

Paying attention to consciousness

John G. Taylor^{a,b,*}

^a *Department of Mathematics, King's College London, London WC2R 2LS, UK*

^b *Lobal Technologies, Huntsworth Mews, London, UK*

Received 28 April 2003; accepted 15 October 2003

Abstract

An engineering control approach is developed for the movement of attention, based on several features: experimental data indicating separate sites for attention modulation and for the creation of that modulation; the resulting analogy with motor control, to which an engineering approach has been applied; simulation and qualitative results supporting the presence of several of the necessary modules. These features are reviewed in the paper and a control model developed for the movement of attention. The engineering control framework is extended to the attended learning of motor control, again with description of support arising from simulations and qualitative analysis of several paradigms. The framework is even further extended to analyze how consciousness could arise during attentive processing, using the COrollary Discharge of Attention Movement (CODAM) model. This model is extended to encompass the temporal development of activity in various brain sites. Particular signals of the CODAM model are described and related to paradigms such as the attentional blink (AB) and features of simultaneous experience in neglect. A program of future explorations of the CODAM model and a set of open questions conclude the paper. © 2003 Elsevier Ltd. All rights reserved.

Contents

1. Introduction	306
2. The nature of attention	308
3. Motor control models in the brain	309
3.1. Engineering control theory	309
3.2. Application to motor control in the brain	311
3.3. Conclusions	312
4. A control model of attention movement	312
4.1. The basic control model	312
4.2. The components of the control model	313
4.3. Simulations of attention control	314
4.4. Further simulations	315
4.5. Relation to other models of attention	315
4.6. Conclusions	317
5. Control of motor responses through attention	317
5.1. A general model of motor attention	317
5.2. Simulating specific motor attention tasks	317
5.3. Extensions	319
5.4. Further work on motor attention	321
5.5. Conclusions	321
6. From attention to consciousness	321
6.1. Introducing consciousness	321
6.2. The creation of consciousness by CODAM	322
6.3. Ownership and agency: components of the self	323
6.4. Support for CODAM	325
6.5. The experience of self	325

* Tel.: +44-207-848-2214; fax: +44-207-848-2017.

E-mail addresses: john.g.taylor@kcl.ac.uk, jtaylor@lobaltech.com (J.G. Taylor).

6.6.	Limitations	326
6.7.	Conclusions.....	326
7.	Relation of CODAM to philosophies of mind	327
7.1.	Western cognitive science.....	327
7.2.	The other side of consciousness	327
7.3.	Conclusions.....	328
8.	Signatures for CODAM	329
8.1.	General signatures.....	329
8.2.	The temporal flow of activity in cortex: the Ni/Pi sequence	329
8.3.	The attentional blink (AB).....	329
8.4.	Conclusions.....	331
9.	Conclusions	331
9.1.	Summary.....	331
9.2.	Unanswered questions for CODAM.....	331
	Acknowledgements	332
	References	332

Nomenclature

AB	attentional blink, the phenomenon that occurs when a second target is attempting to be detected about 300 ms after a first in a rapid stream of targets and distracters.
ANG	angular gyrus, in the parietal lobes.
CD signal	the efference copy or corollary discharge of the signal for the movement of the focus of attention.
CTOA	cue-target stimulus time difference.
EEG	electroencephalography, the process of measuring the electric field around the head due to brain activity.
ERP	event-related potential, obtained by averaging EEG observations over many repetitions of an experience, so as to remove possible noise variations.
FBEL	feedback error learning, so using the error signal as proportional to the learning rate in any adaptive process in the brain.
IMC	inverse model controller for attention movement, generating a control signal for the movement of the focus of attention associated, with a given area of the brain.
MEG	magneto-encephalography, the process of measuring the magnetic field around the head.
MON	monitor module, involved with calculating errors between a desired and actual state of affairs, in the movement of attention.

PCE	pure conscious experience, as experienced in meditation, with no content in the experience.
PRS	The pre-reflective self or ipseity.
SMG	supramarginal gyrus, in the parietal lobes.
TMS	trans-cranial magnetic stimulation.
WMcd	working memory buffer for the efferent copy (corollary discharge) of a signal from the inverse model controller for movement of the focus of attention.
Wmsens	working memory buffer in a given sensory modality.

1. Introduction

There are presently numerous deep problems facing us when we try to understand consciousness. These include the following:

- Reconciling subjective experience, at the heart of consciousness, with objective knowledge, at the root of scientific methodology.
- Achieving interaction between the self, supposedly possessing a non-relational character, with the content of experience, as contained in so-called 'qualia', which do have relational structure.
- Explaining the many different states in which consciousness can be said to be present: in the normal awake state, in dreams, under drugs, in trance states, in meditation, in mental ill-health, such as schizophrenia and Alzheimer's disease.
- Choosing from the large range of models of consciousness that have been proposed with no consensus amongst them.

Because of this array of unsolved problems, as well as the vast range of prior predilections of those interested, confusion reigns in consciousness studies. This is evident from even a cursory perusal of the main journals in the field. Are there new developments that can help us move out of this morass? Promising avenues would be those concerned with new or newly appreciated knowledge of key mental states. Admittedly such knowledge is mainly of a qualitative sort and is not at all of the crisp, quantitative, objective variety sought by science. Even so, it is important to attempt to move our understanding of consciousness forward by bringing in such new knowledge. This includes the following:

- The ‘inner view’ of recent philosophers (Strawson, 1999) and of Western phenomenologists (Sartre, 2001; Henry, 1963; Merleau-Ponty, 1962; Zahavi, 1999; Parnas, 2000).
- The presence of special states arising in meditation (‘pure consciousness’ or ‘stillness’ (Forman, 1990) and now being studied by brain imaging methods (Herzog et al., 1990; Baerentsen et al., 2001).
- Updating of neuro-scientific studies of the modules crucially involved with deficits in conscious experience, as well as more details being obtained on the nature of the related deficits.
- Increasing of understanding of attention as a control system, and the various networks of brain regions involved in various attention-based tasks; it is especially important for consciousness that attention also be understood: many regard attention as necessary for consciousness (Mack and Rock, 1998; Rensink et al., 1997; Cotterill, 2001).

Since there is presently no quantitative theory of mental states, I will attempt, in this paper, to use the qualitative features noted under the above four points to create a tentative neural basis for the creation of aware states in the brain. These states will need to be shown to have similar properties to those experienced in consciousness, a vast project only able to be touched on briefly here.

It is necessary to have a solid information-processing framework for the brain before we are able to embark on such an ambitious project, even at a qualitative level. To obtain such a firm basis, I will start by considering attention, the fourth point noted above. This has been claimed to be essential for consciousness, as I noted earlier. Thus, Mack and Rock conclude (Mack and Rock, 1998, p. 74): “so, in the end, we have been driven by our data to conclude that attention is essential for conscious perception”. In spite of this position being argued against by some (Lamme, 2003), there is now such strong evidence in its favor that I will use it as the basis to make my assault on the citadel of consciousness. I will later turn to consider more closely the relation between attention and consciousness. In summary, I will use attention as a gateway to consciousness.

To make progress, I need to build a suitably specific theory of attention. Only a detailed enough theory will be able to contain within its structure those components that could be recognizable as that specific neural activity creating con-

scious experience. The tactic of this paper is therefore initially to concentrate on attention and only afterwards explore how consciousness might arise from inside the resulting neural model of attention. We will be able to use the qualitative features of consciousness, mentioned under the first three points above, to help guide what to look for in our model of attention.

Such a step, from attention to consciousness, will not, however, be plain sailing. Attention allows for a reasonable translation between brain activity associated with cognition and the associated mental states, as I have discussed earlier (Taylor, 1999, 2001a,b). However, the resulting dictionary had a crucial missing entry: that of the subjective self. How to incorporate that is the ‘hard problem’ of consciousness (Chalmers, 1996) and also discussed perceptively more recently (Noe and Thompson, 2003). Thus, before we can hope to be able to tackle the hard problem, we must develop a suitably detailed and powerful neural model of attention. Only then can we hope to be able to make progress.

Much experimental work has recently been performed on attention. Detailed experiments have measured brain activity while attention was being used to solve specific tasks. Such measurements have been made at both single cell and, at a more global brain level, by brain imaging machines. The resulting experiments have shown, in particular, that many brain areas are involved in attention-based processing. As part of constructing a model of attention, we have to face the difficult task of understanding the nature of the involvement of these different areas across a range of paradigms.

Recently, I introduced an engineering control framework to explain the movement of the focus of attention (Taylor, 2000a,b,c, 2001a,b, 2002a,b). Such an approach has support from the nature of attention, properly regarded as an action of certain parts of the brain on other parts. This is similar to motor control, which also involves the action of parts of the brain on other parts and even more directly on parts of the body. In the motor case, there are prefrontal and parietal sites that are observed to exercise influence on motor cortex and sub-cortical sites, these all being involved in movement planning and learning (Willingham, 1998). Considerable success has been reached in applying control ideas to motor actions (Sabes, 2000; Desmurget and Grafton, 2000; Miall and Wolpert, 1996; Wolpert and Gharharmini, 2000; Kawato, 1999). Moreover, discovery of detailed sites of various components of the motor control components have been claimed by a number of groups (Imamizu et al., 2000; Schweighofer et al., 1998; Desmurget and Grafton, 2000).

Such a control approach to attention is accepted by many who are investigating its neuro-scientific attributes experimentally, either by brain imaging, by single cell, or by deficit studies. It was already noted in an earlier review of attention (Coull, 1998, p. 343), “a distinction is drawn between the brain areas which are crucially involved in the top-down modulation of attention (the ‘sources’ of attention) and those sensory-association areas whose activity is modulated by attention (the ‘site’ of attention)”. The

approach presented here is based on this support. It tries to organize the results, from the numerous attention studies, under the aegis of engineering control. This latter approach has great flexibility, at the same time allowing specific functions to be searched for. Such a framework also leads to predictions, as will be described later. In addition, it is possible to work by analogy with the successful use of such a framework in motor control. In total, the engineering control approach to the movement of attention has allowed tentative assignments to be made for some of the brain areas observed active (Taylor and Rogers, 2002, Taylor and Fragopanagos, 2003). Such assignments have also to be tested by developing simulations of paradigms as well as qualitative analysis of a range of further paradigms. I will report here on the present progress on this program of understanding attention through engineering control theory.

On the basis of the resulting control model for attention in the brain, I will extend the analysis of the model to explore its capability to create conscious experience. I have already reported on preliminary results of that approach to consciousness elsewhere (Taylor, 2000c, 2002a,b,c). The resulting model was termed Corollary Discharge of Attention Movement (CODAM) (Taylor, 2000c). It is such corollary discharge that I proposed as the basis of the experience of the pre-reflective self. I will analyze CODAM further later in the paper so as to determine what experimental signatures could arise from it. These signatures will be important for future experiments to test the model as well as to be related to present experimental evidence.

The paper starts with a brief survey of the nature of attention. I then briefly review engineering control, and more specific motor control concepts, in Section 3. A general control model for sensory attention is developed in Section 4, including simple simulations which support the presence of certain components—the inverse model controller (IMC) and goal modules. Other components—the monitor and forward or observer models—are discussed in terms of further paradigms. In Section 5, I extend this control model to one for joint sensory and motor attention. I again turn to simulations of simple paradigms and discuss further paradigms qualitatively. In the following section, I present and discuss the CODAM model suggested for the creation of consciousness. Some philosophical aspects are discussed in Section 7. Experimental signatures to observe CODAM by brain imaging are discussed in Section 8. Conclusions and open questions to be further explored are considered in the last section.

2. The nature of attention

Attention is most often defined, for example in the visual modality, as the process of selection of a part of a visual scene for further more careful inspection; it involves the partial or complete exclusion of the rest of the scene. Unattended inputs fail to reach awareness, as shown by experiments in which a subject is completely unaware of

the sudden appearance of an object at an unattended point (Mack and Rock, 1998; Rensink et al., 1997). In a very similar way, attention can also be paid to inputs in other modalities or to actions.

