Will and Emotions:
A Machine Model that Shuns Illusions

Igor Aleksander
Dept. Of Electrical and Electronic Engineering,
Imperial College , London SW7 2BT
i.aleksander@imperial.ac.uk

Mercedes Lahnstein
Dept. Of Electrical and Electronic Engineering,
Imperial College , London SW7 2BT
m.lahnstein@imperial.ac.uk

Rabinder Lee
Dept. Of Electrical and Electronic Engineering,
Imperial College , London SW7 2BT
r.lee@imperial.ac.uk

Abstract

Benjamin Libet discovered a neo-cortical 'readiness potential' associated with the spontaneous movement of a finger (Libet et al, 1983). As this happens approximately 350ms before the participant becomes conscious of willing the action, has led Dan Wegner (2002) to develop an illusion-based hypothesis of volition. This paper suggests that the readiness potential is emotional in nature and appropriately unconscious, removing the need to evoke illusions. A machine model is developed which shows how an emotional readiness potential might relate to a legitimate sensation of causation.

1 Introduction

We believe that the problem with Libet's discovery (Libet et al, 1983) of a neo-cortical 'readiness potential' associated with the spontaneous movement of a finger, became controversial because it hinges on a volitional task that does not involve emotion. In general volition does, and it is the objective of this paper to show that if a model incorporating emotion is developed, Libet's Readiness Potential may have an emotional basis. Wegner's (2002) illusion-based hypothesis involves the existence of an unconscious cortical event which both controls a conscious sensation *and* the resulting action which the volitional organism mistakenly interprets as the action being *caused* by the sensation of volition. This, he argues, is akin to believing that traffic lights have changed to green as a result of the will of the observer. Here our addition to Libet and Wegner's reasoning is that choices in an act of volition involve emotional evaluation and this leads to an illusion-free hypothesis. We have developed a mechanistic model which tests a non-illusory hypothesis which still accords with Libet's results and respects a group of basic *synthetic phenomenology* (depictive) rules for conscious mechanisms (Aleksander and Dunmall, 2003 [A&D]).

2 The illusion interpretation

2.1 Libet's data

Libet (1983) wanted to measure the time it took between wanting to do something and doing it. He devised an experiment where the decision to do something had an arbitrariness about it. In line with other similar experiments, he measured the electro-encephalographic recording (eeg) that related to wanting to lift a finger at a time arbitrarily selected by the participant, and the moment of lifting the finger.

Normally this experiment would measure the time of the brain activity when wanting to lift the finger (Bw) and then, the time of both the brain activity that lifts the finger (Bl) and the actual moment of lifting the finger (L). It was known that Bl would occur a small fraction of a second after Bw followed another fraction of a second later by by L. All of this accords with the folk idea that we need brain activity to want something and that a little later this causes other brain activity that activates the muscles that unleash the desired physical action.

To add greater interest to the experiment Libet invented an ingenious way of measuring the moment at which the conscious thought of lifting the finger occurred to the participant. He asked the participant to observe a dot moving in a circular trajectory on the screen, the trajectory being marked with numbers and to note the number when the conscious thought of lifting the finger occurred. This acted like a clock. Of course the experimenter expected this reading to coincide with the brain activity Bw. The surprise came when it transpired that (Bw) occurred half to three-quarters of a second *before* the participant became conscious of 'wanting'. This brain activity then became known as the *Readiness Potential.*

An obvious interpretation of this result is that the conscious will to do something is not the free event we feel, but it is dependent on an unconscious occurrence in the brain that is initiated in a way that is as yet not properly understood.

2.2 Wegner's interpretation

Dan Wegner (2002), suggested an important philosophical significance of Libet's finding. He proclaimed a 'theory of apparent mental causation':

> *"People experience conscious will when they interpret their own thought as the cause of their action"*

So a totally unconscious neural event causes the wanting and it causes the action a little later. This feels as if the action is caused by the wanting, but the link entirely illusory.

We express a little scepticism on the rush to interpret Libet's results as defining free will as an illusion. First one needs to question exactly what it is that is measured and labelled as Bw. Second, we are able to model at least one mechanism in which the kind of delays encountered by Libet would arise in the course of a non-controversial scheme that links consciousness of a desire to eventual action. The missing feature is the evaluation of emotions associated with available choices. For this we appeal to some 'axiomatic neuromodelling'.

