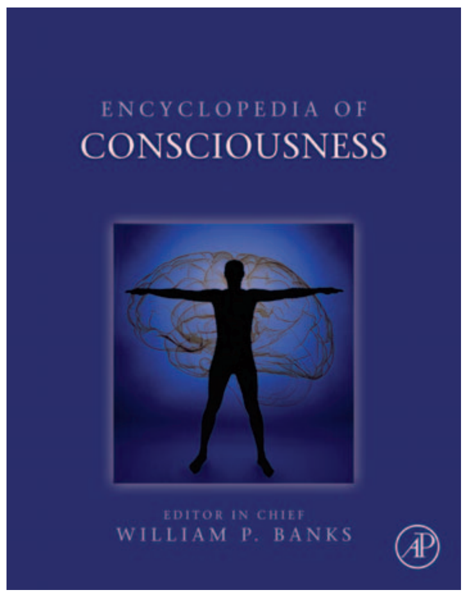

Pockett S (2009), Brain Basis of Voluntary Control. In: William P. Banks, (Editor), Encyclopedia of Consciousness. volume 1, pp. 123-133. Oxford: Elsevier.

Brain Basis of Voluntary Control

S Pockett, University of Auckland, Auckland, New Zealand

Introduction

Philosophical arguments notwithstanding, an ordinary dictionary is usually a pretty good indicator of how words are commonly used. *The Concise Oxford Dictionary* defines a voluntary movement as one that is controlled by the will.

Such a definition suggests that an article on the brain basis of voluntary control should discuss not only the brain basis of movement *per se*, but also (and perhaps more importantly) the brain basis of the will. This makes the whole exercise both more interesting and less straightforward. What is the will, anyway? Does it even have a basis in brain function? Our everyday intuition tends to be that the will may be some sort of dualist entity, capable of acting on the brain, but not itself generated by the brain. More radically, does the will really exist at all? Perhaps will is nothing more than an inference, a mythical will o' the wisp created by conscious selves desperate to believe that they and not their brains control their bodies.

Returning to earth, we might profitably adopt the same strategy with regard to the noun will as we did for the adjective voluntary – ask the dictionary. *The Concise Oxford Dictionary* says the noun will means "faculty by which one decides what to do; fixed desire or intention." This definition implies (although to be fair, it does not explicitly require) that the will is a faculty of consciousness. Certainly, thanks largely to Sigmund Freud it does seem possible nowadays to have desires that are unconscious. Perhaps we might even allow the existence of unconscious intentions, though this at least initially seems more like a contradiction in terms. But the faculty by which one decides what to do must surely be a faculty of consciousness. We may not be conscious of all the factors influencing a given decision, but we do consciously make decisions. Don't we?

Maybe not. The first professor of psychology at Harvard, William James, spent a whole chapter of his 1890 *magnum opus The Principles of Psychology* arguing against his contemporary Thomas Huxley's theory that humans do not consciously control their actions, but are basically automata. In contrast, the current professor of psychology at Harvard, Daniel Wegner, has summarized recent empirical evidence on the matter in a book entitled *The Illusion of Conscious Will.* Contrary to our everyday intuitions, what relatively little science has been done in this area tends to show that not only the neural events controlling bodily movements, but arguably even the neural events initiating bodily movements, are inaccessible to consciousness. Experiments show that we are not very good even at knowing whether we caused any given occurrence ourselves, or whether somebody or something else did. But then, of course, given the tenacious nature of everyday intuitions, it is not surprising that philosophers (specifically, all the philosophers represented in a recent multiauthor book called *Does Consciousness Cause Behavior?*) vehemently disagree with Wegner's conclusions on this matter (although not with the empirical evidence on which those conclusions are based). The century-old argument continues. A pessimist might conclude that a dozen decades of hard work have brought us no closer to any definitive understanding of the innocent-looking word 'voluntary.'

But, such pessimism is rarely justified. We have learned some things over the last 120 years. To ease the reader gently into these deep and treacherous waters, the present article first provides a short and relatively uncontroversial account of current knowledge about the functioning of those brain areas known to be involved in decisions, intentions, and movement. It then strikes out strongly toward the center of the issue, describing a number of lines of experimental evidence that suggest that, at least much of the time, the functioning of these brain areas does not give rise to any conscious sensations or experiences at all. Finally, after skirting delicately around the whirlpool of philosophical

Encyclopedia of Consciousness (2009), vol. 1, pp. 123-133

controversy whipped up by these experimental results, we retreat to the far shore to consider briefly the possible consequences of our findings for the legal system, which presently regards humans as conscious agents.