A number of the features of attention are not new; the study of attention has a long history. It was already noted by Aristotle that ‘Of two movements, the stronger always tends to extrude the latter’. That attention can be directed was also known in antiquity. Thus, Lucretius wrote, in the first century B.C., that one can possess ‘attentive ears and minds’. The division of attention into that directed voluntarily or externally was also known early. The automatic ‘tug of attention’, as compared to its willed direction, was noted by St. Augustine in about 400 A.D. and discussed later by Rene Descartes. Many remarked that attention enhances sensory sensitivity (Lucretius, Descartes and Bonnet), and the ability to move attention covertly was also commented on by Aristotle as well as being studied later by Bonnet. Again Aristotle pointed out that attention was basic to unifying consciousness. More recently William James was much concerned with attention. However, a hiatus arose in analysis of inner experience brought about by the failure of introspection studies and later by behaviorism. Only in the latter part of the previous century has attention come back, not only to being recognized as crucial to consciousness but also in its own right deserving of more detailed study as a crucial component in information processing in the brain. Its introduction into artificial software ‘agents’, for example, is now being seriously attempted (Taylor and Kasperides, 2003).

The mode of action of attention in the brain is now realized as through relative amplification of the neural activity for the attended input and the concomitant inhibition of distracting elements in the visual (or other sensory) neural representations in early cortical sites. Such modulations have been observed at single cell level (Reynolds et al., 1999; McAdams and Maunsell, 1999), at multi-unit level (Mehta et al., 2000), and at multi-modular level by brain imaging (Buechel and Friston, 1997). They occur in various parts of the cortex, both in posterior visual sites and in parietal and frontal sites.

For example, Reynolds et al. (1999) tested the effect of adding a second stimulus to an earlier one being observed by a macaque. They showed that the effect of directing attention to one of a pair of stimuli on single cells in V2 and V4 drove the response to that elicited by the attended stimulus alone. They concluded that, “this finding is consistent with the idea that attention biases competitive interactions among neurons, causing them to respond primarily to the attended stimulus”. Buechel and Friston used structural equation modeling to analyze fMRI time series from three subjects paying different levels of attention to visual motion. They showed that the data could be best explained by the existence of a modulation of the connection between V5 (MT) and the posterior parietal cortex by attention. They further demonstrated, by an extended structural model, that this modulation arises from the prefrontal cortex.

Numerous experiments in brain imaging have demonstrated that there is a network of cortical modules involved in controlling attention (Hopfinger et al., 2000, 2001). This network has been noted to be very similar for attention controlled by external rapidly appearing inputs (so-called exogenous attention) and the top-down or endogenous form controlled by internal goals (Giteleman et al., 1999; Kim et al., 2000; Coull and Nobre, 1998). The most important cortical components of this attention network have been proposed as being at three sites: the parietal lobes for perceptual processing, the anterior cingulate for limbic-based intention, and the prefrontal cortex for rules and templates for guiding response (Mesulam, 1981). A recent review (Kastner and Ungerleider, 2000) concluded that “attention-related activity in frontal and parietal areas does not reflect attention modulation of visually evoked response, rather it reflects the attention operations themselves”. Also (Corbetta and Shulman, 2002, p. 201), in their even more recent wide-ranging review of attention, concluded: “one system, which includes part of the intra-parietal cortex and superior frontal cortex, is involved in preparing and applying goal-directed (top-down) selection for stimuli and responses”.

We may summarize this control feature of attention by a two-module model, as in Fig. 1.

The right-most module in Fig. 1 denotes the regions of the brain that are controlled by attention. They consist of lower regions—especially in occipital, temporal and motor cortex—in which input or response activities are modulated by attention, but these areas do not function so as to control attention in any strong manner (barring areas, such as the amygdala, which code for input saliency, so affecting where attention is directed). The controlled areas are, thus, recipients of attention control signals. The other set of areas are those under the title of controller in Fig. 1; these areas create the attention control signal itself. They are composed of ‘higher’ areas, especially the prefrontal and parietal areas of cortex.

Given that attention functions by signals from a separate control region being sent to modulate the controlled regions, the creation of such modulation itself is expected to involve internal complexity in the relevant neural control regions. To explore this complexity, we note that the variety




Fig. 1. Summary of the experimental results on attention: the brain is divided into two types of areas, the controlled and the controller. The first set of areas, as shown in Fig. 1, are under attention modulation from the sites acting as the controller. The former sites consist of sensory areas across all modalities and motor areas, as well as semantic-level object map areas. The latter, controlling sites, consist mainly of prefrontal and parietal areas.

of distinct functional components in engineering control models (goal sites, forward models or observers, inverse model controllers (IMCs), error monitors, all of which will be discussed in the next section) leads one to expect a similar range of functions being performed by the component brain areas observed in attention. The separate network of cortical modules for overall attention control (prefrontal, parietal), as well as the earlier cortical sites acting as recipients of the control signals needed to achieve the filtering process basic to attention from the higher-level sites, are ripe for analysis by engineering control concepts.

In parallel with deepening of our understanding of attention, the increased structure being brought to bear by engineering control ideas could also contain the missing functional elements providing an explanation of consciousness. This will help guide us to determine what the extra functionality is to move beyond attention and explain why it is only necessary and not sufficient for consciousness.

In conclusion, attention involves brain sites with the following features:

- They amplify or decrease sensory input (in sensory cortices).
- They create control signals for this amplification/inhibition (in parietal and prefrontal cortices).
- Other possible functions could also be performed by brain modules involved in attention, such as error detection, predicting feedback, etc.
- The division of labor involved in attention control, and the large set of brain areas involved in various attention tasks, as noted above, implies that some form of control approach could be helpful to help understand how the brain achieves the overall process of attention.

The use of engineering control theory to understand attention will, thus, be developed next.

3. Motor control models in the brain

To construct a control approach to attention it is useful to start with a brief description of the similar approach to the brain's control of motor actions. There is increasing support for the presence in the brain of such engineering-style control models to achieve effective actions. I will start by giving a brief review of a general control approach before describing how it is being applied to motor control in the brain. I will turn to application to attention in the following section.

3.1. Engineering control theory

The task of engineering control is the construction, by a variety of methods, of a system able to provide a control signal to a given plant so that the plant response is kept within certain specified limits during its operation (Jacobs, 1993; Franklin et al., 1989). Controlled parameters might, for example, be the concentration of a given substance or the temperature of operation of a given preparation process

(such as in a steel mill). Various ‘observables’ are assumed known, such as the temperature of a given part of the plant (say, the surface of the steel). These observables are used to determine how to alter the values of the plant controls, such as the amount of heat supplied to the plant, in order to keep the whole plant at or within desired temperature, or other variable, limits.

To begin with, we assume that the state of the plant is describable by a vector denoted $state(t)$, at a given time t . The components of the state vector will be the values of temperatures, concentrations, etc. of various crucial components of the plant, as measured at time t . This is not a strong assumption since, if velocities or accelerations of various quantities are also involved in describing the state of the plant, these can be added to the total vector so as to extend it to a complete state description. Nor does it assume that the state cannot be given by descriptors of the plant at different times, such as at times t and $t + 1$; that also can be handled by adding all the necessary components together, as in time series analysis, but I will not go into those details further here.


The basic module to be constructed is the IMC. This sends a control signal, denoted $control(t)$, to the plant to modify its operation. The size of the control signal will depend on the actual state of the plant, say $state(t)$ at time t , and the desired state to be attained by the control signal, $state_d(t)$. The larger the difference between the desired and actual state, the larger must be the control signal to change the actual into the desired state. We have the general picture of the IMC shown in Fig. 2a.

A further module of importance in engineering control is the goal module. This has already been used implicitly in the IMC of Fig. 2a. The goal module contains the desired state either as a given constant state or more generally as a state trajectory to be followed by the plant. Such time variation would be the case, for example, for a car-driving controller where the desired trajectory is along the right side of the road (or left side, depending on the traffic rules) following correctly the curves of the road, stopping at traffic signals, etc.


A further important control-engineering component is a forward model or so-called observer. This gives an estimate of the state of the plant. The forward model produces, as output, an estimate of the plant state at the next time step, $state'(t + 1)$, given the state at time t and the associated control signal $control(t)$. This estimated or predicted state is of crucial importance in the creation of fast control in the case that feedback is delayed in response to changes in the plant (as occurs, for instance in motor control with respect to feedback from a limb). A forward model is shown in Fig. 2b.

Another module of relevance is that used to give an estimate of what feedback is to be expected from the plant after a given control signal has been applied to it. It is assumed that observations can be made on the plant, although these may be of limited form. The feedback can be used to update the state estimate from the forward model so as to keep it as close to reality as possible. The forward output module is shown in Fig. 2c.

a. Inverse Model Controller



b. Forward Model/ Observer



c. Forward Output Model



d. Error/Monitor Module




Fig. 2. Components of engineering control systems: (a) the inverse model controller (IMC), produces a control signal so as to achieve a desired state given an initial state for the plant; (b) the forward model or observer, provides an estimate of the next state of the plant, given its present state and the control signal being applied to it; (c) forward output model produces an estimated sensory feedback from a given state, and can be used for fast ‘feedback’ error correction; (d) error or monitor module, which produces an error signal, determined by the difference between the desired state and the actual state. This latter can be obtained either from the estimated state or that arising from feedback to update the state.

Finally, there is the error or monitor module, valuable in training the various control modules as well as being crucial for correcting errors that may have arisen in the plant control signal from the IMC. The output of the monitor module is calculated as the difference between the desired and the actual, or between the desired and the estimated, state values of the plant; the Monitor output is the relevant error signal. This is shown in Fig. 2d. Other errors (such as between the state estimate made by the forward model and that arising from use of observation) are also of value to correct and train the overall system.

We next consider a typical engineering control model, as in Fig. 3, incorporating the various components of Fig. 2, which act in combination. There is feedback from the plant through the observation module, control on the plant by the IMC module, with fast response from the Forward model to generate an expected error in the monitor module, in respect to the desired result from the goal module. Both the error signal and a direct signal from the goal module are shown in Fig. 3 as being used by the IMC to generate the control signal.




Fig. 3. An overall engineering control model of a plant, with the various components of Fig. 2 in place. The plant feedback is measured by an observer, and used to update the forward model after its arrival there. Before that the corollary discharge signal, sent as an efference copy of the control signal from the IMC, is used to give a more rapid update of the state estimate in the observer or forward model, so as to allow for more rapid error control through the monitor.

Several points to note on the control model of Fig. 3 are as follows:

- The observation process is made by a specific module, denoted as the observer module. This provides either a partial or complete description of the state of the plant.
- Any delay in observation feedback can be overcome by a fast ‘efference copy’ of the attention movement control signal being fed from the IMC into the forward model. The updated estimate of the plant state can then be used to correct the IMC response if it is in error.
- The IMC functions by using both a direct goal signal as well as an error signal (from the monitor module); the relative contribution of these two can be determined in a variety of ways, depending on the needs of the system.
- A feed back error learning (FBEL) signal from the monitor, proportional to the error, and shown as a dashed arrowed line in Fig. 3, can be used to train both the IMC and the goals and forward modules.
- It is expected that the IMC is also controlled by feedback from the plant; this is shown as arising from sensory feedback, in the observer module, and is used in the forward model. This feedback can be used in a trivial manner as direct feedback without being combined in any optimal manner with the forward model. On the other hand, it can be incorporated more efficiently with the forward model estimation. In some cases, a forward model may not be present at all, when such a fusion problem is absent.
- If Observer feedback has little delay in reaching the IMC, then a forward model will not be needed to speed up the IMC signal.

The subject of control modeling is far more extensive than the above brief introduction indicates (Jacobs, 1993; Franklin et al., 1989), but the above modules will be the main ones to be used in the following discussion of attention control in the brain.

3.2. Application to motor control in the brain

In motor control, there may be a considerable time delay before feedback comes available to the IMC; this delay may be upwards of several hundred milliseconds (Sabes, 2000;

Desmurget and Grafton, 2000; Kawato, 1999; Wolpert and Gharharmini, 2000). The basic idea is that for rapid and effective control a forward model of the state of a given plant (such as a limb) is needed so that actions can be made on the plant without having to wait for plant feedback to correct movement errors (Miall and Wolpert, 1996; Kawato, 1999). Such a forward model maps the current state estimate, and the action just taken, onto an expected new state, as already described in engineering control terms (see Fig. 2b). An inverse model for producing the action, based on the newly estimated internal state, is also needed. This uses the estimated present state, together with the desired next state, to calculate the requisite action. Updating the forward model is necessary both in real time by error terms arising from so-called corollary discharge (arising from an efferent copy of the control signal) compared to the feedback from the periphery as well as by longer term modification of the parameters of the forward model (Shadmehr and Brashers-Krug, 1997; Imamizu et al., 2000).

