Volition and the Function of Consciousness

(draft for book chapter to appear in Cognitive Neurosciences IV, Ed. Michael Gazzaniga, *MIT press*)

Hakwan Lau
Columbia University

Abstract

What are the psychological functions that could only be performed consciously? People have intuitively assumed that many acts of volition are not influenced by unconscious information. These acts range from simple examples such as making a spontaneous motor movement, to higher cognitive control. However, the available evidence suggests that under suitable conditions, unconscious information can influence these behaviors and the underlying neural mechanisms. One possibility is that stimuli that are consciously perceived tend to yield strong signals in the brain, which makes us think that consciousness has the function of such strong signals. However, if we could create conditions where the stimuli could yield strong signals but not the conscious experience of perception, perhaps we would find that such stimuli are just as effective in influencing volitional behavior. Future studies that focus on clarifying this issue may tell us what the defining functions of consciousness are.

Introduction

Many acts of volition seem to require conscious effort. We consciously initiate spontaneous motor movements. We cancel planned actions at will. We deliberately avoid particular actions. We intentionally shift our action plans in order to pursue different goals. Sometimes, theorists say, these are the functions of consciousness, as if evolution has equipped us with the gift of consciousness just to perform these acts. Without consciousness, presumably, we would only be able to perform much simpler actions that are no more sophisticated than embellished reflexes.

In this chapter we review available evidence to see if these intuitive claims are empirically supported. Recent studies in cognitive neuroscience suggest that many of these complex processes can actually be performed without consciousness. Or at least, many of them can be directly influenced by unconscious information. This calls into question the true function of consciousness, if not to enable us to deliberate our actions. We end by discussing what is logically required for an experiment to demonstrate the true function of consciousness.

Spontaneous Motor Initiation

Motor actions that are made not in immediate or direct response to external stimuli can be said to be spontaneously initiated. These are also sometimes called self-paced or self-generated actions. For instance, one may choose to casually flex one's wrist while sitting in a dark room, out of one's own free choice and timing, not to react to anything in particular. Some philosophers have argued that in cases like that, it should seem obvious that the action is caused by one's conscious intention (Searle 1983). Whereas one may argue that fast reactions to external stimuli may be driven by unconscious reflex (e.g. a runner leaping forward upon hearing the starting whistle), spontaneous actions do not seem to have any immediate cause but the conscious intention itself.

However, it has been shown that there is preparatory activity in the brain that starts at as early as 1-2 seconds before spontaneous actions are executed. This piece of one of the most perplexing findings in cognitive neuroscience was originally reported by Kornhuber and Deecke in the 1960s (Kornhuber and Deecke, 1965). They placed electrodes on the scalp to measure electroencephalography (EEG) while subjects made spontaneous movements at their own timing. The EEG data that were time-locked to the point of motor execution (as measured by muscle contraction indicated by electromyography, EMG) were averaged over many trials, which produced an event-related potential (ERP) known as the *Bereitschaftspotential* (BP) or readiness potential (RP). The readiness potential is slowly rising, peaking at around the point of action execution and starting from 1-2 seconds before that (fig 1). The readiness potential is most pronounced at electrodes near the vertex (Cz in the EEG coordinate system), which is directly above the medial premotor areas (including the supplementary motor area, SMA, pre-supplementary motor area, pre-SMA, and the cingulate motor areas below them). It is generally believed that one major source of the readiness potential lies in the medial premotor areas (Ball et al., 1999; Erdler et al., 2000; Weilke et al., 2001; Cunnington et al., 2003).

The demonstration of the readiness potential calls into question whether spontaneous movements are really caused by the preceding conscious intentions. Intuitively, conscious intentions seem to cause motor actions almost immediately - it seems to take much less time than 1-2 seconds. This could mean that the brain starts to prepare for the actions way before we consciously initiate them.

Benjamin Libet and colleagues empirically studied the timing of the conscious intention in relation to the readiness potential and the action (Libet et al., 1983). To measure the onset of conscious intention, he invented a creative but controversial paradigm which is sometimes called the "Libet clock paradigm". In those studies, subjects watched a dot revolving around a clock face at a speed of 2.56 second per cycle, while they flexed their wrist spontaneously (fig 2). After the action was finished, subjects were required to report the location of the dot when they "first felt the urge" to produce the action, i.e. the onset of intention. The subject might say it was at 3 o'clock or 4 o'clock position when they first felt the intention, for instance. This way the subjects could time and report the onset of their intention, and the experimenter could then work out actually when the action was produced, and hence the temporal distance between the two. Libet and colleagues reported that subjects on average report the onset of intention to be about 250 ms before motor execution.

