to ‘processes’; rather, it is the ‘states’ these processes produce that are consciously represented: We perceive the sensory consequences of an action but not the operations necessary to make them accessible to our experience. Moreover, increasing evidence suggests that conscious experience is restricted to rather global, highly integrative states. This is particularly obvious for the generation of action goals. Even though we are commonly aware of the action goal we currently pursue (at least to

some degree), we commonly have no insight into why it is this goal we chose. This is not to say that we cannot think of justifications for the goals we have after the fact or point to the role of a subgoal in a broader action plan (like boiling water in the process of making tea), but why exactly we ended up having this but not other goals we commonly cannot say for sure. (Again, this is not any different from perception; e.g., while it is evident that we can switch between different views of multistable stimuli, we are unable to say why and how we switch.) Nevertheless, chosen goals often reflect multiple constraints, often more than we can enumerate, which suggests that the generation of a goal considers huge amounts of information collected from the senses, from memory, and from *ad hoc* reasoning. Again, our conscious experience does not have access to all that information and the processes integrating it, but is merely presented with the final outcome.

A given goal can lead to different types of action. Western legal systems are based on the idea of conscious deliberation and, indeed, at least the selection of novel and complicated actions to reach a particular goal is commonly accompanied by conscious experience. Again, the experience is not referring to processes but to states representing the effects the considered actions are expected to have. Whether or not the experience as such has any causal relevance – which given that experience relates to products rather than causes seems questionable, however – it seems clear that intermediate states of deliberation are consciously accessible and often associated with conscious experience. However, it is only the anticipated (perceptual and affective) consequences of the considered actions the experience refers to. Eventually, one alternative is being selected (often for justifiable but actually unknown reasons) and thereby made available for conscious representation. This is where conscious insight into action control ends; the next aspect we experience is only whether the action produced unexpected results (while expected results often go unnoticed).

The strong link between consciousness and action planning and the observation that conscious experience often refers to the consequences of actions have motivated a number of so-called ‘simulation theories’ of cognition. In a certain sense, these theories reverse

the commonsense approach to action control. The commonsense approach suggests that it is consciousness that makes action planning possible by allowing one to mentally simulate and ‘think through’ alternative actions before deciding which action one prefers. Simulation theories acknowledge the value of such mind plays but assume that it is the ability to plan and play through an action that actually came first and conscious experience only emerged as a consequence. Making creative use of action-effect associations that originally evolved for much more basic functional reasons may have allowed extending one’s thoughts beyond the current, perceptually available situation by reasoning in terms of what-if.

To summarize, conscious experience is most closely related to the goal of an action and its intended consequences, and the experience seems to focus on a rather general, highly integrated representation of these aspects. This observation fits with approaches that see consciousness as a kind of global work space, such as the theory of Bernard Baars. The idea is that much cognitive processing is going on in local modules that are highly specialized and fast, but their inner workings and intermediate results are not consciously accessible. Their products however, the information they generate, is sent to a global work space, where information can thus be related to each other, integrated, and acted upon. Even though this approach fails to explain the *qualia* aspect of conscious experience, it is consistent with the picture that emerges from the available evidence, the picture of a surprisingly ignorant and only globally informed voluntary agent.

See also: Intentionality and Consciousness; Perception, Action, and Consciousness; Perception: Subliminal and Implicit.

Suggested Readings

- Baars BJ (1988) *A Cognitive Theory of Consciousness*. New York: Cambridge University Press.
- Blakemore SJ, Wolpert DM, and Frith CD (2002) Abnormalities in the awareness of action. *Trends in Cognitive Sciences* 6: 237–242.
- Clark A (2007) What reaching teaches: Consciousness, control, and the inner zombie. *British Journal of the Philosophy of Science* 58: 563–594.

- Damasio A (1999) *The Feeling of What Happens: Body, Emotion and the Making of Consciousness*. London: Vintage Books.
- Hesslow G (2002) Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences* 6: 242–247.
- Hommel B (2007) Consciousness and control: Not identical twins. *Journal of Consciousness Studies* 14: 155–176.
- Hommel B, Müsseler J, Aschersleben G, and Prinz W (2001) The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences* 24: 849–878.
- James W (1890) *The Principles of Psychology*, vol. 2. Cambridge, MA: Harvard University Press.
- Jeannerod M (2006) *Motor Cognition: What Actions Tell to the Self*. Oxford: Oxford University Press.
- Miller GA, Galanter E, and Pribram KH (1960) *Plans and the Structure of Behavior*. New York: Holt, Rinehart, & Winston.
- Milner AD and Goodale MA (1995) *The Visual Brain in Action*. Oxford: University Press.
- Prinz W (1992) Why don't we perceive our brain states. *European Journal of Cognitive Psychology* 4: 1–20.
- Wegner D (2002) *The Illusion of Conscious Will*. Cambridge: MIT Press.

Biographical Sketch



Prof. Dr. Bernhard Hommel holds the chair of “General Psychology” at Leiden University since 1999, after having worked as senior researcher at the Max-Planck Institute for Psychological Research (PhD at the University of Bielefeld in 1990; Habilitation at the Ludwig-Maximilians University of Munich). His research focuses on cognitive, computational, developmental, neural, and neurochemical mechanisms of human attention and action control, and the role of consciousness therein. He served as Associate Editor of *Acta Psychologica*, the *Journal of Experimental Psychology: Human Perception and Performance*, and the *Quarterly Journal of Psychology*, and currently is editor-in-chief of *Psychological Research*. He has (co-) edited two books on action control and the relationship between perception and action, and co-edited several special issues on attention and action control. He has (co-) authored more than 100 articles in international journals and numerous chapters in readers and psychological textbooks.