It is relevant to note that these more sophisticated control approaches, using forward models, are related to the notion of corollary discharge and the ‘re-afference principle’, in which self-motion effects are removed from afferent signals by subtraction of the corollary discharge, thereby removing the expected re-afferent component (von Holst and Mittelstaedt, 1950; Sperry, 1950). The image of a moving external world, as arising from self-motion, is thereby prevented, and stability of perception is attained. The discovery of such a process in living systems has made the re-afference principle very attractive as a control component; Sperry, for example, permanently rotated one eye of a frog and found it could no longer compensate for self-motion but kept moving round and round to steady its external world. More recently, inability to tickle oneself has been recognized as involving suitable re-afference canceling otherwise ‘foreign’ signals (Blakemore et al., 1998).

The general features of the relevant control model have already been shown in Fig. 3. This can be related to many further discussions on motor control in the brain, beyond that of simple re-afference, such as Sabes (2000), Desmurget and Grafton (2000), Schweighofer et al. (1998) and Wolpert and Kawato (1997). From Fig. 3, an IMC sends its movement

signal to the ‘plant’ (such as the arm), as well as to a forward model for modifying the estimate of the state of the arm by an error signal arising from the monitor. This error signal is based on a comparison between the desired arm state, from the goal module, and that predicted after the new action (from the forward model, updated by the corollary discharge) or actual arm response (by the sensory feedback signal).

The presence of such a control model for motor action in the brain is supported by tests of a variety of subjects on a range of controlled movements (Sirigu et al., 1996). In particular, the subjects were required to make movements of a hand-held pointer to a target of variable size. The time taken to make the movement was found to increase logarithmically with decreasing size of the target following Fitts’ law, as was their time to make similar imagined movements. The detailed nature of the slowing, observed in the case of subjects with motor cortex or parietal lobe deficit, leads to the conclusion that the forward model involves motor cortex while the monitor involves the parietal lobe (Taylor, 2002c). Other sites are also expected to be involved, especially the cerebellum, suggested as holding adaptive values of the parameters of the controller and estimator (Shadmehr and Basher-Krug, 1997; Imamizu et al., 2000; Schweighofer et al., 1998).

Recent results strongly support the existence of a forward model in motor control (Sabes, 2000; Desmurget and Grafton, 2000; Wolpert and Gharharmini, 2000; Kawato, 1999). The experiments analyzed eye movements related to the unseen movement of a subject’s hand towards a given target. It was found that each saccade was to a spot that the hand would move to about 200 ms later. This movement was in terms of several saccades during the smooth hand movement. Perturbation of the hand was followed by a brief inhibition of saccades for about 100 ms with a following saccade to regain the predictive capability of the eye about 100 ms later.

That there is a frontal component of the forward model is to be expected on general grounds from the well-known anatomical recurrence: cortex → basal ganglia → thalamus → cortex as well as the more direct recurrence cortex ↔ thalamus. Such recurrence has been postulated as supporting temporal sequence storage and generation, which is at the basis of neural architectures able to perform the temporal updating required to solve the differential or difference equations of a forward model plant estimator. The parietal involvement in the forward model is supported by many results on working memory activations observed there (Paulesu et al., 1993; Honey et al., 2000, 2002).

3.3. Conclusions

Modern engineering control has proved valuable in many areas of industrial application. It is now proving itself in application to understanding motor control in the brain. This value implies that it deserves to be exploited further in control problems involving living systems. Attention presents such a control problem, so we turn to consider how to use engineering control in that context.

4. A control model of attention movement

4.1. The basic control model

The control model of Fig. 3 is now applied to attention. Past models of attention have been based mainly on posterior sites of amplification/reduction of inputs (Koch and Ullman, 1985; Olshausen et al., 1993; Itti and Koch, 1999; Mozer and Sitton, 1999; Wright and Ward, 1998; Deco, 1999; LaBerge, 1995; Corchs and Deco, 2001), although some models have taken account of frontal control by using templates held in frontal working memory sites for testing input correctness in target search paradigms (Jackson et al., 1994). All of these earlier approaches use a general control approach; however, it is varied to suit the task. We here take the control model of Fig. 3 more seriously and incorporate endogenous effects more fully by explicit treatment of the goal module. At the same time, we also take specific account of the other modules in Fig. 3, especially the monitor and forward modules; these have not been introduced previously into models of attention.

To summarize our approach, by its generality, the engineering control approach to attention gives a unified but flexible framework from which to take maximal account of the underlying control nature of attention as observed by psychological analysis and brain imaging experiments. It includes the majority of previous approaches to modeling attention. To support the engineering control approach, we must now search for evidence in the brain for the various components of the model, both in the endogenous and exogenous cases. Later, we will turn to more detailed architecture, with simulation results, to support our argument.

A general approach to the search for the requisite sites employed in motor control has earlier been delineated by Schweighofer et al. (1998); we will use a similar set of criteria. These criteria were based either on

1. The desired character of the inputs, internal activities, or outputs of the various modules of the overall control model.
2. The effects on performance of degradation of one or other of the modules of the control system.
3. The existence of suitable adaptive character for the connections to a given module, especially by means of feedback error learning (FBEL).
4. The appropriate topographic character of inputs to a given component module to enable it to function in the desired manner (so as to enable control of nearby items, such as joints).

We have already used these criteria, in essence, in presenting the results on the existence and siting of motor control modules in the previous section. Here, we will apply them more explicitly to guide our analysis of various sites as candidates for the components of Fig. 3 but now in control of the movement of attention; we treat this in the context of visual information processing as this has been far

the best-studied modality, but as will be later emphasized, this does not imply a limitation of the model in this context.

4.2. The components of the control model

Considering each of the components of Fig. 3 separately, we arrive at the following conclusions.

- **Plant:** this consists of posterior cortical sites on which input processing at feature level has occurred, and in which selected activity is to be amplified, the remaining activity inhibited, for later higher-level processing. The brain sites involve various feature maps, such as for color and shape in V4 or for motion sensitivity in V5 (MT). In general, these modules are in primary and unimodal associative visual cortex. We also include, in the plant, modules for object representations, sited in the temporal lobes. Both of these assignments are strongly supported by brain imaging, and single cell and optical dye recording. Evidence for amplification/reduction of attended/unattended inputs has already been cited (Reynolds et al., 1999; Treue and Maunsell, 1999; Mehta et al., 2000).
- **IMC:** this is required to produce the desired relative amplification/inhibition of the input signal for higher processing. There is evidence, from brain imaging during covert movements of attention, that control signals are produced in the lateral intra-parietal area LIP (Corbetta, 1998; Corbetta and Shulman, 2002) although the exact siting is still controversial. It is agreed, however, that the IMC is in the parietal lobes since both these and many other brain imaging studies (Coull and Nobre, 1998), as well as neglect studies in humans (Vuilleumeir and Rafal, 2000), support this. There is evidence that spatial attention movement may be controlled by SPL whilst object and feature attention movement is controlled by the Precuneus. This split may arise from the dorsal/ventral split (Vidyasagar, 1999), as involved in a division into spatial/object and feature space differentiation, with dorsal cortical sites being activated more rapidly by the M-stream of neural activity (Hopf et al., 2000), and thereby giving a spatial bias to the control of attention movement.
- **Monitor:** this is an important component of the overall control system of Fig. 1 but not separately included in previous models of attention. It provides a measure of the error occurring during an attention movement and obtained by comparing the actual to the desired attention target. Evidence has been cited earlier for the monitor involved with motor action control being in the parietal lobes and associated thalamic nuclei. The same is also expected to be true for that involved with the movement of attention, from evidence already quoted as well as from the confluence of inputs there from many bodily sites (and needed for effective determination of errors in attention as well as action control). There is also evidence for monitor signals involved with saccades occurring in the thalamus (and in superior colliculus) where visually responsive cells are found to show no response to self-motion (Robinson and Petersen, 1992). As the authors note, “responses of pulvinar visual neurons exhibit a filtering process that distinguishes the movement of retinal images of the external visual environment from a similar movement produced by eye movements themselves”. This effect can, thus, be explained as arising from the cancellation of corollary discharge signals of intended saccades from the afferent input to the thalamus—in other words, the re-afference principle in action. Pulvinar and superior colliculus are well connected to the parietal lobe, so the involvement of these further sites in the monitor is, thus, supported. Further differentiation between the pulvinar, thalamus and superior colliculus is not clear, so they will be lumped together in the monitor.
- There may also be error signals from the cingulate. A number of experiments have shown that the cingulate is involved in conflict resolution and error correction (Botvinik et al., 2001). A separation of these two functions has been suggested between the caudal cingulate (for conflict resolution) and rostral cingulate (for error correction) (van Veen and Carter, 2002).
- **Forward model:** this updates the state estimate of the attention state of the plant. Visual attention is known to be involved in eye-movement preparation, as indicated by the ‘premotor’ theory of attention, for which there is now considerable support (Rizzolatti, 1983; Rizzolatti et al., 1994). Updating of receptive fields, even before a saccade occurs, has been shown to occur in FEF, SEF and superior colliculus as well as in the parietal lobes (see the next point on the CD signal). Even stronger evidence has recently been obtained for the existence of a forward model in saccade direction (Ariff et al., 2002). In this, saccades are shown to precede unseen arm movements by about 200 ms.
- Thus, there is support for the existence of a forward model, possibly sited in the FEF/SEF (along with sub-cortical components), together with possible contributions from the parietal and superior colliculus. Siting of the monitor as partly in the superior colliculus and the parietal leaves the forward model expected to have some frontal basis in the FEF/SEF, as supported by the underlying recurrent architecture needed for dynamical sequence learning and generation.
- **Corollary discharge (CD) (or efference copy):** this is a copy of the attention movement signal from the IMC to be used in updating the error on the monitor and in allowing modification of the state estimate by the forward model. The existence of the CD signal is crucial to our CODAM approach to consciousness (as developed in Section 6). The presence of corollary discharge in eye movement and many action processes is well recorded (Miall and Wolpert, 1996; Sabes, 2000; Desmurget and Grafton, 2000). It also plays an important role in self-actions, such as not allowing one to tickle oneself (Blakemore et al., 1998). The corresponding situation is not so clear for

attention. However, the tight coupling between attention movement and saccade preparation noted above is relevant to help understand, in terms of the presence of attention movement corollary discharge, recordings from FEF and LIP cells that possess pre-saccadic responses with latencies of up to 70 ms before a saccade (Umeno and Goldberg, 1997; Duhamel et al., 1992). Single cell recording in FEF from those references implies, from the predictive response, the presence of a corollary discharge of the intention to make a saccade. Comparison of the response to a stimulus, without a saccade, with the pre-saccadic response, shows that the latter has an additive contribution compared to the former, which begins before the saccade, rises to a maximum after about 150 ms, and then dies away by about 300 ms. Attention and saccades are coupled strongly together, as noted above, which implies that this pre-saccadic activity could be generated by the corollary discharge of the IMC for attention movement. This supports the existence of the corollary discharge for attention movement in various sites (including the LIP and superior colliculus). It also supports the proposal that the FEF is involved in updating the posterior attention state estimator noted above. We will discuss further evidence for existence of the CD signal in Section 8.

- Goals: this module holds, possibly for up to many seconds, the neural activity representing movements required to be made on receipt of a given GO signal, such as to move attention to a particular place when a fixation light is extinguished. The goal module is sited in the dorso-lateral prefrontal cortex (DLPFC) since the second of a pair of sequential saccades cannot be made accurately if there is a DLPFC deficit (Perriot-Desilligny et al., 1995). Numerous other results, such as in the Stroop paradigm, support this site. We do not here consider decisions; the latter are now observed as served by the ventral network, providing limbic-based emotional salience to bias which rules to be used (Corbetta and Shulman, 2002).

The overall assignments outlined above are shown in Table 1, with the sites of the attention control system being given in Fig. 4 as the overall system.