Many people feel uncomfortable with the fact that the onset of the readiness potential seems to be so much earlier than the onset of intention, and some have tried to explain away the gap. Libet and colleagues have tried to study the onset of the readiness potential more carefully, discarding trials which might have been "contaminated" by pre-planning of action well before the action (for instance, by counting to 10 and then triggering the movement), as reported by the subjects. By only looking at the trials where the actions were supposed to be genuinely spontaneous, Libet and colleagues reported that the onset of the readiness potential is only about 500 ms before action execution (Libet et al., 1983). However, this is still clearly earlier than the reported onset of intention. And by discarding so many trials, it may be that the analysis just lacked the power to detect an earlier onset.

Some have argued that the onset of readiness potential might be an artifact due to the averaging needed to produce the ERP (Miller and Trevena, 2002). However, Romo and Schultz (Romo and Schultz, 1987) have recorded from neurons in the medial premotor areas while monkeys made self-paced movements. It was found that some of neurons in this region in fact fired as early as 0.6 – 2.6 seconds before movement onset. From the reported results it was also clear that this pattern of early firing for these neurons was consistent across trials. One other recent study has reported that the spatial pattern of fMRI activity from this region, at up to 5 seconds before action, can statistically predict the timing of action above chance level (Soon et al., 2008).

Others have argued that the readiness potential may not reflect the specific and causal aspects of motor initiation. However, as mentioned earlier, it is likely that the readiness potential partly originates from the medial premotor areas. Lesion to these areas can abolish the production of spontaneous actions (Thaler et al., 1995). These areas also contain neurons that code specific action plans (Shima and Tanji, 1998; Tanji and Shima, 1996). Further, when people use the Libet clock paradigm to time their own intentions, there is attentional modulation of activity in the medial pre-SMA (Lau et al., 2004), as if people were reading information off the area which is likely to be a source of the readiness potential.

The Libet clock method has also received considerable criticism. It involves timing across modalities, and could be susceptible to various biases (Libet, 1985; Gomes, 1998; Joordens et al., 2002; Klein, 2002; Trevena and Miller, 2002). However, it is unlikely that all these biases are in the direction that would help to narrow the gap between the onsets of the readiness potential and intention. Some have actually suggested that the different biases may point to different directions and thus just cancel each other out (Klein 2002). Also, in the original experiments by Libet and colleagues, there were control conditions that tested for the basic accuracy of the clock. They asked subjects to use the clock to time either the onset of movement execution, or in another condition to time the onset of tactile stimuli presented externally by the experimenter. Since the actual onsets of these events are objectively measurable, they could estimate the subjective error of onset reports produced by the clock method. They found the error to be in the order of about 50 ms, e.g. people misestimate the time of action execution to be 50 ms earlier than it actually is. This size of error is considerably much smaller than the gap between the onsets of the readiness potential and intention.

The basic results of Libet and colleagues have also been replicated in several different laboratories (e.g. Lau et al., 2004; Haggard and Eimer, 1998; Soon et al., 2008). In general, the same pattern is found, that the onset of intention is either around or later than 250 ms before action execution, which seems to confirm our intuition that conscious intentions seem to be followed by motor actions almost immediately. In fact, given that the readiness potential starts as early as 1-2 seconds before action execution, it is hard to imagine how the onset of intention could coincide or precede the readiness potential, unless one thinks of intention as a kind of prior intention (Searle, 1983), like the general plan that is formed at the beginning of the experimental session when the subject agrees to produce some actions in the next half an hour or so. We shall discuss this kind of higher-cognitive "intention" later in the chapter. However, the intention we are concerned with here is the immediate "urge" to produce the motor action (Libet et al., 1982).

Taken together, the evidence suggests that conscious intention, i.e. the immediate feeling of motor initiation, is unlikely to be the "first unmoved mover" in triggering spontaneous motor movements. It is likely to be preceded by unconscious brain activity that may contribute to action initiation. What, then, is conscious intention for?

Conscious Veto?

Libet's interpretation of the timing-of-intention results is that although intention may not be early enough to be the first cause of action, the fact that it is before action execution means that it could still be part of the causal chain. Maybe the decision to move is initiated unconsciously, but the awareness of intention may allow us to "veto", i.e. to cancel the action.

This seems to be a possibility. Libet and colleagues (Libet et al., 1983) as well as other researchers (Brass and Haggard, 2007) have performed experiments where subjects prepare for an action and then cancel it in the last moment, just before it is executed. The fact that we have the ability to "veto" an action seems beyond doubt. The question, however, is whether having the conscious intention is critical. Can the choice of veto be preceded by unconscious activity, just like the intention to act is preceded by the readiness potential? Or maybe sometimes actions are unconsciously vetoed, even without our awareness?