The Neuroscience of Voluntary Control

Nearly 150 years ago, the same Huxley whose automaton theory was so vigorously opposed by James famously remarked that "the great end of life is not knowledge but action." The enduring truth of this aphorism is suggested by the fact that a large part of the brain is concerned in one way or another with the production of actions. Figure 1 shows the physical locations in the brain of the areas known to be involved in the initiation and control of voluntary movements. A list of these areas includes the primary motor cortex, the supplementary and presupplementary motor areas (pre-SMAs), the premotor cortex, the frontal eye fields, the cingulate cortex, the posterior parietal cortex, the dorsolateral prefrontal cortex (DLPFC), the basal ganglia, the thalamus, the cerebellum – and of course, much of the spinal cord (not shown).

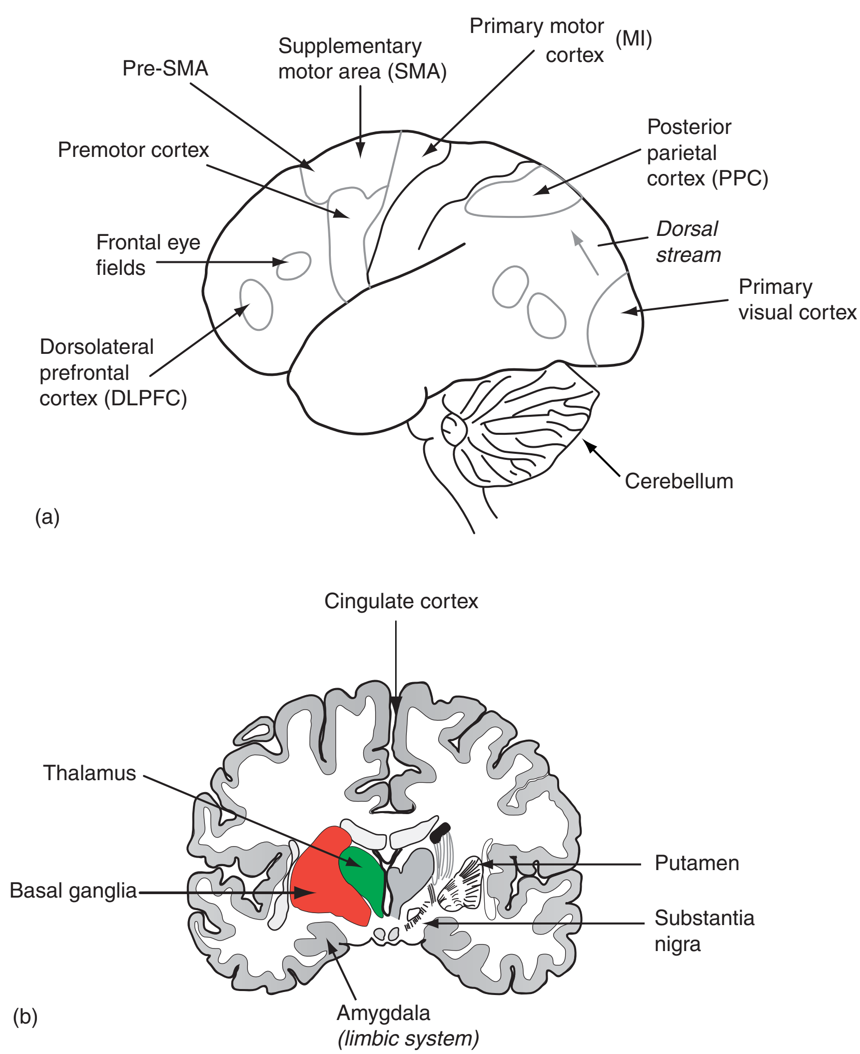

Figure 1 Location of brain areas involved in movement. (a) Surface of left hemisphere. (b) Transverse section of brain in the region of SMA. The putamen and substantia nigra are two of the many components of the basal ganglia. The amygdala is part of the limbic system, which mediates emotion and motivation.

Figure 2 shows a schematic model that attempts to pull together what is currently known about how these brain areas interact to implement voluntary control. At this stage our knowledge about the neuroscience of motor control is far from complete, and so like most (if not all) scientific models, this one must be considered a heuristic rather than a final construction.

According to the model in Figure 2, a willed movement is (not unreasonably) conceptualized as starting with the will. For the moment we can sidestep the whole controversy about the nature and even existence of the will and define that entity operationally as a specific intention.

Neuroscientific evidence suggests that there are two different kinds of intention, subserved by two widely separated areas of brain. Intentions of the first kind are called willed intentions. Willed intentions are abstract, early plans for movement. They specify the goal and type of movement, but not the detail of how the movement will be carried out. Given the general accuracy of the lay notion that the frontal lobes of the brain are involved in thinking, it is not surprising that willed intentions are generated in frontal areas, specifically the DLPFC and its adjacent pre-SMA. One important feature of willed intentions is that they need not necessarily be acted upon – indeed the road to Hell is said to be paved with good ones.