Table 1
Suggested sitings for the modules of the visual attention model of Fig. 4

Modules of Fig. 4	Brain siting
Plant	Associative visual cortex and temporal lobes
Inverse model controller (IMC)	LIP/SPL
Working memory buffer (WM)	IPL
Monitor	Parietal cortex/pulvinar/SC/ACG
Goals	DLPFC
Observer/forward model	FEF/SEF/TPJ

LIP: lateral intraparietal cortex; SPL: superior parietal lobe; IPL: inferior parietal lobe; DLPFC: dorsolateral prefrontal cortex; FEF: frontal eye fields; SEF: supplementary eye fields; SC: superior colliculus; ACG: anterior cingulate cortex; TPJ: temporo-parietal junction.

Points to be noted in association with Table 1 and Fig. 4 are as follows:

- The model can apply to other sensory modalities, such as audition, somato-sensation, olfaction.
- There will be competition/combination between different sensory modalities in the overall model with a multi-modal parietal region very likely devoted to handling the overall competitive process needed to correctly share attention (Downar et al., 2000).
- There may exist separate control of attention to object features (color, shape, etc) and to space although there seems strong evidence that spatial attention is primary and used to guide temporally later feature/object based attention control, as we have already noted (Hillyard and Anillo-Vento, 1998).
- There must be extended learning (especially in childhood), possibly through the FBEL method using the Monitor signal, to lead to effective IMC input modulation.
- Sub-cortical sites will be involved in input representations, and in the setting up of goals, as well as in the Monitor & Forward models.
- Sub-cortical activation, through acetylcholine (from the Nucleus Basalis Meynert), together with other neuromodulators, such as noradrenaline from locus coeruleus, is also to be expected to be crucial in the spread of attention control (Doya et al., 2002).
- There will also be use of dopamine in learning (in basal ganglia/pre-frontal cortex/amygdala/hippocampus), as is well documented.

The model of Fig. 4 and Table 1 is extendable to modalities other than vision, as noted earlier, although there is not as much experimental evidence. There is also the question of how executive control is achieved for attention across different modalities. It has been suggested by a number of researchers that parietal lobe is able to support an extended competitive interaction (as noted above). Thus, overall control may be fought out mainly in the parietal lobes. Further experiments are needed to demonstrate that conclusively.

4.3. Simulations of attention control

Simulations have been performed to demonstrate that data from various attention paradigms can be explained by the control architecture of Fig. 4. In particular, paradigms concerning the effects of the top-down modulatory attention signals to lower sites (Reynolds et al., 1999) and of the dependence of the Posner benefit (Posner, 1980) on cue-target asynchrony have been carried out (Taylor and Rogers, 2002). In particular, the Posner benefit shows the value of using a control approach to attention.

The Posner benefit involves measuring the benefit, in speeded reaction time, obtained by directing attention to a target input, as compared to when the input is not being attended to. Thus, a subject is asked to fixate a central cross. Their attention is then directed to one or other side of the




Fig. 4. Application of the control model of Fig. 2 to the movement of visual attention: The plant of Fig. 2 is identified with the visual cortex and relevant object representation. Much feature analysis is assumed occurring in the visual CX, not being shown explicitly. Similar structures are expected to be present in other modalities.

viewing screen, to an outlined target box, by an exogenous signal (sudden brightening of the target box) or by a central arrow cue directing attention to be moved in the direction of the arrow. A target then appears some hundred or so milliseconds later in either the cued target box (valid case) or in the un-cued box (invalid case). The differences of the reaction times of the subject to the invalid target (say recognition as to the target being an \times rather than a $+$) as compared to the valid target are denoted as the attention benefit. This has been studied under a variety of situations and in its dependence on the cue-target time difference (CTOA).

A simulation was constructed and reported in Taylor and Rogers (2002) of the dependence of the attention benefit on the CTOA. It was a very simple simulation, able to take account of both the exogenous and endogenous cue cases. The model consisted of a goals module, for directing attention to the visual space, an inverse controller to move attention to left, right or central positions, and an input module for these three regions. Each module was taken to contain only three neurons, as dedicated nodes, these being average firing rate neurons with a sigmoid response function, with corresponding temporal sensitivity. Inhibition was taken to exist between the neurons in the attention movement controller, so slowing down attention movement in the invalidly cued case. Results from the simulation, in the two cases of exogenous and endogenous invalid cueing, are shown in Fig. 5a and b, respectively. In the first case, the CTOA was chosen to be 0.8 s, the second 0.2 s. As seen in Fig. 5a there is a rise of the incorrect (left-most) IMC neuron, directing attention incorrectly to the right. The correct right IMC neuron is then activated by the target input; this neuron then causes the incorrect IMC left-most neuron to be turned off by lateral inhibition. It also amplifies the input to the right plant neuron, so leading to recognition response. In the endogenous case (Fig. 5b), the incorrect IMC neuron is more rapidly turned off (due to the shorter CTOA), and the resulting plant neuron is turned on (and stays on, due to the goal module working memory). Varying the CTOA for the exogenous and endogenous cases leads to the two curves shown in Fig. 5c; these curves give a result close to that observed in humans (Wright and Ward, 1998).

4.4. Further simulations

In this sub-section, I will briefly review several recent simulations of other attention paradigms, which contribute to the overall understanding of how attention can work as a modulating process. The manner of this modulation is not yet unambiguous: it could act either in a multiplicative manner on input to a module or by a threshold variation. Simulations have been performed using either (Mozer and Sitton, 1999; Deco, 1999; LaBerge, 1995; Jackson et al., 1994).

There is vast range of attention experiments concerned with serial search through a set of distracters for a particular target. This has many results about the dependence of the search rate on the nature of the task: the rate is almost independent of the number of distracters if they differ along only one feature aspect from the target, while search is slow and dependent on the number of distracters if there are conjunctions of features that are needed to distinguish a target from the distracters. Thus, it is fast to pick out a red T from a set of blue ones but not so if one has to pick out a red T from a set of green T's and red L's, say. Simulations, under the general control framework discussed so far, have been able to replicate these search results (Deco, 1999; Corchs and Deco, 2001; Mozer and Sitton, 1999).

4.5. Relation to other models of attention

Besides the above quantitative models, there are also valuable models of attention of psychological form (essentially boxes and arrows); these have become increasingly sophisticated over the years.

First was the debate between the 'early' and 'late' models of attention; the early model supposed that attention modulation was applied to input selection at an early stage in the processing stream before object perception had occurred. The late models supported a late use of attention filtering. The controversy between early and late selection has now been somewhat settled by various detailed psychological paradigms, such as the attentional blink (AB), that will be described later. These paradigms showed that attention is, in the main, applied late. Objects can be analyzed up to




Fig. 5. The results of a simulation of the attention benefit valid/invalid cue paradigm, using the control model of Fig. 4 (Taylor and Rogers, 2002). (a) The temporal development of activity in the neurons of the IMC and the visual cortex, in the invalidly cued case, for an exogenous cue with CTOA = 500 ms. (b) Ditto, but now for the endogenous case, with CTOA = 200 ms. (c) The dependence of the attention benefit (invalid cue RT – valid cue RT) for different CTOAs, and separately for the exogenously and endogenously cues cases. Note the close similarity of this result to that for humans (Wright and Ward, 1998).

semantic level without attention being deployed in their direction.

Second was the debate as to the presence of top-down control systems as important and independent components of the attention system. On one side have been the results from brain imaging, single cell, and deficit data that there is indeed a separate control system for moving attention around, as discussed earlier in the paper (especially in Section 2). The other approach considers the most important part of attention as arising from 'biased competition' (Desimone and Duncan, 1995). The control system providing the bias is not considered as important here; the competitive process, taking place mainly in lower cortical sites, is the crucial determiner of the speed and accuracy of attention control. However, the conflict between the top-down and biased competition approaches is not very deep; the competition can occur, say in parietal lobe, in determining what area is to be modulated by attention but may need to draw on the additional competitive powers of earlier cortical regions to achieve highest efficiency and speed. That there is competition throughout the brain is well supported by experimental evidence, so there need be no contradiction between the biased compe-

tion and the control framework approach presented here. In particular, the bias itself in the former approach needs to be defined, which fits exactly with the goal module of the engineering approach. On the other hand, the biased competition approach indicates that IMC functionality may be spread about amongst various 'plant' sites.

Thirdly, there is much discussion over how objects are represented. The original and very influential 'feature integration' model of (Treisman, 1988) has been used in many discussions of psychological paradigms of attention; it is implemented in the control language in terms of the top-down control systems of Fig. 1, acting on the feature and object file maps of the controlled module in that figure. In terms of the more specific attention control model of Fig. 4, the object module there is to be identified with the feature and object file maps of feature integration theory. Binding of the various features of an object is achieved in Fig. 4, as in feature integration Theory, by attention from the spatial map, conjointly with bias from the goal module (if there is ambiguity, as in attending to an object when two objects are at the same point in space). There are further modules in the model of Fig. 4 (Forward and Monitor);

modulo these the feature integration theory can fit naturally into the control framework of Fig. 4.

4.6. Conclusions

We have been able to give support to a number of the control modules in the architecture of Fig. 2 from our discussion and the simulations. This has also been supported by the work of many others, as reported in Section 4.4. While the range of experiments being simulated has been relatively restricted, it has included the important analysis of the benefit gained in response time from attention being paid to an input. However, so far we have not concerned ourselves with attention to responses, where we have to bring in motor control. We have to take that further step, since without response the system has no ultimate value. Thus, we cannot leave this aspect out of our discussion. We turn, therefore, in the next section, to formulating an extension of our attention control model for motor control and motor response learning. Only then will we feel ready to develop an attention-based model of consciousness able to face up to its numerous problems.

5. Control of motor responses through attention

5.1. A general model of motor attention

It is now recognized that attention is crucial both in control of motor response and in learning to attain better motor control. Willingham has described (Willingham, 1998) how attention is used to set up conscious goals to achieve a desired motor response. If these goals fail, then further processing must occur: more detailed use of attention will be required to set spatial objectives (through parietal and prefrontal sites) as well as to learn the most effective sequences of motor actions on spatially defined objectives (in the supplementary motor area). Thus, attention and motor response are inextricably linked, at least in setting up and attaining good motor control. Ultimately response becomes automatic, when attention is no longer needed, unless errors then creep in to the motor responses due to changed context or other unexpected reasons.

It is presently unclear whether motor control under attention guidance is attained by a separate control system, sited mainly in the right hemisphere (Rushworth et al., 1997, 2001a,b), or by the action of attention control systems on symmetrically-placed motor 'intention' representations (Snyder et al., 2000; Andersen, 2002). We will assume that the difference between these points of view is only a semantic one so consider both possibilities as coming under the more general model of Fig. 6.

In Fig. 6, we use the same notation of separate modules for the attention control of sensory and of motor responses but can replace the term 'motor attention' by 'intention' with no loss of explanatory power. The model of Fig. 6 takes account of the facts that:

1. Attention is crucial to motor response, and its improvement by learning.
2. There may exist a relative separation of
 - (a) spatial attention in the right hemisphere, where ICM_{AV} , the inverse control module for attention to visual inputs, will be sited,
 - (b) motor attention in the left hemisphere, with the site of ICM_{AM} , the inverse control module of motor responses.
3. Attention is used to learn, by FBEL, response rules (in DLPFC) and correct motor response control signals from the IMC (in posterior parietal cortex).
4. The IMC for attention to motor response, IMC_{AM} , can also be regarded as coding for motor intentions of object \rightarrow action sequences.
5. Learning of the IMC for motor actions and the goal module structures are achieved by DA/ACh/NA error-based learning, very likely from sub-cortical sites.
6. The cerebellum may be used for
 - (a) the response-decision system,
 - (b) acting as a 'glue', to insert function 'words', both in language and other sequential processes, depending on cortex,
 - (c) control of the timing of sequences of responses.

The model of Fig. 6 has been used in simulations of several motor paradigms (Taylor and Fragopanagos, 2003), which we turn to next.

5.2. Simulating specific motor attention tasks

We will describe here two simulations, using the overall model of Fig. 6, which have been able to explain various psychological paradigms. The first (Schluter et al., 2001) required subjects to respond to objects appearing on a screen with specific finger movements (pressing a button with the first finger if one of two shapes appeared and pressing another button with the second finger if one of two other shapes was shown); one shape was presented on each trial with 3.75 s between trials. The average reaction time for this choice response task was about 570 ms. Reaction time to the simpler task, where only the same first finger was used to respond to all four shapes, was only about 210 ms. The two paradigms were simulated (Taylor and Fragopanagos, 2003) by means of the overall architecture of Fig. 6. The visual attention processing stage of Fig. 4 produced, as output from the object module, classification of the inputs into one or other of the possible object shapes.