Some recent evidence suggests that the conscious intention may not facilitate a veto. As mentioned earlier, when people were using the Libet clock to time the onset of their intentions, there was attentional modulation of activity in the pre-SMA (Lau et al., 2004). These data have been subsequently further analyzed (Lau et al., 2006), and it has been shown that subjects who showed large degree of attentional modulation tended to also report the onset of intention to be early. One interpretation could be that attention biases the judgment of onset to be earlier. It was found in another experiment that this was also true when people used the Libet clock to time the onset of the motor execution. The higher the level of fMRI activity modulated by attention, the earlier subjects reported the onset to be, even though on average subjects reported the onsets to be earlier than they actually were, which means a bias to the negative (i.e. early) direction produced more erroneous rather more precise reports. In general, the principle of attentional prior entry (Shore et al., 2001) suggests that attention to an event speeds up its perception and negatively biases the reported onset. If this were true in the case of the Libet experiments, this could mean that attention might have exaggerated the 250ms onset, i.e. had subjects not been required to attend to their intentions in order to perform the timing tasks, the true onset of conscious intention may well be much later than 250ms prior to action execution. This calls into question whether we have enough time to consider the veto.

Another study reported that some patients with lesion to the parietal cortex reported the onset of intention to be late as 50 ms prior to action execution (Sirigu et al., 2004). If the awareness of intention allows one to veto actions, one might expect these patients to have much less time to consciously evaluate spontaneous intentions and cancel the inappropriate ones. This could be quite disastrous to daily life functioning. Yet there were no such reports about these patients.

Finally, in another study (Lau et al., 2007), single pulses of transcranial magnetic stimulation (TMS) were sent to the medial premotor areas (targeting the pre-SMA). Again, subjects were instructed to produce spontaneous movements and to time the onset of intentions and movement execution using the Libet clock. Surprisingly, although TMS was applied *after* motor execution, it has an effect on the reported onsets. No matter whether TMS was applied immediately after action execution or with a 200 ms delay, the stimulation exaggerated the temporal distance between the reported onsets of intention and movement, as if people reported a prolonged period of conscious intending. One interpretation may be that TMS injected noisy activity into the area and the intention monitoring mechanism did not distinguish this from endogenously generated activity that is supposed to represent intention. However, what is crucial is the fact that the reported onsets can be manipulated even after the action is finished. This seems to suggest that our awareness of intention may be constructed after the facts, or at least not completely determined before the action is finished. If conscious intentions are not formed before the action, they certainly cannot play any role in facilitating veto, let alone causing it.

This interpretation may seem wild, but it is consistent with other proposals. For instance, on the basis of many ingenious experiments manipulating subject’s sense of agency, Wegner (2002) has suggested that the conscious will is an illusion. The sense of agency is often inferred *post hoc*, based on many contextual factors. Wegner cites experiments to support these claims. One example is a study on "facilitated communication" (Wegner et al., 2003). Subjects (playing the role of "facilitators") were asked to place their fingers on two keys of a keyboard, while a confederate (playing the role of "communicator") placed his or her fingers on top of those of the subject. Subjects were given headphones with which they listened to questions of varying difficulty. Confederates were given headphones as well, and subjects were led to believe that the confederates would be hearing the same questions, although in fact the confederates heard nothing. Subjects were told to detect subtle, unconscious movements in the confederate's fingers following each question. When such movements were detected, the subject should press the corresponding key in order to answer on the confederate's behalf. It was found that subjects answered easy questions well above chance levels. If they had performed the task strictly according to the instructions, however, they should have performed at chance. Therefore, subjects must have been directing their own key presses. Nonetheless, they attributed a significant causal role for the key presses to the confederate. The degree to which subjects answered easy questions correctly was not correlated with the degree to which they attributed causal responsibility to confederates, suggesting that the generation of action and attribution of action to an agent are independent processes.

To summarize, although theorists have speculated that the awareness of intention may play some role in allowing us to cancel or edit our actions, considerable doubt has been cast by recent empirical evidence.

Exclusion and Inhibition

Another kind of situation that seems to require conscious deliberation involves the need to avoid a particular action or response. This is related to "vetoing" as described above, except that the action being inhibited is not necessarily self-paced, and may be specified externally. One example would be to perform stem completion while avoiding a particular word. So for instance, the experimenter may ask the subjects to produce any word starting with letter D (i.e. completing a 'stem'), but avoid the word 'dinner'. So subjects can produce 'dog', 'danger', 'dear', etc., but if they produce the word 'dinner', it would be counted as an error. This is called the exclusion task (Jacoby et al., 1992).

One interesting aspect of the exclusion task is that people can perform well only if they clearly see and remember the target of exclusion (i.e. the word ‘dinner’ in the foregoing example). If the target of exclusion is presented very briefly and followed by a mask, such that it was only very weakly perceived, people may fail to exclude it (Debner and Jacoby, 1994; Merikle et al., 1995). In fact, they tend to produce exactly the word they should be avoiding with higher likelihood than if they were not presented with the word at all. It has been argued that this exclusion failure phenomenon is the hallmark of unconscious processing (Jacoby et al., 1992). The weak perception of the target probably produced a representation for the word, but because the signal is not strong enough to reach the level of conscious processing, subjects are unable to inhibit the corresponding response.