The second kind of intention is called a sensorimotor intention (or motor representation, or stimulus intention, depending on the author). This kind of intention specifies the detail of how an intended movement is to be carried out. Sensorimotor intentions are located not in the frontal areas, but toward the back of the brain in the posterior parietal cortex. Such a location is advantageous because visual input is very important in the construction of specific plans for interacting with the world, and the posterior parietal cortex is directly connected to the dorsal stream of the visual system. (Visual input from the outside

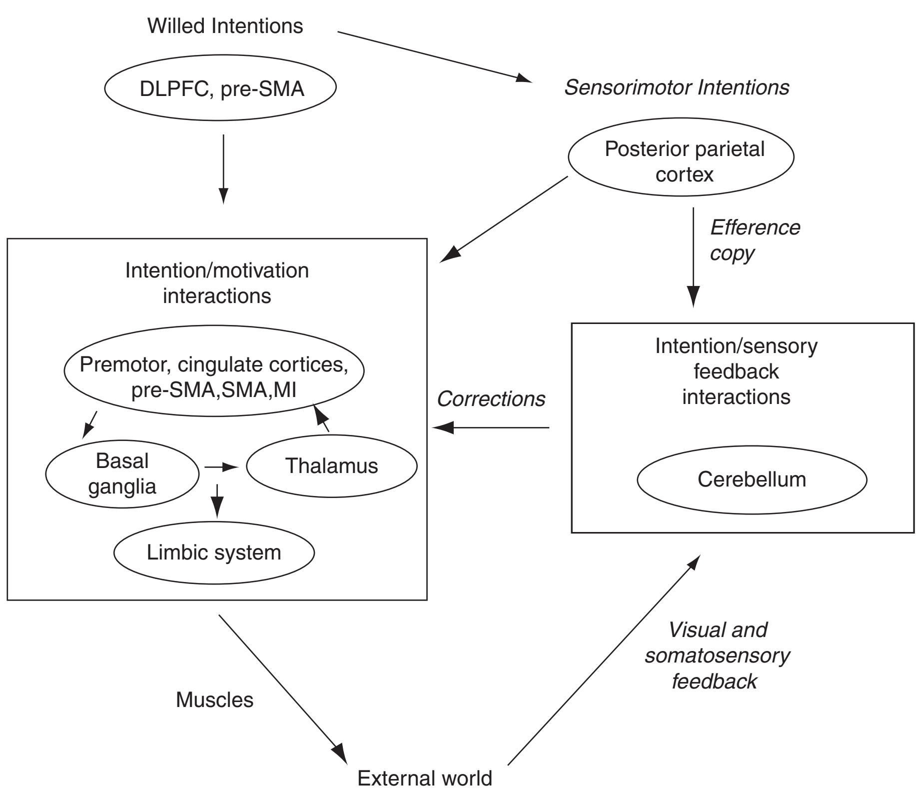

Figure 2 Model of anatomical and functional relationships in the motor system. DLPFC, dorsolateral prefrontal cortex; pre-SMA, presupplementary motor area; SMA, supplementary motor area; MI, primary motor cortex. See Figure 1 for physical locations of these areas.

world enters at the eyes and immediately flows caudally through the large fiber tract of the optic radiation to the primary visual cortex, at the very back of the brain. After traversing through the secondary visual areas slightly rostral to (in front of) the primary visual cortex, visual information then flows forward toward the nose in two streams. The ventral stream is believed by many to culminate in visual perception in the infratemporal area just above the ears. The dorsal stream is not associated with consciousness, leading instead to the sensorimotor intention areas in the posterior parietal cortex.)

Unlike willed intentions, it is not clear whether sensorimotor intentions can be formed and then not acted upon. If sensorimotor intentions can be formed but not acted upon, the initiation of an action must logically be taken as occurring after the formation of its sensorimotor intention. If sensorimotor intentions are necessarily carried out once they are formed, the initiation of action could be considered to occur before the formation of sensorimotor intentions.

This basic uncertainty about the timing of the all-important action-initiation step is reflected by the lack of a box in Figure 2 labeled action initiation. At this stage, so little work has been done on the neuroscience of action initiation that almost nothing is known about its neural basis. It is likely that the basal ganglia are involved, as shown both by (1) functional magnetic resonance imaging (fMRI) studies and (2) the fact that Parkinson's patients often find it difficult to initiate actions. (Probably the primary lesion in Parkinsonism is in the substantia nigra of the basal ganglia) (see Figure 1). However, the SMA has also been suggested as a primary site of action initiation. The truth is that the complex loops sketched in the 'Intention/Motivation Interaction' box in Figure 2 are presently little understood and at this stage no particular brain area or areas can definitely be credited with the function of initiating voluntary actions.