The motor attention stage was also modeled on that of Fig. 6. The motor goal module coded for the eight possible object-response pairings; only the correct pairings produced a spontaneous output, so acting as a bias (a working memory) onto the motor attention control module. This latter again had eight neurons, also representing the two possible responses to each of the four shapes. It had inputs from the motor goal module, as well as from the object module as to which object was being presented. The output from the mo-

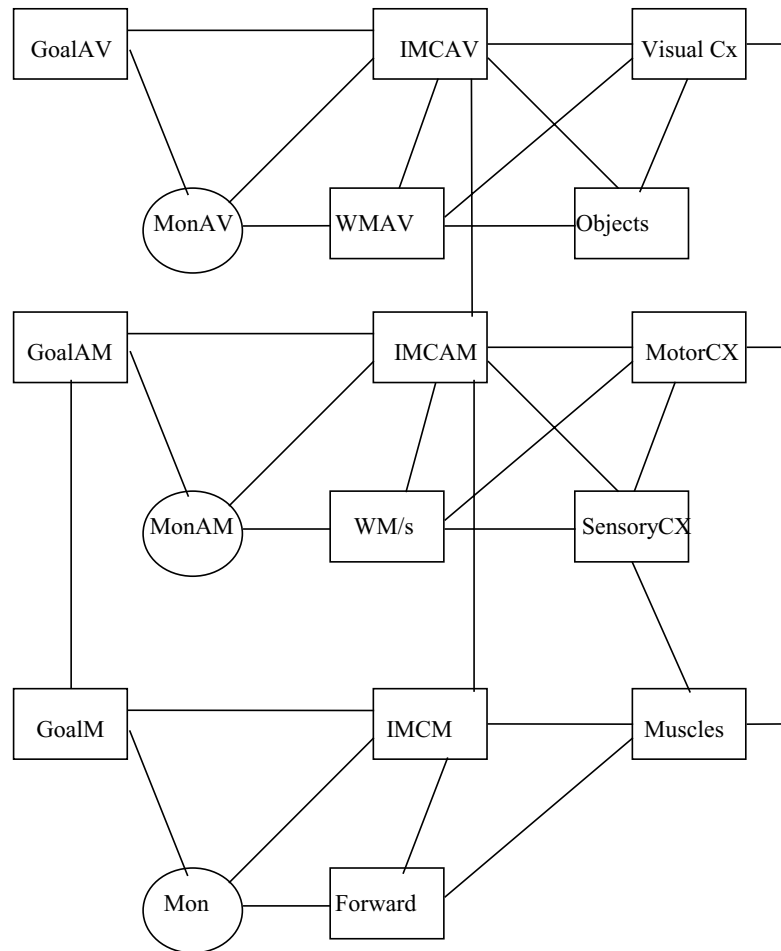


Fig. 6. The extension of the control model of Fig. 4 for the movement of visual attention to the joint case of visual and motor response attention: the AV postscript denoted attention to the visual modality, while AM denoted that to motor response. There is, in addition, a motor control system outside attention, which functions for automatic motor control, with its own smaller control system. However, the goals for it and for the attended motor responses are assumed equal (following Willingham, 1998). By symmetry we would expect a visual control system outside attention, but this is not specifically included here. Finally, we note that the visual attention controller is assumed to bias the controller for attended motor response; that is still to be tested. Note also that the forward models in both motor and visual attention are identified with working memory buffers, denoted WM. Note further that only one control system is shown for vision and motor response, while it is known that these are in both left and right hemispheres. However, there is a degree of asymmetry in functionality, as shown by Rushworth et al. (2001a,b), at least for humans.

tor attention control module was directed to the motor control module so as to determine which motor response should occur. Finally, the motor control module guided the motor response plant (the muscles) to give an action with one or other of the two relevant fingers. Learning occurred at the connections from the motor attention controller to the motor controller, determined by the error in response, by a suitable error monitor module. In all the simulations, suitable time lags were included between modules to model those occurring in the brain; they do not occur otherwise in the simple simulation due to the simplicity of the internal construction of each of the modules.

In the simple reaction time task, the visual input was an on-off switch with response independent of the particular shape of the input but depending only on whether there was an input at all. Thus, it was assumed that the visual input by-passed the visual and motor attention processing stages,

being fed directly onto the motor control processing stages by way of the motor control module.

After training had been completed, the simulation showed that the choice reaction time was about 550 ms while the simple reaction time was about 260 ms. This relatively good agreement between the experimental and simulation results is encouraging. The reason for the slower reaction in the choice response task to the simple task was because of the increased complexity of the object classification required for the complex task response. The architecture also used inhibitory connections internally in the modules, between competing object classes, which added to the time taken.

The architecture of Fig. 6 needs to be tested on further paradigms. In particular, degradation of some of the modules gave a better test of the overall model. To achieve that, the paradigm of Rushworth et al. (1997) was then considered (Taylor and Fragopanagos, 2003). This task is a motor

extension of the Posner paradigm, in that a cue is given to prepare the subject to make a cue-dependent motor action, but the cue itself may be valid or invalid. In the former case, there is a speed-up of response; in the latter, there is a retardation of that response. To model this process, a similar architecture to that of Fig. 6 was again used, to allow for the influence of the visual input to determine the motor response, at the same time including the effects of the invalid cueing in slowing response.

In the task, subjects fixated a white cross in the center of the screen. At equal distances above and below the cross, there appeared two hexagons, acting as pre-cues. The turning red of the border of one of these hexagons (upper or lower) acted as a cue as to which of the forthcoming targets would appear; the inner region of the corresponding hexagon then turned green to set the response in motion. The subject had then to make a speeded response to the relevant hexagon with the index finger pressing a button for the upper hexagon turning green and a middle finger button press for the lower hexagon. A valid cue corresponded to the green coloration of the hexagon occurring for the relevant red-colored border; an invalid cue arising from an upper green coloration for the hexagon after the lower hexagon's border had been colored red (or vice versa). Varying levels of cue validity were used: 85, 60, 40 and 15%.

The simulation of this paradigm was very similar in outline to that of the previous one:

- The visual goal module encoded the visual rules to help correctly classify the position of the cue (above or below the fixation point). On input of an invalid cue, the target position required modification of the goal output to redirect visual attention to the opposite hexagon. This was achieved by inhibition in the goal module, which added to the time for response.
- The visual attention movement controller modulated the input to the object module from the screen through being biased by input from the goal module (as specified above).
- The object module processed the cues as inputs and gave, as output, a goal-biased classification ('up' or 'down').
- The object module sent its output to the motor attention control module. This latter was biased by its goal module, which maintained a template of the motor response rules (index finger response to the top cue, middle finger response to the lower cue). All possible cue-response combinations were included in the goal module with learning to achieve correct responses from the motor attention controller by FBEL, as in the previous simulation of the choice response task.
- The inverse motor attention control module had 4 neurons for the 4 cue responses (upper or lower cue and the two possible fingers).
- The motor controller had two responses only: index or middle finger.

The experimental results were more testing of the overall model than the previous simulation, since they involved

Table 2

Comparison of results obtained on defects in attention benefit due to brain damage in cued motor decision making, following the paradigm of (Rushworth et al., 1997), with simulations using the attention control model of Fig. 6 (Taylor and Fragopanagos, 2003)

	Cue validity (%)			
	0.85	0.6	0.4	0.15
Control subject	350	370	450	400
Deficit R patient	450	470	550	530
Deficit L patient	570	650	980	920
Control simulation	400	380	480	430
Degraded R simulation	500	580	420	550
Degraded L simulation	700	650	1000	1050

both normal subjects and those with separate left or right parietal deficits. The manner in which these deficits affected the reaction times also depended non-trivially on the level of invalidity. Such deficits were simulated by modifying the neurons in the IMC either for the visual attention system (in the case of a right deficit subject) or for the motor attention system (when the left parietal was in deficit). The simulations were in relatively good agreement with those of the experiment, especially in overall trends; this is shown in Table 2.

The results follow the general trend of slowing of response when there is increased invalidity; more especially there is even slower response time for the left deficit subjects than either of the other two subject classes. This increased slowing is to be expected from the overall greater difficulty in directing response, even when validly cued, for the left parietal subjects, given the left bias of cortical control of motor response noted above.

5.3. Extensions

In the previous two subsections, the results of many experimental paradigms on orienting attention in vision, and in motor response learning, were used to develop and support the existence, and the siting in the brain, of the various modules of the overall engineering control architecture of Fig. 6. We described, in the previous subsection, simulations of two specific paradigms using this control architecture. The first was for the comparison of choice to simple reaction times to specific visual shapes (Schluter et al., 2001). The second used an extension, to motor responses to specific cues, of the Posner benefit paradigm (Rushworth et al., 1997). The second of these paradigms was also tested on subjects with either left or right hemisphere deficits. Both paradigms, including the deficits, were simulated with good agreement to the experimental results. In this subsection, we will discuss results expected from the attention control architecture of Fig. 6 for several further paradigms. We can only do this in a qualitative manner here but expect that the results will be in general agreement with what would arise from a more detailed quantitative simulation.

We start with a discussion of the results of transcranial magnetic stimulation (TMS) applied to the Posner benefit paradigm of [Rushworth et al. \(2001b\)](#). Localized TMS has been used in a variety of paradigms in the past to determine the relative contribution of various brain sites in achieving task solution by subjects. In [Rushworth et al. \(2001b\)](#), it was applied to the angular gyrus (ANG) and the supramarginal gyrus (SMG) in both a visual and a motor version of the Posner benefit paradigm. It was shown that only when TMS was applied to the right ANG was there an increase in the reaction time to orienting to invalid visual inputs (as compared to the left ANG or to the SMG in either hemisphere). By contrast, only TMS applied to the left SMG caused an increase in reaction time to invalid cues for motor response. This supports our proposed site of the IMC(AV) for controlling the orienting of attention in vision, to the right ANG, and the IMC(AM) for orienting motor attention, to the left SMG.

TMS could also have been applied to other sites in the brain during these paradigms. Let us consider what would be expected in terms of our attention control model of [Fig. 6](#). In particular, the modules for plant, for monitor, and for the goal holding could each have been accessed. This could have been done separately for the visual and for the motor components of the overall architecture. We have already considered the effects of deficits on the visual component ([Taylor and Rogers, 2002](#)), so let us only consider here the effect of TMS on the motor component of [Fig. 6](#). That leaves us to consider the plant, identified with the IMC(M) for automatic motor control, the Monitor for attention orientation, and the goal module. We suggested that these were sited in the network of motor cortex, cerebellum, and basal ganglia; in the anterior cingulate and cerebellum; and the pre-supplementary and pre-motor cortex respectively. How will TMS applied to each of these affect the various reaction times in the paradigm?

Applied to the plant, we would expect that TMS would cause a slowing of all response, to the valid, to the neutral, and to the invalid cues. This is because the plant, composed of distributed neural representations, would be degraded. We cannot determine the relative importance of the separate components of the plant network from our present position. That would only be possible by a more complete analysis, using the different architectures and suggested functionalities of these components.

When we turn to the monitor for attention orientation, we would expect that TMS, if applied at all stages of the paradigm, would initially cause a slowing of the process of learning the correct responses. For it is the monitor signal that is claimed to be used as the learning rate in FBEL and sent to the various sites (goal module, IMC(M), IMC(AV), etc.) to achieve the correct response. If this correct response had already been learnt, before application of TMS, then correct goal and inverse model control signal would have been generated, without the need for the monitor to guide learning. Thus, from our architecture of [Fig. 6](#), there would be

no degradation of performance in such a TMS application. If it were found that effects did occur, then one could conclude that some feedback activity were being used, through a forward model, to create an error signal to allow better responses. Such guidance would then be damaged by the applied TMS. Thus, we would expect longer reaction times in this situation again. However, here we predict that the invalid cueing case would suffer most since we expect that the monitor would be used to help correct, through visual feedback to the IMC, a better response. Thus, in this case, we would expect a differential change in reaction times on comparing valid to invalid to neutral.