In addition to the intuitive appeal, the notion that consciousness is required for exclusion is also supported by a case study of a blindsight patient (Persaud and Cowey, 2007). Subject GY has a lesion to the left primary visual cortex (V1), and reports that most of his right visual field is subjectively blind. However, in forced-choice situation he can discriminate simple stimuli well above chance level in his “blind” field (Weiskrantz, 1986; Weiskrantz, 1999). In one study he was required to perform an exclusion task (Persaud and Cowey, 2007), i.e. to say the location (up or down) where the target was *not* presented. Whereas he could do this easily in the normal field, he failed the task when stimuli were presented to his blind field. Note that he was significantly worse than chance in the blind field, as if the unconscious signal drove the response directly and inflexibly, defying exclusion control. This seems to support the account that consciousness is required for exclusion.

The general idea that inhibition requires consciousness seems to be supported by other studies too, including those that do not employ the exclusion paradigm. One study tested subjects' ability to ignore distracting moving dots, while doing a central task that has nothing to do with the distractors (Tsushima et al., 2006). It was found that if the motion of the distractor was above the perceptual threshold, people could ignore the dots and inhibit the distraction successfully. Somewhat paradoxically, when the motion was below perceptual threshold, people could not ignore the dots and were distracted. The results from brain imaging seem to suggest that when the motion of the stimuli was strong, it activated the prefrontal cortex, and triggered it to suppress the motion signal. When the motion of the stimuli was below perceptual threshold, however, the signal failed to trigger the inhibitory functions in the prefrontal cortex, and therefore the motion signal were not suppressed and thus remained distracting.

However, the notion that flexible control or inhibition of perceptual signal requires consciousness is not without its critics (Snodgrass 2002; Haase and Fisk, 2001; Visser and Merikle, 1999). One problem becomes clear when we consider the motion distractor example above. "Conscious signal" here seems to be the same thing as a strong signal, driven by larger motion strength in the stimuli. Obviously, signals have to be strong enough to reach the prefrontal cortex in order to trigger the associating executions functions. Do unconscious stimuli fail to be excluded because we are not conscious of them, or is it just because the signal is not strong enough? Or, are the two explanations one and the same? Not all study are subject to this argument. For instance, in the blindsight study mentioned above (Persaud and Cowey 2007), the subject failed to exclude in the blindfield even when the contrast level would give a performance that was similar to that in the normal visual field. So if we take forced-choice performance as an index of signal strength, the signal from the blindfield was not weak in this sense. However, in most other cases we often take awareness to be the same as good performance. Are we justified to do so? This is an important issue and we will come back to this in the final section of the chapter.

Other researchers have reported evidence that seems to support unconscious inhibition. For instance, in one study (Snodgrass and Shevrin, 2006) people were asked to detect visually presented words. In certain conditions, some subjects showed detection performance that was significantly *worse* than chance. These words were presented so briefly that typically detection performance would be near chance. We usually take chance-level as the objective threshold for conscious perception. Below chance-level performance could be taken as evidence that the subjects did not consciously perceive the words. And yet, if they had no information at all regarding the words, performance should just be exactly at chance rather than below. It seems that these subjects were actively suppressing the words.

These are unusual cases and are somewhat hard to interpret. We take chance-level as the objective threshold for conscious perception because when people perform at chance, it indicates that they do not have the explicit information regarding the target of perception. However, if people perform significantly below chance, it means that somehow they have the information regarding the detection, which violates the very logic we adopt to label perception unconscious. But in any case, the stimuli were supposed to be really weak, and it is intriguing that some subjects seem to be automatically suppressing the words. Are we to take these somewhat unusual cases as evidence to reject the notion that exclusion or inhibition requires consciousness? It seems that, logically, if we claim that a certain function *requires* consciousness, we should predict there will never be a case where one could perform such function unconsciously. How seriously are we to take this logic and reject functions as requiring consciousness by a single experiment? We will return to this argument in the last section of the chapter.

Top-Down Cognitive Control

So far we have discussed acts of volition that are relatively simple, like starting a motor movement, or avoiding a particular action. Sometimes we also voluntarily prepare for a set of rules or action plans in order to satisfy a more abstract goal in mind. For instance, a telephone ring may usually trigger a particularly action, e.g. to pick up the phone. However, when one visits friends at their homes, one may deliberately change the mapping between the stimulus (telephone ring) and action, i.e. it would be more appropriate to sit still, or ask the host to pick up the phone, rather than picking it up oneself. This volitional change of stimulus-response contingency is an example of top-down cognitive control.

It has been suggested that top-down cognitive control may require consciousness (Dehaene and Naccache, 2001). The idea is that unconscious stimuli can trigger certain prepared actions, as demonstrated in studies in subliminal priming (Kouider and Dehaene, 2007). However, the preparation or setting up of the stimulus-response contingency may require consciousness.