However, whether or not we understand when and where it happens, at some point any given action must be initiated, or launched. The action then needs to be seen through to its completion. This means that the movement must be controlled and if necessary corrected. With very fast, so-called ballistic movements, there is no time for correction during the movement itself, but if the end result of the movement is not as planned, at least the next such movement can be adjusted. Slower movements can be adjusted on the basis of visual and other feedback during the course of the movement.

Much of the function of control and adjustment of voluntary movements is subserved by the cerebellum (see Figure 1). The cerebellum is an evolutionarily ancient structure with a unique anatomical structure. More functions for it are being discovered all the time, but probably its main task is to compare a copy of the original motor instructions from the posterior parietal cortex (a so-called efferent copy) with feedback from the environment, in order to compute error and error correction messages. These error correction messages are then fed back to the posterior parietal cortex and/or forward to the frontal motor areas, the spinal cord, and eventually the muscles.

Consciousness and Voluntary Control

The brain activity described in the 'Introduction' section serves to explain the generation of first general and then specific intentions to act, and the control of movements once they have been initiated. But so far, no mention has been made of where consciousness does or does not fit into this essentially deterministic chain of events. Our everyday intuitions would probably say that the formation of a specific intention to act is conscious. The initiation of a voluntary act or movement would also generally be considered to be conscious. The dictionary definition of 'voluntary' might suggest that the control of voluntary movements is conscious as well. But intuitions and definitions are always trumped by empirical evidence. What does the empirical evidence say?

Control of Movements

It is a truism that many voluntary movements, particularly those which could be called overlearned, are carried out without the engagement of very much at all in the way of consciousness.

A skilled typist can produce a manuscript without any awareness of the multiple complex motor decisions made during the act of typing. A skilled driver may navigate the entire trip from work to home without ever thinking consciously about the act of driving, and without remembering any of the particulars of the journey on arrival. These are not acts that could be described as automatisms, such as sleep-walking, or reactions, such as withdrawal of the hand from a painful stimulus. They are complex sequences of deliberate movements, directed by both internal and external stimuli. When performing such acts, we can become very conscious of what we are doing if something interrupts the flow of events. But even then, if a sudden reaction becomes necessary – a child runs out in front of the car, for example – we act first and become conscious of the stimulus only after the action is completed. We find our foot on the brake pedal and then become aware of the child. Such anecdotal experiences of the unconscious control of ongoing movements have been amply confirmed by experimental data, as described in the article by Marc Jeannerod elsewhere in this volume.

What is the brain basis of this kind of relatively unconscious motor activity? The key to answering this question is to examine the differences between neuroimaging data obtained during the performance of novel actions and neuroimaging data obtained during performance of the same actions after they have been well learned. When learning a new action, we are conscious of every little decision. As the skill is learned, the action becomes less and less conscious. Karl Friston and Richard Frackowiak at University College London and Richard Passingham at Oxford have shown that the brain areas that become less active as a motor skill is learned are the prefrontal cortex and the SMA. This suggests that the prefrontal cortex and the SMA are somehow involved in the production of consciousness. It will be recalled from the discussion in 'The neuroscience of voluntary control' section that the prefrontal cortex is implicated in decision making and the SMA has been suggested as one of the sites involved in action initiation.

In contrast, the brain sites concerned with the unconscious control of voluntary actions are very likely to be the posterior parietal cortex and the cerebellum. As mentioned earlier, these are the sites involved in producing sensorimotor intentions (the specific plans for how a movement will be carried out) and fine-grained motor control, which the article by Jeannerod shows to be largely unconscious. This seems eminently reasonable from a biological point of view. If it were not so, we would not be able to walk and carry on a conversation at the same time. Consciousness has a very limited capacity. It would not do to be conscious of every little calculation involved in keeping one's balance or placing one's feet. We need to be able to think about other concerns while still moving efficiently about the world. Babies and people relearning the arts of movement after a neurological injury have to think about the details of walking, but the rest of us do not.

But surely we do need to have conscious control over the initiation of sequences of motor actions. We have to decide consciously to get into the car and start driving. If all our acts were automatic at that level, we could not count ourselves as conscious selves – could we?