Finally, we consider TMS applied to the goal module. If this were done during learning the paradigm, then we would expect an increase in error rates, or a slowing of learning to criterion, in all conditions. For in this case there would have been only degraded goal structures created in the goal module. If it were done after learning had been achieved, then again an increase in errors would be expected, especially under the condition of degraded inputs. There is expected to be a concomitant increase in response times since the degraded goal representations would only give a reduced bias to the IMC(AM) module. There may be a difference between the valid, neutral and invalid conditions, however: in the invalid case, the bias brought about by the degraded goal representation will be less strong, so modification to achieve the correct response will be expected to be faster than normal.

Let us consider the explicit learning paradigm of [Eliassen et al. \(2001\)](#). This paradigm used the serial reaction time task, in which the participants made two different button presses to the visual stimuli, an L or an R, presented to them on a screen. There was explicit learning (learning with the participants being advised that the stimuli followed a specific sequential order, which they should learn) or implicit learning (when they were not so advised). Alternatively, the stimuli appeared in random order. There was only found to be a reduction of response time to the stimuli in the explicit situation. Significant increased activity was observed, in the explicit learning situation, in the prefrontal and parietal (IPL) cortices and in the cerebellum. Here we note that if TMS had been applied during the learning process, we would expect that the reaction time reduction would have been slower. Also if there had been subjects with relative deficits in each of these sites, our architectural assignments, given earlier, would predict that (a) for prefrontal deficit, there would have been goal degradation; (b) for parietal (IPL) deficit there would have been IMC(AM) deficit, so slowing of orienting across all the valid, invalid and neutral paradigms of motor attention orienting; and (c) for cerebellar deficit, there would be either Mon(AM) or IMC(M) degradation. In the first case, this leads to slower learning, and possibly slower response, in the case of a feedback control error being used by an observer. In the second case again there would be slower responses, as well as increased errors.

Let us extend the results of [Sakai et al. \(2000\)](#) to this discussion, in particular using that the IPS is involved in the

IMC(AM) network. The paradigm of Sakai et al. (2000) emphasized choice reaction time under either response selection or timing selection respectively. The IPL was observed to be jointly active under these conditions as was the lateral pre-motor cortex PMCL. The latter was assigned by the authors to performing goal choice. At the same time, there were a number of other sites observed active separately, such as the supplementary motor cortex for response selection and the posterior lobe of the cerebellum for timing control. Again we expect that these assignments can be tested by TMS or in the case of subject deficit, leading to the deficits we have already discussed above.

Finally, we note that at least two recent publications have discussed how attention to motor response leads to modulation of motor cortex (Johansen-Berg and Matthews, 2002; Binkofski et al., 2002). That is why we have assigned the motor attention control network to act on the automatic motor plant network as composed of motor cortex, cerebellum and basal ganglia (plus possibly some parietal sites). Again we can predict what our model would predict to be the effects of TMS or deficits as applied to this automatic network. However, we still cannot disentangle the different contributions of the various subcomponents of this network.

5.4. Further work on motor attention

How can we extend the overall architecture of Fig. 6 to take account of either possible additional engineering control modules or to known extra features of brain processing? The most important of these extra modules is the observer module, which acts as a model of the plant. This can thereby update the IMC more rapidly than waiting for slow feedback to arrive. There is strong evidence for the presence of such a module from motor control paradigms (Sabes, 2000; Desmurget and Grafton, 2000; Wolpert and Gharharmini, 2000). Since these paradigms involve both automatic and attention control processes, such evidence may also be relevant to the presence of an observer in attended motor control. We will consider this later in more detail when we turn to analyze paradigms relevant to the CODAM model; this depends most crucially on the existence of the observer. In any case, there is an enormous amount of further work to develop detailed simulations of the many further experimental paradigms on attended motor responses, such as serial reaction time tasks or learning motor sequence response with various levels of attention applied to the task.

5.5. Conclusions

We have developed an overall model of attention control that covers both sensory and motor response modalities. The model still has a number of modules that have not been clearly recognized by experiment, especially the forward or observer module. However, we have not yet tried to determine in what manner consciousness might be embedded in this model. In the process we may obtain guidance as to

the possible function of the further modules. We turn to that search now.

6. From attention to consciousness

6.1. Introducing consciousness

In the previous sections, I presented a control model of attention. It was developed in a general form in Section 3 and then applied to particular aspects of visual attention in Section 4. The model was further extended to include motor learning in the previous section. Whilst there has not been a complete justification of the presence of all the modules in the overall control model of Fig. 6, there is support for many of the components. We will, thus, use the architecture of Fig. 6 as a basis for the further search for consciousness. We would expect that, if our control model of attention is of real worth, it will allow us a glimpse of the way that consciousness might arise as a component of the overall movement of attention. It will be through this movement of attention that awareness of an external object or of an action arises. We must now consider how that could occur.

We must be careful in the expectations we place on attention. In particular we must realize that whilst attention to an event is necessary for consciousness of it, attention is not in itself sufficient for consciousness: a blind-sighted subject benefits from a prior cue directing his attention to a blind-sighted part of his visual field in spite of having no awareness of that cue (Kentridge et al., 1999). There is also controversy as to whether in a very rapid stream of objects there is total unawareness of the second of the two objects (unavailable for later report) or only short-term awareness but ensuing amnesia of the object—so-called ‘attentional blindness’ (Wolfe, 1997; Resnik et al., 1997). There is also subliminal attention capture of objects, with coding up to semantic level, as numerous psychological experiments have shown (Rolke et al., 2001; Vogel et al., 1998).

In addition, a recent paper has proposed that attention and consciousness are distinct (Lamme, 2003). However, an important part of this claim, in which feedback is regarded as the only necessary ingredient to create consciousness, is rebuffed by the detailed account given in Taylor (2001a) of the large number of feedback systems in the brain that function outside consciousness. A simple example is that of the densely amnesic subject HM who had both hippocampi removed, yet he has no apparent defects of conscious processing except for long-term memory deficits. The hippocampus has remarkable lateral connectivity in CA3. This has now been shown crucial for pattern completion, since it is lost in transgenic mice with no such lateral connections. Thus, at least the hippocampus is an example of a neural network in which relaxation allows for pattern completion but is not helpful for ongoing consciousness.

The subtleties of the relation of attention to consciousness, mentioned above in association with blindsight,

do not change the fact that attention is the closest neuro-physiologically and psychologically well-researched faculty we possess to consciousness. It is *the* gateway to consciousness. As such, we will use it to guide our modeling of consciousness. The attention ‘gate’ has subtleties about it that will require careful consideration, but these can only be usefully discussed, as I have already pointed out, after an initial, suitably powerful model of attention has been created.

To start, we will consider consciousness itself. It is much harder than attention to define. An approachable definition (Bisiach, 1988) is that of “monitoring of internal representation”, so allowing access to one’s own brain activity, and relating it to an attention type of processing. This might be thought to leave out phenomenal experience, associated with the ‘what it is like to be’ character of consciousness that we are also trying to model; such phenomenal experience need not be involved with the process of active monitoring mentioned earlier, which is of introspective character. Both components of consciousness—that of content and that of the phenomenal self—are needed to be present in any overall model of consciousness. At the same time, the effects of brain damage and mental disease indicate further dissociations in consciousness. Thus, we see that there are aspects of consciousness that may need more careful discussion. In the light of the dangers of ‘premature definition’ (Crick and Koch, 1990), it is appropriate to leave further definition at this point.

Before so doing, it is important for later analysis to note that I have earlier noted that the perceptual component of attention involves activity in the parietal lobes, as associated with certain of the control components of attention discussed above. This, and related deficits of awareness arising from parietal damage, has led to the parietal lobes being regarded as a crucial region for consciousness creation (Taylor, 2001a). This has been supported by recent brain imaging results, such as when there is awareness of emotionally-relevant inputs during extinction and neglect experiments (Vuilleumeir et al., 2002); the parietal lobes were especially important in those tests, increasing activity with awareness of the inputs.

6.2. *The creation of consciousness by CODAM*

The control model of visual attention discussed in Sections 4 and 5 is now ready to be analyzed for its ability to support the additional experience of consciousness of an input being attended. For that we develop ideas, current since William James’ important work (James, 1890), on the need for neural activity to have a suitable duration in order for consciousness to arise from the relevant input. This has been explored and updated in more detail elsewhere (Taylor, 1999), but to summarize: a competition on a working memory buffer, with attendant ‘bubbles’—thereby temporally extended—of neural activity (Taylor, 2001a), leads to the creation of the necessary activity, suitably based

on past context, salience and long-term experience. There is, in particular, in the confluent ‘central representation’ of the buffer working memory sites of parietal lobes, information about the state of the body, as well as of intentions for action, from the excellent connectivity possessed by the parietal lobes with the requisite sites.

In spite of this connectedness and extended duration of neural activity, there is no obvious reason for inner experience to arise. A robot built with the above internal neural machinery would be expected to be a zombie. Since this position is one at odds with all the models of consciousness produced till now, including my own (Taylor, 1999), it needs to be discussed at more length.

What has been missing from all of these earlier models, in spite of the claims of their authors that they are sufficient for consciousness? It is that any neural system, constructed on previously proposed architectures and mechanisms, would have no sense of ownership. For example, in many models the neural activity arises from a set of coupled attractor nets. However, there is no mechanism for the existence of an associated ‘inner feel’ to such activity. There is no mechanism to provide the sense of ‘holding’ or of ‘possession’ of it. That, I claim, is what is missing from all of the models proposed so far. It is a fatal flaw, since there is nothing ‘it is like to be’ those systems. They have content but no breath of conscious experience.

In order to insert some form of inner experience involved with ownership, let us return to the control models of Figs. 4 and 6. An input to be attended to (chosen by its salience, either from high-level goals or due to external demand) needs to have a suitable attention movement signal to be created by the sensory IMC of Figs. 4 and 6, to focus attention on the input and catch and hold it, at the same time repelling distracters. There is a corollary discharge (CD) signal arising from this movement of attention. That will arrive at the monitor some hundred or so milliseconds before any signal from the amplified input, whose amplification is turned on by the original signal from the controller. Such amplification requires some 100 or so milliseconds to be achieved (Reynolds et al., 1999; Mehta et al., 2000; Hopf et al., 2000). Thus, the CD needs to rest on a buffer before the comparison can be made in the monitor. We could have assumed that this buffer is the monitor itself. However, when the amplified afferent signal arrives at the monitor, there must be cancellation there of the CD activity: there is then no error. Related conscious experience, in that approach, would then be annihilated, which would be too rapid for experimental results on the development of consciousness (Libet, 2001). It would also have too brief an existence, ceasing after only about two hundred milliseconds. As shown in Fig. 7 (only given for sensory attention), we need to introduce separate components of the forward model of Fig. 4 for storage of activity from the afferent input and from the IMC. These are denoted in Fig. 7 as WMsens and WMcd respectively (WM: working memory buffer here, ‘sens’ denotes the relevant sensory modality: see the ‘Glossary’ for further




Fig. 7. The CODAM model for the pre-reflective self. The working memory buffers are divided into one for amplified input (suitably thresholded), denoted WMV, and one for holding the efference copy/CODAM, denoted by WMcd. It is this latter buffered signal which is regarded, in CODAM, as providing the experience of ownership of the about-to-be-experienced attended input.

definitions), which could be adjacent in the parietal lobes (or even in separate cortical layers in the same region). It is this extension of the buffer-monitor components of Fig. 4 that we explore further.

The mechanism of the creation of consciousness in the coupled modules of Fig. 7 is proposed to be as follows. Firstly, the IMC produces an attention movement control signal, under either exogenous or endogenous control. It also emits a CD signal, which is buffered on the specialized working memory site WMcd so as to be available for error assessment in the monitor. After some period of time, the plant control signal from the IMC will have done its work, with amplification of the afferent signal being achieved and concomitant inhibition of distracters in the visual scene. One possible scenario is as follows. The amplified afferent plant signal arrives at the monitor and turns it off. It is reasonable to assume that, before then, inhibition of relevant sites on the WMsens buffer had occurred (to avoid error occurring on it through incorrect activation) or a high enough threshold of WMsens neurons exists. Those sites are now freed from inhibition by cancellation of the inhibitory signal by the incoming amplified input or are now strong enough to overcome the threshold. Thus, WMsens can now be accessed by the attentionally amplified input. Awareness of the input (at least spatially) is now supposed to ensue, along with more detailed ‘filling out’ of the content by suitably coupled activations to associated feature and object maps. Yet perspectivalness still appears missing from our analysis. To determine how perspectivalness might be finally incorporated, let us probe the system of Fig. 7 further. It will lead to a slightly different, but more complete, scenario than above.