3 Axiomatic neuromodelling

3.1 Axioms: a resumé

While details of our axiomatic approach are fully set out in [A&D], we include a brief resumé here for the sake of completeness. The five axioms are a result of answering the question "What is important to me about my consciousness" and then asking what known informational mechanisms might exist which *are necessary* to sustain such sensations. Sufficiency for these mechanisms is not claimed, but necessity leads to composite systems in which these mechanisms interact to suggest a design for a conscious organism. In this paper they are only used to the extent that we wish to suggest a non-illusory structure that links conscious thought to action. The word 'axiom' is used not in the sense that one might in the starting point of a logical proof, but more in a sense of a set of starting points that are inwardly felt to be fundamental for the design of a model that reflects introspective features of being conscious.

3.1.1 Axiom 1: Being in an out-there-world

What seems a central feature of being conscious is the fact that sensations appear to be situated where they are in the world rather than like pictures or representations in our head. The property required for this is that some neurons not only react to the presence or absence of minimal source events in the world, but they do this conditionally on the position occupied by such an element. The mechanism required for this relies on the presence of sensors that are sensitive to such position as, for example, binocular vision, eye, neck and body movement in vision. As indicated in [A&D] there are neurons in the brain that are selective in this way. This special representation of being in an out-there-world has been called *depictive.*

3.1.2 Axiom 2: Imagining

Referring to visual sensation for a while, it is clear (speaking introspectively) that, if I close my eyes, the visual world does not go away: I can imagine what things look like, that is, what they looked like at some time in the past. The sensation is not quite as vivid as when I am actually looking at something, but there nonetheless.

These 'visions' need not go away when I do open my eyes. Indeed they are part of my visual interaction with the world out there. I sometimes loose my keys and look for them. I form a mental image of what they might look like when I eventually do see them. If I should see a different bunch of keys, the differences between the depiction of these and the mental image are intensely, almost painfully felt. When seeing a well known face, it is known that I can form a sufficiently appropriate mental image of the person even before my fovea has had a chance to look at every feature. That is, the mental image snaps in.

There is another aspect to these inner sensations: they can construct something we may never have seen or experienced as when reading noverls. This is a case where visions are generated by words, but visions could be generated by any of the sensory modalities: the smell of freshly baked bread can trigger scenes from childhood, touching a slimy surface in the dark can create nightmarish visions of unpleasant gutters.

The material implication of these inner visions and memories is that of *feedback* or *re-entry* in depictive neural structures. Having a mental image of something that has happened in the past has a strong material implication: closed information paths in depictive networks must exist which can *sustain* depictive firing patterns. Indeed, axiom 2 could be simply rephrased to say: 'no depiction and no feedback – then no imagination'.

3.1.3 Axiom 3: Attending to input

So far, we have spoken of worlds out there as if the conscious organism just blunders around in them. Nothing is further from the truth. Selecting what we experience in the world and how we think about the world in our imagination, requires some selection mechanisms: attention. The technical detail of how attention is achieved in an external and internal sense is beyond the needs of this paper. Suffice it to say that such mechanisms are largely unconscious. The most telling are eye movements: largely driven by the superior colliculus, they could be determined by the content of the perifovea of the eye (high spatial frequency) by the extrastriate cortex (supply missing meaning) or even the auditory cortex (eyes move in the direction of a sudden noise). But all we feel and experience is the foveal depictive reconstruction of axiom 1.

3.1.4 Axiom 4: Thinking ahead

This axiom and the next are central to this paper, as the paper is directed towards their elaboration and assessment with respect to Libet's data. Thought is not just a process of having static depictions. It is a highly dynamic process. We are constantly thinking ahead, considering alternatives and, every now and then, deciding what to do next. What are the material implications of this possibility? TIn fact, no new machinery needs to be evoked over and above that which we have seen in axiom 2: re-entrant neural networks. It is well known that these are capable of sequence recall as well as the recall of more static experiences.

Again, speaking introspectively, I am looking at a pencil on my desk and deciding that I want to pick it up. This thought is a sensation of my actually doing it in my head, before I do it for real. My depictive areas are producing a kind of film show in my head in anticipation of the real act. This comes from the fact that the depictive areas can learn appropriate depictive sequences as part of the build-up of experience as a sequence of depictive states. That is, as a child I learn to pick things up by trial and error. When I succeed reliably, my visual, tactile and muscular neurons have, together, learned to go from state to state by the same axiom 2 mechanism that allows them to remain stable in one state. There is very little technical difference between learning sequences and learning single stable states. But if there are many possibilities how are these controlled? What is it to want to execute one of the possible plans? This leads to axiom 5 and then to the main meat of this paper.