Initiation of Movements

As mentioned earlier, considerable confusion presently surrounds the initiation of voluntary movements. One of the few experiments done so far that has examined the act of initiating a previously willed intention was first carried out a quarter of a century ago by Benjamin Libet. Libet (who died in 2007 at the age of 91) asked his subjects to watch a spot of light rotating on a clock face while they made a series of spontaneously generated finger movements. After each movement, the subject had to report exactly where the spot was on the clock face at the instant they had felt the urge or 'wanting' to move their finger. Libet called this reported instant time *W*. The clock method of measuring the objective time of a subjective event is now generally called the Libet clock, but it is actually a modification of a general method invented by Wilhelm Wundt almost a century previously. Libet's conceptual breakthrough was to compare the mean *W* times from groups of 40 movements with the electroencephalography (EEG) event-related potential extracted by back-averaging the EEG immediately preceding those 40 movements. (It was necessary to average 40 movements because event-related potentials are so small that they are buried in the biological noise and cannot be seen for individual trials. Averaging is a standard technique for pulling small signals out of noise. It works because at any given instant after (or in this case before) the event in question (in this case the movement) random noise is equally likely to have a positive or negative voltage. This means that in the mean of many trials, the noise averages out to close to zero. In contrast, the signal is always the same at any given instant before (or after) the movement, and so averaging does not change its amplitude.) The event-related potential in question is a slow, negative-going waveform that had been discovered 20 years before Libet's experiment by Kornhuber and Deecke, who named it, in German, the Bereitschaftspotential. The Bereitschaftspotential is now generally known by its English name of readiness potential (RP). An example of an RP is shown in Figure 3.

The outcome of Libet's experiment was reported in 1983 in the journal *Brain*, and has been the subject of vigorous debate ever since. The main finding was that the RP started about 350 ms before time *W* (see Figure 3). The startling implication of this is that the brain initiates voluntary movements before the subject is conscious of

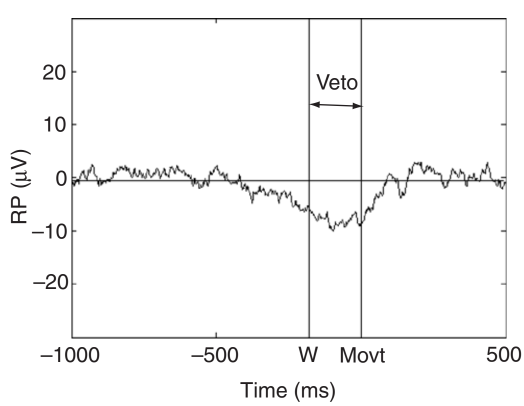

Figure 3 Relationship of Libet's time *W* and veto window to the EEG RP. RP is evident in average of 50 trials recorded from the vertex, presented with negative voltage shown in the downward direction (in contrast to the earlier convention of presenting negative as up). In Libet's terminology this is a Type II RP, produced by spontaneous rather than preplanned movements. Type II RPs start about 500 ms before the movement.

having willed them. In other words, the brain does all the work of starting a movement un- or preconsciously, and keeps consciousness informed of what it has done only as a sort of professional courtesy. On this interpretation, it seems that consciousness has no function when it comes to voluntary movement. If there is such a thing as a will that causes bodily movements, this will cannot be a conscious entity.

Libet himself believed that his results did demonstrate the unconscious initiation of voluntary acts. However, he remained convinced that consciousness must have a function. To reconcile these two opposing beliefs, he proposed that consciousness has the opportunity to veto the action initiated by the brain, in the approximately two-tenths of a second between time *W* and the actual movement (see Figure 3). This idea presently remains hypothetical, in part because it is not possible to study the neural activity associated with a veto by back-averaging off vetoed movements, since a vetoed movement by definition never takes place.

Other workers have been less convinced that Libet's original result does demonstrate the unconscious initiation of voluntary acts. The result itself is not in question – the main experimental finding has now been replicated in a number of different laboratories. RPs do, undeniably, start before time *W*. But controversy still surrounds the interpretation of this fact. The basic problems are

1. The relationship between RPs and the neural activity leading up to a movement remains to be elucidated. Specifically, it is not clear whether the start of the RP does represent the initiation of the movement, or whether it simply indicates a readiness for the movement to be initiated at some later point in time. There do exist at least two types of expectation-related waveform that resemble RPs, but precede events that are not movements.
2. It is not entirely clear that subjects do actually experience at time *W* subjective events that could be called urges or wantings. Recent experiments show that reports of *W* can be influenced by external events that take place after the movement. This suggests that *W* reports may represent something more akin to

a cognitive reconstruction after the event of what the subject infers must have happened. Perhaps the subject never actually experienced an urge to move, but in an effort to please the experimenter simply indicated that he or she had done so, at a time subjectively perceived as being a little before the movement itself.

Work is actively proceeding on these issues, but at this point it is not far enough advanced for its conclusions to be enshrined in an encyclopedia.