There is a gap of time before the amplified afferent signal arrives at the monitor to cancel the CD signal waiting there. This will be so both for the exogenous movement (driven by a direct sub-cortical input to the IMC) or the endogenous case (with a signal from the goal module in prefrontal cortex used to bias the IMC). During this time period the CD signal will be buffered on WMcd. The basic proposal of the CODAM model (Taylor, 2000c, 2002a,b) is that this briefly buffered signal (possibly together with suitably bound relevant activity in other sites) provides the conscious experience of ownership and gives the sense of inner perspective.

We turn to consider this process, and related points about ownership and agency, next.

6.3. Ownership and agency: components of the self

The self has a variety of components, involving a subjective as well as an objective character (Mitchell, 1994). Here, we are concerned with subjective aspects of it going under the descriptions of ownership and agency (Gallagher, 1999, 2000). It is important to distinguish between these. I can consider that I am the agent of a movement of my arm as I move it. On the other hand, if my hand is moved passively by someone else, I can realize that is so, but still know that it is my arm that is being moved. Thus, agency and ownership are distinct. This is especially so in terms of the possibility of error. Thus, I cannot be in error when I claim that my arm moved. Similarly I cannot be in error that it is I who feel pain. It is not sensible to ask “are you sure it is you who feels pain?”. This important feature of conscious experience is what has been termed ‘immunity to error through misidentification with respect to the first person pronoun’ (Shoemaker, 1984). Such an error can occur over the attribution of agency, as is known by tests on subjects viewing their gloved hand moving (Fournier and Jeannerod, 1998). An experimenter’s hand can replace (by suitable image tricks) the subject’s own hand without them realizing, provided the movements made by the experimenter are not too different from those of the subject. On the other hand it is difficult to conceive of attributing to someone else the inner conscious experience involved with ownership.

It is a reasonable thesis that the most primitive form of self-knowledge is that of the ownership of the movement of attention. This can arise in any animal with the most simple attention control system. It is not even necessary to possess a frontal goal module (possibly due to lack of frontal cortex), since attention could still be moved exogenously. As noted earlier, such low-level attention movement control occurs for rapid inputs: they gain nearly automatic access to the attention movement controller, the IMC (it should be called the AMC, but we stick to the control term ‘IMC’), as known by many studies of attention shifting (Wright and Ward, 1998). Such exogenous movement can also occur without

the need for peripheral feedback. It is therefore much more rapid than the endogenous variety (Wright and Ward, 1998). This mechanism of ownership, proposed here as arising from the buffered attention movement corollary discharge, helps explain the continued consciousness of subjects who have lost all sense of such feedback by de-afferentation (Cole and Paillard, 1995). They should have suffered severe deficits in their sense of self if there was a solely bodily basis for that; they did not.

The claim now being made here is that the buffered CD signal produces a conscious experience of ownership, breathing the light of inner experience into temporally extended neural activity. The robot thereby equipped would no longer be a zombie. It would have an inner feel, of the ownership of the amplified input brought about by the short-lived CD signal on WMcd. The claim is based on the related one that non-perspectival content-full consciousness supposedly arises by temporally-extended activity on a suitably well-connected buffer site, already posited in locating the creation of the experience of non-perspectival consciousness in the buffer site of Fig. 7 (Taylor, 1999, 2001a). For the coupled monitor/buffer sites of Fig. 7 the new item of information is that of one's own movement of one's attention. It is this item that is the spark of 'how it is like to be conscious'.

An important feature of the mechanism of Fig. 7, as described in more detail above, is that the sense of ownership, claimed to be created by the CD signal buffered on WMcd for a suitable length of time before being annihilated by the amplified plant signal arriving at the monitor, is indeed immune to error through misidentification of the first person pronoun, without considerable, if not total, breakdown of attention. This can be seen for the undamaged system: the sense of ownership engendered by the CD signal on WMcd can only be cancelled by the appropriate afferent input amplified by the original attention movement signal itself. Thus, no conscious experience can arise until the CD signal arrives at WMcd. Content-full conscious experience of the amplified input can only then be achieved by cancellation (or reduction) of this buffered CD signal, which can only happen if there is identity of the two attention control signals, that for the amplification and that given by a non-degraded copy of that signal. Since these two signals are identical, then attribution of ownership through the CD signal experienced on WMcd can only be to the correctly amplified input produced by the original attention movement signal. There is, thus, no chance of error in attributing such an experience to oneself.

We note that this semi-identification of the representations of WMcd and WMsens must occur in a manner in which activations in WMcd are not bound to those in lower sites in the same manner in which WMsens presumably is (so as to give consciousness detailed content). Such binding can occur either by amplitude modulation or by phase synchronisation. Content is not supposed to enter the pre-reflective self created, in the CODAM model, by activa-

tions on WMcd; the latter can therefore have no binding to lower-level content. How this important 'hiding' of WMcd from the outside world is achieved is presently unknown.

There can be damage to this subtle circuitry. For example, it is possible that the CD signal could be lost, due to destruction of WMcd or of the monitor, but there was still an amplified input to WMsens. If the monitor was damaged then there would be no error correction, and so the attention system could have severely compromised top-down control. If the corollary discharge component of the buffer site case were lost (so no CD signal preserved), but with an intact monitor, there would still be severe problems of attention control, since error values could be incorrect in the short term (since no forward model to speed up attention-based response) but be corrected by feedback information. In either case, not only the sense of ownership but also any top-down attention control would be damaged, leading to deficits in any resulting experience.

Is there someone actually reading off the movement of attention? In other words, is there still a homunculus hidden in the works? Do we need a 'reader' to keep track of the persistent neural activity claimed to be at the root of consciousness? If there were such an agent, possessed with intentional powers, then we would be back on the treadmill of the infinite regress: the homunculus would then need another homunculus inside it to cause it to act, and so on. However, the CODAM model does not assume that the activity on WMcd has agentive powers. The homunculus has been deconstructed so completely that it has no 'will'; it is no center of power. It only responds to incoming neural activity and then interacts in a suitable manner to grant immunity to error of misidentification of the first person pronoun mentioned earlier. The WMcd activity grants solely ownership, not agency. The origin of movement is from the goal modules in the prefrontal cortex, themselves being driven by salience in orbito-frontal cortex. The deconstruction of the homunculus is complete: only its most vestigial powers are left: only the experience of 'what it is like to be'. In spite of this deconstruction there is continuity of experience, given by short- and long-term memory representations. The former grant continuity over the short-term—of orders of seconds, while the latter provide a permanent store of 'self' descriptors that provide a basis for personality.

Finally, we consider features of consciousness left out so far: how are co-occurring events in different modalities, say in touch and hearing, brought into a unified consciousness when attention is drawn to them? Further, how is spontaneous conscious thought brought about? The first of these questions can begin to be answered in terms of results arising from brain imaging of multimodal attention processing (Downar et al., 2000): besides areas of cortex involved in the control of attention in separate modalities, there are sites observed that are common to all of them. Thus, we can suppose there are higher-order control sites for multimodal attention movement, especially buffer sites for the corollary discharge of the attention movement controlled by these sites: these,

by CODAM, would bring about co-occurring conscious experience across modalities. Alternatively there may be apparent co-occurrent conscious experience that in truth arises from a rapid switch from one modality to another, too fast for it to arise in experience. Only detailed experiments could distinguish these two (or other) alternatives.

The second question, of spontaneous conscious experience, can be explained in terms of internally generated neural activity. This could arise in an associative manner from previous activity (as by ‘chaining’) or by noise in the brain bringing about completely new neural activity. Such novel activity would bring attention to bear on it, and hence arise in consciousness, by CODAM. This virgin activity would use previous chains or schemata to allow for more rapid development (attended from thought to thought, and by CODAM arising into consciousness), especially if the thought was goal directed; in general, already present memory representations would be used to give sense and direction to the thought stream. In this way, we see that CODAM can rise to further challenges about experience.

6.4. Support for CODAM

There is support for this model from various forms of deficits and from other sources. For example, anosognosia, the denial that one has a deficit brought about by some form of brain damage, dissociates into various forms (Berti et al., 1996). These forms were separated into impairments for upper and for lower limbs, and separately for cognitive loss. It was concluded that the only suitable model for such dissociations involved the loss of “the functioning of monitoring modules (possibly located in the parietal lobes)”. A similar model had been proposed independently (Heilman, 1991).

Further support for the presence of some form of monitor at the basis of phenomenological experience comes from brain imaging experiments using ‘introspective’ types of paradigm, where a subject has to internally monitor their sensory experience. One such is the motion after-effect (MAE: Taylor et al., 2000). Here, the subject has to press a button when their MAE experience ends. Results of fMRI brain imaging of such subjects showed two networks of active cortical areas, one involving posterior sites, the other those more anterior. The posterior sites are mainly involved as motion processing filters, as in the motion sensitive area MT. The anterior sites appear to act as templates set up for testing the experience of the MAE, being active mainly during that period, as seen from the time series. A crucial site in the posterior network was an area in the inferior parietal lobes, suggested as acting as the site of the creation of the perceptual experience of the MAE.

A similarly important inferior parietal site was detected during binocular rivalry switching (Lumer et al., 1998). In this case, the subjects made no motor response, the fMRI time series being used to detect similar temporal structure across a variety of areas. Other attention-shifting paradigms have also supported a crucial role for parietal sites (as well as

related modules in the attention network mentioned earlier) (Konishi et al., 1998). There is also trans-cranial magnetic stimulation (TMS) support for the key role of parietal lobes in conjunction search, since application of TMS about 100 milliseconds after stimulus onset causes a significant delay in response time in the paradigm. Finally the GO-NOGO paradigm also provides evidence for a monitoring role of the parietal lobes (Shibata and Ioannides, 2001).

In conclusion, evidence has been presented that there exists a set of monitors in the parietal lobes for attention paid to various modalities. As importantly, evidence was presented for there being a corollary discharge signal of the movement of attention being sent to various sites, a key requirement for the model presented here, to add to the (single cell) evidence for its existence presented in the previous section. This CD signal is generated as output from the IMC, proposed sited in the intra-parietal sulcus in the parietal lobes, and possibly in other sites, such as the SPL. For example it is observed as pre-saccadic eye movement activity in FEF, SEF and LIP cells. The well-observed buffer working memory sites are also placed in parietal lobes. We therefore consider there to be sufficient foundation for the consciousness model of Fig. 7, for it to be used as the basis of further analysis. Analysis will be given later (in Section 8) of specific experimental paradigms relevant to the validity of the CODAM model. We next turn to extend the consciousness model of Fig. 7 to self-consciousness, where frontal sites do play a crucial role.

6.5. The experience of self

As noted earlier, self has a number of aspects. The one of concern in this section is that of agency, which is more advanced than the most primitive of all, that of ownership. Agency has been analyzed empirically through the fact that there is loss of the sense of agency in schizophrenics suffering from delusions of control (Feinberg, 1978; Frith, 1992; Gallagher, 1999, 2000). For them there is degradation of signals from the site of intentionality of a particular motor action or of language production in the SMA. This approach has led to numerous important insights into the disease (Duggan and Frith, 2000). However, it does not deal directly with consciousness itself but only with internal monitoring of motor actions (which can occur unconsciously). For example, one does not think about having a thought in order to have it: agency is not enough for consciousness. But experience of agency is important. In particular, let me consider more directly, in terms of the model of the previous section, how consciousness of self at the level of agency or intention could arise.

One possibility is that the experience of agency arises as that of the movement of motor attention. In other words it is the experience of motor attention ownership. This could arise from an extension of Fig. 4 by inclusion of WM structures of motor attention control and its corollary discharge, in a similar way that we extended Fig. 4 to Fig. 7 to create

the pre-reflective self. It functions by buffering the motor attention control signal (on the motor equivalent of WMcd) with later receipt of the motor feedback signal on the motor equivalent of WMsens.