3.1.5 Axiom 5: Emotions

One of the criticisms levelled at those who speak of conscious machines is that this is one element of humanity that machines cannot have: feelings and emotions. We argue that as these seem to be essential to being a conscious human being they must be essential to being a conscious machine on account of their aid to survival. One should be suspicious of the consciousness of a machine were it not to have mechanisms that play the role of emotions in living organisms.

In the first instance emotions are related to the evaluation of depictive input. Children not more than a few hours old will show signs of fear (facial expression and a retreating action) if a large object moves towards them. The same occurs if the child is allowed to move freely over a glass surface that appears to stretch over a precipice. The child avoids the precipice and shows signs of fear. On the other hand the child shows contentment on being fed when hungry. So, basic emotions such as fear and pleasure, are neural activities that are pre-wired, through evolution, at birth. They have obvious survival value. Others in this innate group are anger, surprise, disgust and distress.

Other emotions and feelings are developed during perceptual life. Feeling hurt after being rebuked or being jealous of the attention someone else is getting are examples of a vast group of such subtle phenomena. On the basis that every scrap of our sensation is due to some neuro-chemical activity, the axiom suggests that such patterns have distinct characteristics that both adapt to be attached to perceptual depictive events as well as imagined events. As planning proceeds according to the mechanisms of axiom 4, predicted states of the world trigger emotional neural firing which determines which plans are preferred for execution and which might lead to unwanted consequences.

This is an area where a great deal of study still needs to be carried out and this paper is an example of such develoment. But one thing is sure, an organism without neural mechanisms for conscious emotional evaluation of thoughts and plans would have its capacity for survival strongly curtailed.

3.2 Comment on the axioms

The main reason for presenting the axioms as sequences going from a felt inner sensation to a generating mechanism that may both be found in the brain and act as a design principle for a conscious machine is to stress an important point. It shows that there need not be any insurmountable gaps in this sequence. If sensation implies mechanism then it at least seems feasible to assume mechanism implies sensation. To seek a science that separates sensation from the action of its material mechanism seems unnecessary.

The axioms should be seen as a necessary set. There is no claim here for sufficiency and the research community is invited to add to the set.

4. An axiomatic 'kernel' structure

Fig. 1 shows a minimal architecture implied by the axiomatic/depictive properties. The perceptual module directly depicts sensory input and can be influenced by bodily input such as pain and hunger. The memory module implements non-perceptual thought for planning and recall of experience. The memory and perceptual modules overlap in awareness as they are both contain locked cells. The emotion module evaluates the 'thoughts' in the memory module and the action module causes the best plan to reach the actions of the organism.

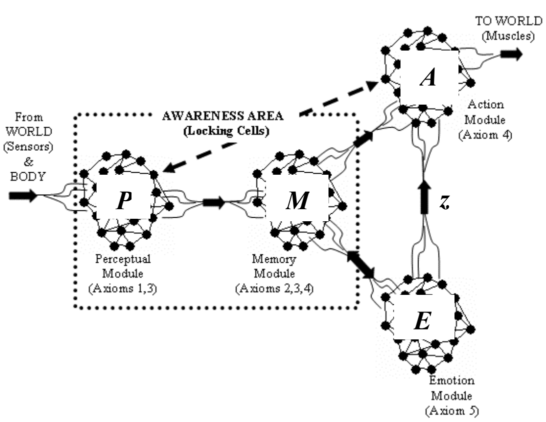

Figure 1. The axiomatic kernel architecture

Each module has its own set of states as follows:

Perceptual states:
$\boldsymbol{P}$: $\{p_1, p_2, p_3, \ldots\}$

The memory states:
$\boldsymbol{M}$: $\{m_1, m_2, m_3, \ldots\}$
Where each state is further subdivided into components:

m_j is a triple $(\alpha_j, \rho_j, \varepsilon_j)$ where
α_j is a remembered *action*
ρ_j is a remembered *result of the above action* and
ε_j is a remembered *emotion associated with the above result.*

Emotional states $\boldsymbol{E}$ are 'wantedness' evaluations that return an intensity ε_j ***(i)*** for state ε_j in $\boldsymbol{M}$.

Action states $\boldsymbol{A} = \{a_1, a_2, a_3, \ldots\}$ are vectors of *decided* muscular activity that may define the (largely) outward behaviour of the organism.

z is a very important signal: it is generated when the wantedness values for a particular imagined result exceed a threshold $\boldsymbol{t}$. This will be clarified in the example below.