However, recent studies suggest that this might not be true, in the sense that unconscious information seems to be able to influence or even trigger top-down cognitive control too (Mattler, 2003; Lau and Passingham, 2007). In one study subjects had to prepare to do a phonological or semantic judgment, based on the orientation of a figure they saw (fig 3). In every trial, if they saw a square, they had to prepare to judge whether an upcoming word has two syllables (e.g. "table") or not (e.g. "milk"). If they saw a diamond, they had to prepare to judge whether an upcoming word refers to a concrete object (e.g. "chair") or an abstract idea (e.g. "love"). In other words, they had to perform top-down cognitive control based on the instruction figure (square or diamond). However, before the instruction figure was presented, there was actually an invisible prime figure, which could also be a diamond or a square. It was found that the prime could impair subjects' performance when it suggested the alternative (i.e. wrong) task to the subjects (incongruent condition). One could argue that this was only because the prime distracted the subjects on a perceptual level, and did not really trigger cognitive control. However, the experiment was performed in the fMRI scanner, and the brain recordings suggest that when being primed to perform the wrong task, subjects used more of the wrong neural resources too (Lau and Passingham, 2007). That is, areas that are more sensitive to phonological or semantic processing showed increased activity when the explicit instruction figure made subjects perform the phonological and semantic tasks respectively. The invisible primes also seem to be able to trigger activations in task sensitive areas. This seems to suggest that they can influence or exercise top-down cognitive control.

Another study examines how unconscious information affects our high-level objectives by focusing on how the potential reward influences our level of motivation (Pessiglione et al., 2007). Subjects squeezed a device to win a certain amount of money. The harder they squeezed, the more money they would win. However, the size of the stake in question for a particular trial was announced in the beginning by presenting the photo of a coin. The coin could either be a British pound (~2 US dollars) or a penny (~2 US cents), and it signified the monetary value of the maximal reward for that trial. Not surprisingly, people squeezed harder when the stakes were high, but interestingly, the same pattern of behavior was observed even when the figure of the coin was masked such that subjects reported not seeing it. This suggests that unconscious information can influence our level of motivation as well.

If unconscious information alone is sufficient to exercise all these sophisticated top-down control functions, why do we need to be conscious at all?

How to find the true function of consciousness?

The foregoing is not meant to be an exhaustive review of all studies on the potential functions of consciousness. We select some examples from a few areas that are particularly related to volition, and discuss what role consciousness may play. It may, of course, be that there are other psychological functions that require consciousness.

Yet, one cannot help but feel that there seems to be some inherent limitation to this whole enterprise of research. If we claim that a certain function requires consciousness, strictly speaking, the interpretation could be that the function should never be able to be performed unconsciously. Of course, one could make the weaker claim that a certain function is usually or most suitably performed consciously, and when consciousness fails, unconscious processing can act as a backup. This is similar to arguing that one function of having legs is to facilitate locomotion; if we lose our legs, we could still move around, albeit poorly. However, let us assume that one is to make the stronger prediction that such functions should never be able to be performed unconsciously. In principle, it would only take a single experiment to falsify that. This explains why this review may seem biased in that we focus on studies that show the power of the unconscious, rather than studies demonstrating what functions definitely require consciousness. In principle, falsifying the claim that a certain function requires consciousness is straightforward. But this is not the case for demonstrating functions that would always require consciousness.

One can, of course, try to show that subjects could normally do a task if the relevant information is consciously perceived. And then one tries to 'knock-out' the conscious perception for such information, and try to show that the task could no longer be performed, or that it is performed at an additional cost, i.e. slower or with more errors. But how would one know that in 'knocking-out' the conscious perception, one does not 'knock-out' too much? One typically suppresses conscious perception by visual masking, by using brief presentation, by distracting the subject, by applying transcranial magnetic stimulation, by pharmacological manipulations, etc. But all of these could potentially impair the unconscious as well as the conscious signal. Maybe in cases where the perception has been rendered unconscious, the signal is just no longer strong enough to drive the function in question? This would mean that, in principle, it would be possible for a future study to find the optimal procedure or setup to just render the information unconscious, without reducing the signal strength too much. And in that case the subjects may be able to perform the task in question. That would falsify our claim.

This means that in looking for functions that require consciousness, we need to adopt some different strategies. One potentially useful approach is to try to demonstrate something akin to a "double dissociation". When conscious perception is suppressed, we often find that a sophisticated function (e.g. top-down cognitive control) can no longer be performed, though some simpler function (e.g. priming for a prepared motor response) may still be activated by unconscious information. From the foregoing discussion, one could see that this may not be as surprising or informative as it seems. It could be just that the unconscious signal is just too weak to drive the relatively sophisticated function. A demonstration of the opposite would, however, be much more convincing: If after suppression of conscious perception, the subjects can still perform a rather sophisticated function, but fail to perform a simple function, that would suggest that the simple function really requires consciousness. In this case, it could not be that the suppression of conscious perception has taken away too much of the signal strength, because if that were the case then the subjects should not be able to perform the relatively sophisticated function (fig 4). Understanding this "double dissociation" approach help us to see the logic behind how we could deal with signal strength as a confounding variable. However, one problem is that it is unclear what is most convincing way to define "sophisticated/complicated" functions vs "simple" functions.