One obvious question in this area that has not so far been adequately addressed concerns the nature of the neural activity that is going on at time *W.* If there is a particular neural event that always coincides with the subjectively reported urge, it seems a reasonable bet that the urge might be a real subjective event, generated by (or identical with, if you espouse the neural identity theory of consciousness) the neural event. It might seem that it should be fairly simple to ascertain exactly what is happening in the brain at time *W.* Unfortunately, technical considerations mean that this is far from the case.

The problems with measuring the nature and location of the brain activity occurring at any particular time arise from the characteristics of the three noninvasive techniques that are currently used to measure brain activity in humans. The newest of these noninvasive techniques is fMRI. This measures blood oxygen level dependent (BOLD) signals, which essentially means it measures the increased blood flow that accompanies increased neural activity. The spatial resolution of BOLD is of the order of millimeters, but the temporal resolution is very bad, owing to the variable and indeterminate period of time (on the order of a couple of seconds) necessary for blood flow to a given brain area to increase once that area becomes active. The oldest of the noninvasive techniques is EEG. This measures the voltage difference between chosen points on the scalp. The temporal resolution of EEG can be of the order of microseconds (although millisecond resolution is more usual), but the spatial resolution is of the order of tens of millimeters, simply because the scalp is so far away from the brain. It is often thought that the main problem contributing to the poor spatial resolution of EEG is the smearing action of the skull, but in fact the real difficulty is the point spread function due purely to the distance between the site of waveform generation in the cortex and the measuring electrodes on the scalp. It is not widely appreciated that this distance is 15–20 mm, while the width of cortex generating most EEG waveforms is only 2–3 mm. The third noninvasive technique in our arsenal is magnetoencephalography (MEG). This uses superconducting quantum interference device (SQUID) sensors to measure directly the magnetic component of the electromagnetic fields generated by the brain. Since magnetic fields are unaffected by the skull, it is widely believed that the spatial resolution of MEG is substantially better than that of EEG. In fact it is about the same, because the main smearing influence in any electromagnetic measurement is the distance between the field generation site and the sensors, and this distance is actually slightly increased in MEG because the necessity to cool SQUIDs with liquid helium means that MEG sensors can not be placed directly on the scalp. MEG does however have the advantage that its measurements are reference-free. The requirement to measure EEG signals between recording sites and a reference site is the source of many problems. The spatial resolution of both EEG and MEG can theoretically be increased by mathematical solution of the inverse problem, which allows localization of the source of the electromagnetic fields in the brain. But the inverse problem is notoriously underdetermined, which means that in principle there are an infinite number of solutions that fit any given dataset.

Returning to our question, it can be seen that none of these noninvasive techniques is optimal for determining exactly what brain areas are active at time *W.* Either the spatial resolution is good but the temporal resolution is too low (fMRI) or the temporal resolution is good but the spatial resolution is too low (EEG and MEG). More precise spatial localization with preserved temporal resolution could potentially be achieved using electrocorticography, which is essentially the recording of EEG data from the surface of the brain instead of from the scalp. This is occasionally done for the clinical purpose of localizing epileptic foci prior to their excision, but to date published data from this technique do not contain enough detail to show exactly what brain areas are active 150–200 ms before a voluntary movement.

What data are available from the noninvasive techniques suggest that at time *W* activity is occurring in the SMA and/or the primary motor area (MI). The suggestion that SMA activity might be the neural correlate of an urge to move is supported by the finding that low-level electrical stimulation of the SMA in awake human patients does sometimes elicit verbal reports of an urge to move. However, higher intensity stimulation of the same areas invariably causes actual movement, and so it is possible that downstream activation of the primary motor area by the low-level stimulation might be the real correlate of the reported urges, or even that very small actual movements might be misinterpreted by the patients as urges.

But whatever the eventual outcome here, the neural activity underlying (or at least occurring at the same time as) the urges in Libet's experiments can only be related to the question of when a predetermined movement should be made. This sort of putative urge has nothing to do with the preceding decisions about what movement should be made, or that a movement should be made at all. In the case of Libet's experiments, those decisions were made many hours before any of the particular movements was initiated. The desires (to please the experimenter), choices (to participate in the experiment), and intentions (to follow the experimental instructions) occurred anything up to several weeks before the 'spontaneous' urge to make any particular movement. What is the relationship of consciousness to these earlier events?

The Will (Desires, Choices, and Intentions) – Experienced or Inferred?

Given that it is lamentably common to form a willed intention to do something (one's tax return springs to mind) – but then somehow never to get around to actually doing it – it seems clear that action initiation is not the same thing as willed intention. Libet-style experiments study action initiation, but they do not study the willed intention that preceded it (in this case the intention to do as the experimenter asks and move one finger repeatedly while watching a clock).