5. Emotions and volition in the model – an example

To illustrate the way that this architecture works we imagine a simple scenario of having to make choices: looking at the menu of a pizza restaurant. The menu reveals only three items of food
Pizza; Pasta; Salad.

These, by a process of depiction (ax1) and attention (ax2) become, in turn (say), the states of the perceptual neural module, $\boldsymbol{P}$. These act as inputs to the memory/imagination module $\boldsymbol{M}$. Now, it is in this module that the 'thinking' goes on. In this case the thinking has to do with imagining the action that might be taken: eating pizza, pasta or salad. But this is not all one imagines. One recalls the taste of these dishes, and a whole lot of emotions associated with them. This and the operation of the emotion module $\boldsymbol{E}$, needs some explanation.

5. The emotion process

5.1 Wantedness generation

Emotions are taken to have the following character: first, they are remembered in the context of a predicted action and result of that action. For example, speaking introspectively, in imagining in $\boldsymbol{M}$ the action of eating pizza, I remember the result of this (the taste) and I also remember the associated emotions. That is, it is the presence of the action state, the taste state and the emotion states that make up the total state of $\boldsymbol{M}$. There could be several emotion states present at the same time: say, a gustatory pleasure, and also guilt because this is bad for my weight. That is, the predicted result of an action can have a collection of emotions associated with it, some positive and some negative. The second character of emotions is that they have a value, that is, an intensity, a strength with which an outcome is wanted. Third, under normal conditions, it is possible to resolve unclear combinations of emotions and make decisions in any case.

In the model we are here developing, the role of the ***E*** module is to recognise the imagined emotion and to evaluate it, that is create a value of 'wantedness'. This is not an instantaneous process and may involve attending to the available actions several times. The feeling is a familiar one: it takes a while to ponder the content of a menu and the actions and their results are visited several times. Whenever a choice of action is visited in ***M*****,** the total of all the related emotional values is summed up to give a 'how much wanted' value for each action.

So, in our emotional model, ***E*** generates a signal ***z*** when a sufficient level of wantedness is achieved. This is fed back to the imagination/memory module and holds that module in the action that has given rise to this high level. At the same time this signal is also sent to the action module enabling it to be set up to drive muscles as required by the action that is currently in ***M*** that the organism has decided to take. Of course this scheme allows for the ***E*** module never to reach an adequate 'wanting' intensity. To deal with this and conflicting or unclear decisions implies that in ***E*** there should be a random process which modulates that computed value of ***z***. The object of the random process is to provide an arbitrary enhancement of the 'wantedness' score in such a way that higher scores are more likely to trigger the action, but that lower score can also achieve this but with a lower likelihood. Even where there is no wantedness, just through the need to make an arbitrary decision (as we shall see in Libet's experiments) the random process can make the arbitrary decision of when to act.

5.2 The Random Process

Neural networks are very good at random processes, and it is often assumed that the neural structures of the brain are capable of this. For example, in our simulation of the example discussed here, we used a very small 10-neuron net which, due to lack of training, was simply producing binary states. The 'score' of the net was the count of neurons at 1 at any one time. This has normal distribution characteristics that are,

Probability of scores 0 and 10 = .097%,
1 and 9 = 0.97%, 8 and 2 = 4. 4% , 7 and 3 = 11.7%, 6 and 4 = 20.5%, 5 = 24.6%.

Now we can go back to the example and list the emotional value that might be in force at one particular moment that our organism is trying to decide what to eat. This can be set out as table 1 :

Table 1: Total wantedness values for three dishes

Menu item	Pleasure	Lack of Guilt	Total
Pizza	3	0	3
Pasta	2	0	2
Salad	1	1	2

Now say that to make a decision (generate signal ***z***) the ***E*** machinery has to equal or exceed a value of 7 (chosen arbitrarily for illustration) when adding the total emotional value to that of the random process net. So, the decision for pizza will be made when the ***M*** module is in the Pizza state and the random process generates 4 or more. This is obtained by adding the probabilities of random values from 4 to 10 all of which will cause the action. This sum turns out to be 82.8%. Similarly, when ***M*** is "thinking" of Pasta or Salad the random process must generate a value of at least 5. Summing the probabilities of generating 5 or greater indicates that the decision to eat the dish currently thought of will be taken in 62.3% cases.