An alternative approach may be to directly match for signal strength between the conscious and the unconscious conditions. This might seem difficult because conscious signals may seem to be strong in general. However, as discussed above, blindsight subjects can perform forced-choice discrimination on visual stimuli well above chance, even when they claim that conscious awareness is missing. Forced-choice performance is often taken as an objective estimate of signal strength; the detection theoretical measure d' is mathematically just the signal-to-noise ratio. In blindsight subject GY, where only half of the visual field lacks awareness, we can imagine presenting weak stimuli to the normal visual field such that forced-choice performance would match that in the blind field (Weiskrantz et al., 1995). This way we can test if certain functions cannot be performed based on information presented to the blind field, which may shed light on when consciousness is required.

One may argue that blindsight patients are rare and the way their brains process visual information may not generalize to intact brains. However, there are other paradigms where in normal subjects one could match for forced-choice performance, and yet produce a difference in the level of conscious awareness. For instance, in one study (Lau and Passingham, 2006) metacontrast masking was used to create similar conditions where forced-choice discrimination accuracy for the visual targets were matched, and yet the subjective reports of how often subjects saw the identity of the targets differed (fig 5). One could imagine presenting these stimuli to subjects and seeing if they drive a certain function with different effectiveness. If the subjects perform better in the condition where subjective conscious awareness of the stimuli is more frequent, one could argue that this function is likely to depend critically on consciousness.

Conclusion

Acts of volition are accompanied by a sense of conscious effort or intention. The fact that we feel the conscious effort is not in doubt. What is less clear is whether the processes underlying the conscious experience directly contribute to the execution of the actions, in a way that is not accomplished by unconscious processes just as effectively. The general picture seems to be that many sophisticated functions can be performed unconsciously or driven by unconscious information.

Does this mean that consciousness has no special function at all? The answer is not yet clear. It is likely that some psychological functions do require consciousness. That is, there may be some functions that can only be performed poorly with unconscious information. Or, there may even be functions that can never be performed unconsciously. But experiments have not yet been able to convincingly pin them down.

They will have to overcome the following problem. If we assume that conscious perception is always accompanied by stronger and longer-lasting signals that are more effective than unconscious signals in propagating themselves throughout the brain, then certainty, consciousness would certainly be associated with the functions of these strong signals. However in studies of blindsight (Weiskrantz et al., 1995) as well as in normals (Lau and Passingham, 2006) it has been shown that signal strength as indicated by forced-choice performance is not always one and the same as conscious awareness. Therefore, future studies may need to focus on identifying the functions that really cannot be performed unconsciously, even when the signal strength is sufficiently strong. This may help to reveal the true function of consciousness.

Acknowledgements

The author thanks David Rosenthal and Uriah Kriegel for comments.

Figures

Figure 1. A schematic depiction of the readiness potential (RP) preceding spontaneous movements. The RP is usually recorded at the top of the scalp, above medial frontal premotor areas. It gradually ramps up, beginning about 1-2 seconds before movement and peaking around the time of movement execution (marked as time = 0 above).

Figure 2. The Libet clock paradigm. **A.** The subject views a dot rotating slowly (2.56 seconds per cycle) around a clock face and waits for an urge to move to arise spontaneously. When the urge arrives, the subject makes a movement (e.g. a key press). **B.** After making the movement, the subject estimates the earliest time at which the intention to move was experienced. To carry out this time estimate, the subject either verbally indicate the location of the dot where the intention was first felt, or move a cursor to that location (as in this example). In a common control condition, the subject uses the clock to estimate the time of movement rather than the onset of intention. Figure edited and adapted from Lau et al., 2007.

Figure 3. Experimental paradigm of Lau and Passingham (2007). Subjects view

briefly presented words and perform either a phonological task (is the word one syllable or two syllables?) or a semantic task (does the word name something concrete or abstract?). Before word presentation, subjects are instructed which task to perform on a given trial by a visual symbol (a square for the phonological task, or a diamond for the semantic task). The symbolic instruction itself acts as a metacontrast mask for an earlier prime, also a square or a diamond. Because the prime is briefly presented and masked, it is not consciously perceived. On half of trials, the prime is congruent with the instruction and on the other half, incongruent. Behavioral and imaging results suggest that the unconscious primes affected top-down task switching. When primes were incongruent with instructions, accuracy fell, reaction time increased, and brain regions corresponding to the task indicated by the prime were partially activated (all relative to the prime-congruent condition). But when the stimulus onset asynchrony (SOA) between prime and instruction was lowered, such that primes became visible, the priming effect was not evident. This double dissociation suggests that the interference of incongruent primes on task switching cannot be attributed to conscious processing. Figure adapted from Lau and Passingham, 2007.