Wegner carried out a different set of experiments, aimed at finding out how conscious we are of our willed intentions. His thesis was that the experience of having caused an action is not a direct introspection of any particular brain events, but rather an inference like any other inference of cause and effect. The suggestion is that we think that A causes B if and only if

1. A occurs just before B,
2. A is consistent with B, and
3. there is no other apparent cause of B.

Likewise, we believe we have caused a given event if and only if

i. we think about the event just before it happens,
ii. our thought is consistent with what happens, and
iii. there is no obvious external cause for what happens.

In a test of condition (i), Wegner and colleagues showed that subjects could readily be fooled into thinking they had caused a computer cursor to stop over a particular object (when in fact the experimenter had caused the cursor to stop) if the name of the object was played into the subject's earphones just before the cursor stopped. Actually the subjects were only 56% sure that they had caused the stop both when the trick above was played and when they themselves actually had caused the stop, which again suggests that we are quite bad at introspecting our own intentions. In a test of condition (ii), subjects who viewed the experimenter's gloved hands in the position where their own would normally be could be fooled into thinking they were controlling the gloved hands, provided the gloved hands moved in accordance with a set of instructions played into the subject's earphones. Again, people were not very good at knowing whether or not they were controlling the hands even in the baseline condition. But when the subjects were led to think about what the hands were doing just before they did it, the perception of control increased significantly (although still only to about 3 on a scale from 1 = no control to 7 = complete control).

These and other experiments have led Wegner to propose the model in **Figure 4**. The suggestion is that actions and thoughts about actions are

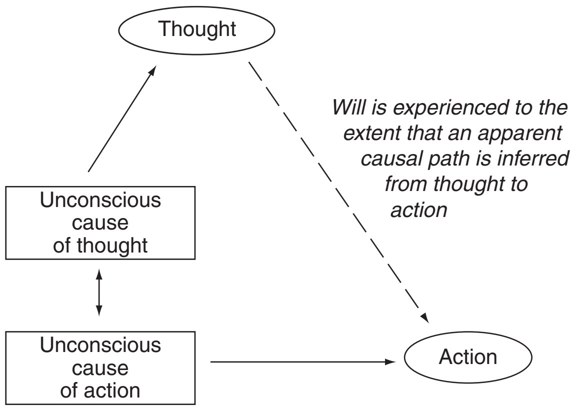

Figure 4 Wegner's model of the generation of voluntary actions. Unconscious events are shown in rectangular boxes. Conscious events are shown in ellipses.

generated by related but separate unconscious processes that act in parallel. The sensation of having caused the action by an effort of will is then generated by a third process, which infers causation of an action by a thought in the same way it might infer causation of tree movement by the wind.

Ironically, if Wegner's model is accurate, it has the contradictory effects of

1. suggesting that conscious intentions are often inferred rather than introspected, and thus that consciousness does not cause actions, and
2. providing at least indirect support for the idea ('Initiation of movements,' section point 2) that Libetian urges to move are also inferred rather than introspected, and thus throwing doubt on Libet's conclusion that consciousness does not cause actions.

This is all very confusing! Perhaps the safest conclusions at this point are (1) it is doubtful whether we really do have conscious access to either our urges or our intentions, and thus (2) it is doubtful whether consciousness plays any vital role at all in voluntary movement.

Philosophical Arguments

Although this is primarily an article about the brain basis of voluntary movement, it is instructive to consider briefly some of the philosophical arguments that have been put forward in opposition to the empirically derived conclusions in the preceding sections – not least because at root, these philosophical complaints constitute arguments against ever accepting any conclusions based on empirical data.

For example, one contributor to the book *Does Consciousness Cause Behavior?* says that we should not conclude anything about whether or not consciousness causes behavior because we do not yet know enough of the physiological facts. This would be fair enough if he did not himself then conclude that therefore there is no reason to throw out our prescientific intuitions that we do consciously cause our own acts. Another argument put forward in the same book is that Wegner's experiments are irrelevant because they put the subjects in abnormal situations. But all scientific experiments put their subjects in abnormal situations. Other contributors complain that while the experimental data that are available clearly show that some voluntary movements are not caused by consciousness, they do not show that all voluntary movements are not caused by consciousness. Such complaints raise the ancient specter of induction. The problem with induction is that one can never prove by induction (i.e., by making a finite number of experimental observations) that, for example, all swans are white, because there is always the chance that tomorrow one will meet a black swan. (This famous northern-hemisphere-grown example is particularly enjoyed by philosophers who hail from the southern hemisphere, where black swans are endemic). The attentive reader will by now have noted that the problem with all of these arguments is that science in general proceeds exactly by induction and experimentation, and so far science has proven itself a remarkably effective method of learning about the natural world. In fact, scientific experiments are the source of much, if not most, of our current understanding of how things work.