A way of interpreting these results is to think of 100 (99 to be precise to have a number divisible by 3) people in this restaurant. They all have exactly the same feelings about the three dishes and each has the same random process. Each will be in one of the three food states, say about a third (33) in each. So of the group thinking 'pizza' (0.828x99/3) will make their decision to eat pizza that is, about 28 will make a firm decision to eat pizza. We note that 5 will make no decision at all, and return to consider the next item on the menu. Similarly, of the other two groups of 33 about 20 will choose pasta in the pasta group and 20 will choose salad in the salad group.

Of course these figures depend on the threshold of 7 that has been chosen to do the calculations. Had a higher threshold been chosen, the restaurant clients would be seen as being far less prepared to make decisions. So this threshold can be seen as a sort of 'mood' emotion. A hungry, relaxed mood (low threshold) will lead to a quick decision whereas, an anxious, picky mood (high threshold) will cause the organism to agonise longer before taking a decision.

The point of all this is to show that uncertainties and conflicting emotional values *could* be represented in machinery as shown in fig.1. We stress that it is not the case that this precise machinery is thought to exist in human brains. It is more an expression of the axiomatic stance which suggests that an architecture might at least be envisaged which appears to have the characteristics of emotional evaluations that go on in the head when one is trying to make a decision. But this needs a bit more discussion.

5.3 How does it feel?

The main tenet of the axiomatic approach is that only that which is depicted is experienced as a meaningful sensation. In figure 1 only the ***P*** and the ***M*** modules are depictive, ***P***, directly and ***M*** indirectly as a memory of the states in ***P***. Therefore it is quite true to say that we are fully conscious of what our options are and the nature of the emotion that comes with them. We have argued earlier that some such feelings are 'wired in' as, for example, fear, pain and pleasure are internally generated neural signals the reach the depictive areas in order to come into consciousness as 'visceral' sensations.

What is new here is that we have imagined a little further than in the bland statement of the axioms how the emotion module could use a random process to control the generation of action. This in situations where a direct reactive response (such as swiping my hand at a fly that has settled on my nose) is not possible due to there being several choices. While this process is not in the depictive part of the mechanism and one would not be conscious of it, the generation of *z* comes into consciousness as it holds the state of ***M*** for long enough to transfer the imagined action to the action module. This 'freezing' is what we would describe as the moment of consciousness that a particular action would be taken. This now puts us into a position where we can re-visit Libet's findings, and decide that perhaps will is not an illusion after all.

6. Lifting Libet's finger

In Libet's experimental setting, the decision to be made is not 'what' but 'when'. While there may be thoughts and emotions present about what one is meant to do, they do not have a direct bearing on the process. We suggest that the only thing that remains at work is the random process. Everything in the ***M*** module is set up to lift the finger, that is the intention is depicted, but the random 'wanting' machinery is on its own as there are no emotions to evaluate. In such a situation the threshold for generating *z* should be within the range of values produced by the random process on it's own. For the sake of an explanation, we assume that the same random process is at work and this means that the threshold 7 will be reached if the random process produces *z* with the sum of the probabilities of generating 7, 8, 9 and 10. Just out of interest this turns out to be 17.2%, that is, somewhat lower than the 'choice' decisions in which emotions are evaluated.

It is now possible to create a hypothesis that makes the unconscious generation of the Readiness Potential less mysterious and will less of an illusion. First, we submit that the generation of something like *z* in ***E*** *is* the readiness potential. But this is not the source of the willed action, just an emotion-like trigger that the action should take place. So it is hardly surprising that this trigger should be generated before the depiction of the action in ***M*** freezes, which is the moment at which the participant would look at the clock.

In other words the sequence of events goes like this: the desire to lift the finger at some point is fully depicted in ***M***, but without input from ***E***. This activates the random process which with some delay determined by the probability (i.e. not the same each time) of exceeding the threshold, generates *z* (the readiness potential) to which ***M*** reacts, say, half a second later and the action is transferred to the muscles a little while after that. But there seems little doubt to me that it is the initial desire depicted in ***M*** that is the cause of these events. So, it now becomes possible to summarise a perspective on the concept of will and all it entails.

7. Free will: a summary

It is possible to point at a folk theory of volition: we visualise something we desire – we act to get it. But our cultural inheritance from philosophy and religion and some recent neurological measurements would not leave it at that. Important questions have to be answered, and here we attempt to summarise what might have been gleaned through the depictive/axiomatic approach.