Figure 4. (A) The normal situation for conscious perception. Stimuli are strong enough to drive processes of different complexity. (B) A typical situation for

unconscious perception. Stimuli are weak such that complicated processes are no longer activated, though simple processes can still be triggered. It could be argued that this is not surprising as we may expect that complicated processes require a stronger signal. (C) A potentially more informative situation. If one could find a stimulus that is not consciously perceived, but yet is sufficiently strong to trigger a complicated process, then the relatively simple process that the stimulus does not drive would seem to critically depend on consciousness.

Figure 5. Inducing "relative blindsight" in normal observers using metacontrast masking. **A.** Metacontrast masking paradigm. The subject is presented with a visual target (in this case, either a square or diamond). Afterwards, a metacontrast mask is presented. The mask differentially affects discrimination accuracy and visual awareness of the target as a function of stimulus onset asynchrony (SOA). **B.** Discrimination accuracy and visual awareness as a function of metacontrast mask SOA. The metacontrast mask creates a characteristic U-shaped function of performance vs. SOA. At shorter and longer SOAs, discrimination accuracy is high, but it dips at intermediate SOAs. The same is true for visual awareness, but the shape of the awareness masking function is not perfectly symmetrical with respect to the performance masking function. That is, there are certain SOAs at which forced choice performance is matched, but visual awareness differs

significantly (e.g. as illustrated in the SOAs of 33 ms and 100 ms). Such performance-matched conditions could be used to investigate the functions of consciousness. If some task can be performed better in the condition of higher subjective visibility, it can plausibly be said to require visual awareness. Because forced-choice discrimination accuracy is matched across the two conditions, the superior performance of the task in the high visibility condition cannot be attributed to a difference in signal strength. Figure adapted from Lau and Passingham, 2006.

References

Ball, T., Schreiber, A., Feige, B., Wagner, M., Lücking, C. H., & Kristeva-Feige, R. (1999). The role of higher-order motor areas in voluntary movement as revealed by high-resolution EEG and fMRI. *NeuroImage*, *10*(6), 682-94.

Brass, M., & Haggard, P. (2007). To do or not to do: the neural signature of self-control. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *27*(34), 9141-5.

Cunnington, R., Windischberger, C., Deecke, L., & Moser, E. (2003). The preparation and readiness for voluntary movement: a high-field event-related fMRI study of the Bereitschafts-BOLD response. *NeuroImage*, *20*(1), 404-12.

Debner, J. A., & Jacoby, L. L. (1994). Unconscious perception: attention, awareness, and control. *Journal of experimental psychology. Learning, memory, and cognition*, *20*(2), 304-17.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*(1-2), 1-37.

Erdler, M., Beisteiner, R., Mayer, D., Kaindl, T., Edward, V., Windischberger, C., et al. (2000). Supplementary Motor Area Activation Preceding Voluntary Movement Is Detectable with a Whole-Scalp Magnetoencephalography System. *NeuroImage*, *11*(6), 697-707.

Gomes, G. (2002). The interpretation of Libet's results on the timing of conscious events: a commentary. *Consciousness and cognition*, *11*(2), 221-30; discussion 308-13, 314-25.

Haase, S. J., & Fisk, G. (2001). Confidence in word detection predicts word identification: implications for an unconscious perception paradigm. *The American journal of psychology*, *114*(3), 439-68.

Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, *126*(1), 128-33.

Jacoby, L. L., Lindsay, D. S., & Toth, J. P. (1992). Unconscious influences revealed. Attention, awareness, and control. *The American psychologist*, *47*(6), 802-9.

Joordens, S., van Duijn, M., & Spalek, T. M. (2002). When timing the mind one should also mind the timing: biases in the measurement of voluntary actions. *Consciousness and cognition*, *11*(2), 231-40; discussion 308-13.

Klein, S. (2002). Libet's research on the timing of conscious intention to act: a commentary. *Consciousness and cognition*, *11*(2), 273-9; discussion 304-25.

Kornhuber, H., & Deecke, L. (1965). Hirnpotentialänderungen bei Willkurbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. *Pflügers Archive*, *284*, 1-17.

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *362*(1481), 857-75.

Lau, H. C., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(49), 18763-8.

Lau, H. C., & Passingham, R. E. (2007). Unconscious activation of the cognitive control system in the human prefrontal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *27*(21), 5805-11.

Lau, H. C., Rogers, R. D., Haggard, P., & Passingham, R. E. (2004). Attention to intention. *Science (New York, N.Y.)*, *303*(5661), 1208-10.

Lau, H. C., Rogers, R. D., & Passingham, R. E. (2006). On measuring the perceived onsets of spontaneous actions. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *26*(27), 7265-71.