That being said, it is certainly fair to point out that science is never finished, and in the present case more work needs to be done before it is reasonable to conclude even on a heuristic basis that consciousness never has any role in voluntary movement. Before we can make such a sweeping conclusion, we need to know more about the brain basis of movement initiation. We need to know what causes the transition from willed intention to movement initiation. And we need to know whether or not it is possible in principle accurately to introspect our own urges, decisions, and intentions.

At present, the most we can say for sure is that the role of consciousness in voluntary movement is in question, and that this role is certainly smaller than that previously thought.

Legal Implications

The machinations of the Western legal system might seem remote from the day-to-day activities of brain scientists, but knowledge about the brain basis of the will is actually very important for the law. At present, most legal jurisdictions require the perpetrator of a crime to have consciously intended to do whatever he or she did in order for the act to be regarded as culpable. People are not generally put in jail because of their involvement in wholly accidental occurrences, when they did not consciously even put themselves in a position that could reasonably have been expected to lead to the accident (e.g., get into the car in an intoxicated state in the first place and start driving down the road).

This requirement for conscious intent makes the relationship of consciousness to action vitally important to the legal system. If it were to be accepted as a scientific fact that we never have conscious access to our intentions and/or to whatever neural activity initiates our actions, but simply have to infer these after the fact in the same way that we infer intent on the part of other people, then either the law would have to be changed, or nobody could ever be found guilty of anything.

So the question becomes, how sure are we that we do not have immediate access to our own intentions, or to the decisions that initiate our acts? As mentioned earlier, we are certainly not completely sure. But then no scientific conclusion is ever completely secure. Are we sure enough to recommend that the law be changed?

There are two points at issue here, which have different legal implications and thus should probably be considered separately:

1. Do we have conscious access to our long-term intentions?
2. Do we have conscious access to the events that initiate a given voluntary act?

Conscious Access to Long-Term Intentions?

How sure are we that we do (or do not) have conscious access to our long-term intentions? It is true that Wegner's experiments are of limited scope. But they build on a long tradition of research indicating that introspection of one's motives, intentions, and desires is significantly unreliable. People readily answer questions about why they did things, but as often as not their answers indicate that they are actually inferring rather than experiencing their own motives – indeed inferring them with little more accuracy than they could infer the motives of other people. Certainly, we are sometimes accurately aware of our own intentions and motives – but then we are sometimes accurate about other people's intentions and motives, too. The critical point is that we seem to have little direct introspective access to the thought processes involved in our own evaluations, judgments, and problem solving. We often do not know why we do what we do, or even that we intended to do it.

There may by now be enough data on this to render prudent a removal of the word conscious from the law relating to intent.

Conscious Access to Action Initiation

On our present scientific understanding it is conceivable that we also lack direct introspective access to the initiation of actions. However, less experimental evidence is available on this. More work needs to be done before it can justifiably be concluded that action initiation is not under direct conscious control.

See also: Free Will; Intentionality and Consciousness; Neuroscience of Volition and Action; Perception, Action, and Consciousness.

Suggested Readings

Huxley TH (1874) On the hypothesis that animals are automata, and its history. *The Fortnightly Review* 16: 555–580. Reprinted in *Method and Results: Essays by Thomas H. Huxley*. New York: D. Appleton and Company, 1898.

James W (1890) The Principles of Psychology. New York: Henry Holt & Co.

Libet B, Gleason CA, Wright EW, and Pearl DK (1983) Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). *Brain* 106: 623–642.

Nadel L and Sinnott-Armstrong W (in press) *Libet, Free Will and Responsibility*. Oxford: Oxford University Press.

Nisbett RE and Wilson TDeC (1977) Telling more than we can know: Verbal reports on mental processes. *Physiological Review* 84(3): 231–259.

Pockett S, Banks WP, and Gallagher S (2006) *Does Consciousness Cause Behavior?* Cambridge, MA: MIT Press.

Wegner DM (2002) *The Illusion of Conscious Will*. Cambridge, MA: MIT Press.

Biographical Sketch

Susan Pockett earned her PhD in neurophysiology from the University of Otago (New Zealand) and subsequently worked on synaptic plasticity in the physiology departments of the University of Oslo (Norway), the University of Auckland (New Zealand), University College London (England), the University of New South Wales (Australia), and the University of Manitoba (Canada). In 1994, she changed fields to work on consciousness, and is presently in the Physics Department at the University of Auckland.