The first of the important questions is that of *freedom*. In what sense can the system in fig. 1 convey the notion of freedom? More pertinently, say that some elaborate form of this system were present in my brain, how is it that it makes me *feel* free to make my decisions in an unfettered way? We suggest that this freedom is felt at least at two levels. Back in the pizza restaurant scenario, (once again, using introspective language) I know, as I can know anything else, that in going to the restaurant, I shall be offered a choice of food from which I will be able to exercise my power of choice, with their emotional overtones and all. This is knowledge like any other: like that when I go into my study I know that will find a computer on the desk or like that when I go to Venice I will have to leave my car in a garage outside the watery city. All this is due to the natural mechanistic implications of axiom 2: areas such as ***M*** provide access to knowledge and experience, and my knowledge of restaurants tells me that I will be able to make a choice that suits me at the time, dependent on my moods and emotions.

Of course, someone could say that it is all predetermined. But it does not *feel* that way just because I am aware of having made different choices under similar conditions. So, predetermined or not, the feeling that what I will choose will be best for me at the time, is good enough not to feel constrained. What would *not* feel free would be the prediction from my prison cell that the same slop as always will arrive at midday.

The above is the first, higher, level of feeling the freedom of will. The second, lower, level, is the mechanism of evaluating emotional states for a series of attentional phases directed (say) at the restaurant menu. The cycling and eventual freezing of the state all occur in ***M***, which, according to the depictive axioms, is felt by the organism. Given language and choice of the most wanted item the organism would describe this as "I felt free to look at the choices offered on the menu and chose the one that appealed to me most". Or if a less wanted item was chosen this might be described as "I chose salad despite the fact that I don't like it all that much, but I know is good for me." In the finger lifting exercise most of us would admit that we didn't know what made us lift a finger at a particular time and that that moment seemed arbitrarily chosen.

The second question about what factors influence my choice is answered by the way we have suggested that emotional evaluations work. Emotions are recalled in ***M*** and the factors of 'wantedness' are computed in ***E***. Of course this process has not been properly elaborated here and is the subject of current research. Open questions relate to how the evaluations get developed through learning and how 'thresholds' develop and change with moods. Finally, the mechanisms described here clarify the involvement of axiom 4 and 5 as the basis of free will. Axiom 4 is the cycling in ***M*** and axiom 5 is the operation of the ***E*** machinery.

7. Will: a philosophical coda

While engineering arguments have been heavily employed in the last few paragraphs above, philosophers may not be happy with this. We conclude this chapter by setting out the logic of the argument in a series of assertions.

For an organism with a 'brain' to have a sensation of free will

1. There exist areas of the brain that support consciousness through having the depictive property (axioms 1 & 2)
2. There exists in the brain a depictive mechanism for cycling through the choices prompted by a perceived external event or an internal imagined event.
3. Cycling through the states in 2 includes memories of emotions associated with the choice states.
4. There exists in the brain a non-depictive evaluational mechanism (non-conscious therefore) that accumulates 'wantedness' values for the emotions associated with each choice. When wantedness exceeds some threshold the current choice state is translated into action.
5. As part of the evaluational mechanism, there exists a random process which adds to the wantedness values helping to resolve situations of conflict or lack of emotional value.
6. Through the depiction of the freezing of the cycling mechanism due to a wantedness trigger, the organism feels that actions are taken among choices according to how much something is wanted.

Conditions 1 and 2 are fundamental postulates which, if denied, block proceeding with the rest. Denying 3 requires a denial that emotions are involved in making choices. Such a denial would contradict common experience and documented material such as Damasio (1995). 4 and 5 are then the basis of the main hypothesis presented here, their denial or confirmation is a matter for both neurological and modelling research.

Acknowledgements

I. Aleksander wishes to thank the Leverhulme Trust for having provided an Emeritus Fellowship that makes involvement with this work possible. Rabinder Lee is supported by an EPSRC grant.

References

Benjamin Libet, and E.W.Wright, and D. K. Pearl,'Time of conscious intention to act in relation to onset of cerebral activity (readiness potential): The unconscious initiation of a free voluntary act', *Brain* , **106**, pp 623 – 42, 1983

Dan. M. Wegner. *The Illusion of Conscious Will*, Cambridge MA: MIT Press, 2002

Igor Aleksander, and Barry Dunmall: 'Axioms and Tests for the Presence of Minimal Consciousness in Agents' *Journal of Consciousness Studies.* **10**, pp 7-18, 2003

Antonio Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain,* London: Picador, 1995