Lau, H. C., Rogers, R. D., & Passingham, R. E. (2007). Manipulating the experienced onset of intention after action execution. *Journal of cognitive neuroscience*, *19*(1), 81-90.

Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, *8*, 529-566.

Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain : a journal of neurology*, *106 (Pt 3)*, 623-42.

Libet, B., Wright, E. W., & Gleason, C. A. (1982). Readiness-potentials preceding unrestricted 'spontaneous' vs. pre-planned voluntary acts. *Electroencephalography and clinical neurophysiology*, *54*(3), 322-35.

Libet, B., Wright, E. W., & Gleason, C. A. (1983). Preparation- or intention-to-act, in relation to pre-event potentials recorded at the vertex. *Electroencephalography and clinical neurophysiology*, *56*(4), 367-72.

Mattler, U. (2003). Priming of mental operations by masked stimuli. *Perception & psychophysics*, *65*(2), 167-87.

Merikle, P. M., Joordens, S., & Stolz, J. A. (1995). Measuring the relative magnitude of unconscious influences. *Consciousness and cognition*, *4*(4), 422-39.

Miller, J., & Trevena, J. A. (2002). Cortical Movement Preparation and Conscious Decisions: Averaging Artifacts and Timing Biases. *Consciousness and Cognition*, *11*(2), 308-313.

Persaud, N., & Cowey, A. (2008). Blindsight is unlike normal conscious vision: evidence from an exclusion task. *Consciousness and Cognition*, *17*(3), 1050-5.

Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., et al. (2007). How the brain translates money into force: a neuroimaging study of subliminal motivation. *Science (New York, N.Y.)*, *316*(5826), 904-6.

Romo, R., & Schultz, W. (1987). Neuronal activity preceding self-initiated or externally timed arm movements in area 6 of monkey cortex. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, *67*(3), 656-62.

Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind* (p. 278). Cambridge University Press.

Shima, K., & Tanji, J. (1998). Both supplementary and presupplementary motor areas are crucial for the temporal organization of multiple movements. *Journal of neurophysiology*, *80*(6), 3247-60.

Shore, D. I., Spence, C., & Klein, R. M. (2001). Visual prior entry. *Psychological science : a journal of the American Psychological Society / APS*, *12*(3), 205-12.

Sirigu, A., Daprati, E., Ciancia, S., Giraux, P., Nighoghossian, N., Posada, A., et al. (2004). Altered awareness of voluntary action after damage to the parietal cortex. *Nature neuroscience*, *7*(1), 80-4.

Snodgrass, M. (2002). Disambiguating conscious and unconscious influences: do exclusion paradigms demonstrate unconscious perception? *The American journal of psychology*, *115*(4), 545-79.

Snodgrass, M., & Shevrin, H. (2006). Unconscious inhibition and facilitation at the objective detection threshold: replicable and qualitatively different unconscious perceptual effects. *Cognition*, *101*(1), 43-79.

Soon, C. S., Brass, M., Heinze, H., & Haynes, J. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11*(5), 543-5.

Tanji, J., & Shima, K. (1996). Supplementary motor cortex in organization of movement. *European neurology*, *36 Suppl 1*, 13-9.

Thaler, D., Chen, Y. C., Nixon, P. D., Stern, C. E., & Passingham, R. E. (1995). The functions of the medial premotor cortex. I. Simple learned movements. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, *102*(3), 445-60.

Trevena, J. A., & Miller, J. (2002). Cortical movement preparation before and after a conscious decision to move. *Consciousness and cognition*, *11*(2), 162-90; discussion 314-25.

Tsushima, Y., Sasaki, Y., & Watanabe, T. (2006). Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science (New York, N.Y.)*, *314*(5806), 1786-8.

Visser, T. A., & Merikle, P. M. (1999). Conscious and unconscious processes: the effects of motivation. *Consciousness and cognition*, *8*(1), 94-113.

Wegner, D. M. (2002). *The Illusion of Conscious Will*. MIT Press.

Wegner, D. M., Fuller, V. A., & Sparrow, B. (2003). Clever hands: uncontrolled intelligence in facilitated communication. *Journal of personality and social psychology*, *85*(1), 5-19.

Weilke, F., Spiegel, S., Boecker, H., von Einsiedel, H. G., Conrad, B., Schwaiger, M., et al. (2001). Time-resolved fMRI of activation patterns in M1 and SMA during complex voluntary movement. *Journal of neurophysiology*, *85*(5), 1858-63.

Weiskrantz, L., Barbur, J. L., & Sahraie, A. (1995). Parameters affecting conscious versus unconscious visual discrimination with damage to the visual cortex (V1). *Proceedings of the National Academy of Sciences of the United States of America*, *92*(13), 6122-6.

Weiskrantz, L. (1986). *Blindsight: A Case Study and Implications* (p. 187). Oxford University Press.

Weiskrantz, L. (1997). *Consciousness Lost and Found: A Neuropsychological Exploration* (1st ed., p. 304). Oxford University Press, USA.