The existence of the latter is shown by a recent paper (Johnson et al., 2002). In the paradigm, movements are made by subjects to a target light in front of them, which is moved rapidly, a short distance to left or right, 25 ms after the start of their movement. The subject has either to reach to the final position of the target as rapidly as follows, or to reach in the opposite direction. The subject then repeats their movement shortly after the movement has been completed, to show how much ‘motor awareness’ of it they possessed (so assumedly from activity buffered on an appropriate WM site). Rapid errors (those errors towards the moved target) were both made and experienced by the subject, as were the voluntary corrections made later. Both these movement modifications indicate that the motor attention WM is activated, the early movement being identifiable as under exogenous control, the later one being endogenously controlled, assumedly by a frontal goal module.

This experiment shows the existence of a separate motor awareness system, which we identify with that supposedly predominantly in the left hemisphere according to (Rushworth et al., 1997, 2001a,b). In total, we can consider motor agency as arising from activation in an attention motor control buffer equivalent to the buffer WMsens for sensory awareness. But there is clearly considerable further work to perform before the nature of motor attention/intention and of motor awareness becomes well founded.

6.6. Limitations

We next need to consider what fundamental limitation any scientific model of consciousness possesses by virtue of the nature of consciousness itself. For consciousness arises from activity, be it very subtle, in the brain. In the same way that a model of the weather is not the weather itself—it does not rain or have hot sun beating down by means of the equations of flow of moist air or whatever—so CODAM cannot ‘be’ conscious. It can only try to capture the bare principles of the experience. As discussed elsewhere in this paper, at best we can only capture the essential features of conscious experience—its two components of ipseity and content, of ownership and owned qualia, and so on. The various experiential components can only be shown to be present, in principle, in CODAM, not be actually experienced themselves by you, the reader.

These points are relatively obvious, but the further question then arises as to how could CODAM help guide the construction of machine-like consciousness? This is a deeply sensitive question to many, since it could be claimed that there is no way of probing the inner experience of such a machine, built on CODAM principles. For how would we ever be able to have the so-called inner experience of the

machine, however it is constructed? But in spite of our inability to give a complete definition, we do not, for example, stop having children because we cannot justify that they are in fact, or will become, conscious. We know, by close contact, that they are so. We may sadly find, after birth, that such consciousness is limited due to a birth or developmental defect. But that does not mean we cannot intuit this consciousness, even of a more limited form, in others.

There is a similar argument with machines. With constantly increasing computing power, there is both the possibility and the need for such an experiment to be made: to construct a system on CODAM lines. Would it then be conscious?

The most persuasive approach would be to construct a system composed of ‘active’ neurons, with membrane potential continuously active, and spike activity being transferred between such neurons in a similar manner to those in our own brains. A system so constructed, with attention capabilities along the CODAM lines, would be a good candidate for a conscious system. To discuss how to prove (beyond the above-mentioned intuition) that it so goes beyond this paper, requiring a detailed discussion of the set of required ‘probing’ experiments (as in the Turing test). However, in a nutshell, the main thrust would be that all possible experiments that could be performed on the system would lead to observation of two types:

1. Responses expected by a conscious system, including verbal report of apparently autonomous origin.
2. Inner activity (probed by suitable ‘electrodes’) that corresponds to that observed in similar regions in our own brains.

That is how we tell that our children are conscious, especially if they appear to have some defect in their conscious experience, so must be studied closely. It should be the same for a CODAM-based conscious machine.

Thus, there is no reason, in principle, why CODAM could not be used to develop a machine that could be investigated for possible conscious-like activity. However, the description of the development of such a machine is far beyond the scope of this paper.

6.7. Conclusions

In this section, we have extended the attention control models of Sections 4 and 5 to incorporate further control signals, suitably buffered, as well as a buffer to hold inputs being amplified by attention. The first of these, the corollary discharge buffer, is claimed here to provide the experience of ownership of the about-to-be-experienced content of the input to which attention is being paid. It is this inner experience that, it is claimed, provides the experience of ‘what it is like to be’ and bridges the ‘explanatory gap’. It does so without resorting to non-material or non-brain-like elements. As such, it needs to be compared with the

ongoing streams of attack on the problem of consciousness from other directions. That we turn to next.

7. Relation of CODAM to philosophies of mind

It is impossible, in all appropriate details, to relate the CODAM approach to the mind to the enormous philosophical literature on the nature of mind. I will discuss this very briefly only in terms of two distinct strands: (1) Western cognitive science, based on recent discussions (Damasio, 1994; Dennett, 1991; Crick and Koch, 1990, 1998, 2003) and (2) Western phenomenology and Eastern mysticism.

7.1. Western cognitive science

According to Western cognitive science, there is no ‘Ghost in the Machine’ (Ryle, 1952). By this is meant that there are not two distinct worlds, one an inner, mental one, the other that of the observable bodily functions, including the activities of the brain in all its subtlety (Dennett, 1991; Cotterill, 1989; Damasio, 1994; Crick, 1994; Crick and Koch, 2003). However, the claims and theories of the various approaches to consciousness advocated with great heat and strength by the various advocates above (and many others) all stem from the doctrine that there is nothing other than ‘intentionality’ (as the philosopher Brentano put it): all is content in consciousness; there is nothing else. In other words there can only be ‘consciousness of’. There cannot just be ‘consciousness’. Even earlier, the British Empiricist philosopher David Hume wrote that he could find ‘nothing other than a bundle of perceptions’ when he tried to catch his ‘inner self’. It is only that bundle of perceptions that has been considered by Western cognitive science and the writers referred to above. Anything else would seem, in the writings of these researchers, to smack of dualism. Strict attention can therefore only be paid to the contents of consciousness.

But this approach leads to the basic and deep problem: how can a sense of self, answering the question of ‘what it is like to be’, arise? Such a possibility is either ignored (Crick, 1994; Damasio, 1994; Edelman, 1992) or its relevance denied very strongly (Dennett, 1991). Those who realize that there is a missing element (Levine, 1983; Nagel, 1974; Chalmers, 1996) have attempted to fill it with an independent, dualistic-like ‘stuff’, such as information (Chalmers, 1996). If they accept the real challenge of consciousness, concerning the ‘explanatory gap’ (Levine, 1983) or ‘the what it is like to be’ problem (Nagel, 1974), but are still strongly shackled by Western cognitive science’s denial of all but content in consciousness, they can only add extra elements to those of the material world. They can only become dualists in all but name. All the other researchers in that galimaufry of the ‘consciousness explained’ camp cannot bite the bullet of the self; they are either left to claim that there is no inner self at all (Dennett, 1991; Crick and Koch, 2003) or

are left floundering as to what is still missing (Crick, 1994; Damasio, 1994).

7.2. The other side of consciousness

There is a vast richness of the understanding of the nature of consciousness from the two areas of Western phenomenology and Eastern meditation, which are noticeably absent from the Western cognitive science approach. It is difficult to select from the panoply of insight, but some brief comments are in order, in the light of the important discussions in Gallagher and Shear (1999). Thus, Strawson, 1999 reaches the conclusion, from a phenomenological study of his own consciousness, that “there are eruptions of consciousness from a substrate of apparent non-consciousness” (p. 21). From a different philosophical viewpoint, Hayward (1999) notes from meditation experience, “between moments of experience of self are gaps in which there is no sense of self or of separateness from what is being experienced” (p. 390). These are of great importance to the model presented here. Such features arise from the above CODAM model.

To see how the ‘gappy’ nature of experience would be present, let us focus once again on the arrival of the CD signal on its buffer site WMcd. This persists for only a hundred or so milliseconds before it is annihilated (or reduced). It was claimed earlier that the ‘ownership experience’, crucial to provide consciousness with the feeling of ‘being there’, was provided solely by that signal. Thus, the experience of self at this most primitive level is expected to occur only in fits and starts; that is the essence of the gappy experience itself. This is supported by the fMRI observation of inferior parietal (and other) sites solely activated by the switch in binocular rivalry, as well as the MAE data reported earlier, where again there is a switch of perception, and so of attention (Taylor et al., 2000).

The relevance of the briefly occurring CD signal to phenomenology is made even stronger in terms of the analysis of Shear (Gallagher and Shear, 1999). He discusses Kant’s demonstration that the self cannot be experienced or even defined in terms of any empirical quality at all. This can be reconciled with the model presented above, as well as with descriptions from Eastern writings. These claim that meditation training can ultimately lead to a state of consciousness termed ‘pure’. This is empty of any content, but is not empty. It is sometimes called the pure consciousness experience, or PCE, an acronym I will use here. As quoted from the writings of a Japanese Zen mystic (Gallagher, 1999, p. 413),

But it is not vacant emptiness. Rather it is the purest condition of our existence.

In ancient Indian writings (Gallagher, 1999, p. 413), the pure consciousness state is described as

It is unseen, incapable of being spoken, ungraspable, without any distinctive marks, unthinkable, un-nameable . . .

He who knows it, thus, enters the self with his self.

What can be so indefinable yet be called a state of ‘nirvana’ and be attempted to be attained only by considerable concentration and practice? To be able to understand such a level of consciousness, let me summarize the process of the creation of inner experience according to the discussion of [Section 6](#):

- The input scene is analyzed by a salience net.
- An appropriate control signal is created by the goal module, thence biasing the IMC to send out an attention amplification/inhibition signal to the input and distracters.
- A CD signal of this control signal is buffered on the WMcd, which preserves the signal, and creates the experience of ownership of what is to come.
- In the meantime the control signal is sent to the feature maps to selectively amplify the chosen input features.
- The amplified signal arrives at WMcd, and annihilates (or reduces) the CD signal there.
- The amplified signal is now accepted on its own buffer site WMsens, and experience of the attended input now arises in close succession to the sense of ownership of the experience.
- A new input is then chosen for selective amplification, and this goes through the above steps, replacing the previous buffered input after a competition.

I have already noted how the above description leads to the expectation of ‘gaps’ in the activation levels of the buffer modules between the various steps, in the changeover process from a particular winner on WMsens to the next winner taking its place. In more detail, there will be a portion of that interval during which there is buffered activity on WMcd from the CD of the signal creating the new amplified input; that may be relatively short. There is, therefore, a gap of ‘ownership’ activity from the moment of winning on WMsens, and annihilation of the signal on WMcd, to the next signal on WMcd. It is consistent with the articles of Strawson and Hayward to propose that perspectival consciousness only arises from new activity emerging, through attention modulation and competition, to be a winner, with a copy of the attention control signal buffered on WMcd. Only a non-perspectival form of experience arises from activity on the WMsens site over the alternate periods of time, when there is no (or reduced) WMcd activity.

The state of nirvana mentioned earlier can now be understood as suitably extended activation of the WMcd buffer, obtained with no amplified input of any sort. This would very likely be achieved by lengthy meditation training of the input, from a suitable frontal goal state, so as to cause it to generate the self-perpetuating state on WMcd. At the same time, the activation brought about on WMcd, or on the goal module in prefrontal cortex set up by meditation, has been developed to be strong enough to fight off all competitors arriving from associative cortices in all modalities, a property that must be a part of the complex goal state to

arrive at PCE. The CD signal thereby continues its life on its buffer, so leading to the pure consciousness experience, with no concomitant WMsens activity.

The exact nature of the temporally extended activation on WMcd during pure consciousness is unclear. It may be a uniform one, occurring across the whole WMcd module. On the other hand there may be alternate activations able to produce this experience of nirvana. In any of these activations there will be no empirical quality, only continued experience of ‘pure’ ownership of one’s own attention apparatus (although with some hint of spatial extent). That is close to the desired goal of a self-referential system. However, the referent has no content since there is no external input coupled to the WMcd site, not even arising from the internal bodily environment.

The conclusion of this discussion is that the discreteness of conscious experience is a natural component of a version of the model. Moreover, the pure conscious state can be understood as arising from activation solely on the WMcd buffer brought about by internal intention with no amplified inputs of any sort.

The above explanation of pure consciousness leads to expectation that a further ‘pure’ consciousness experience could arise by learning to activate an extended goal state composed of the PCE goal state, plus separately those for normal life.

There is some hint of the existence of a further state, differentiated by a switch over between the two states. This is denoted as state IX: ‘The Stage of Ongoing Enlightened Traits’ in the table of Advanced Extraordinary States of Consciousness in [Austin \(1998\)](#). In this state, there is a higher level of sensate perception, as compared to no registration of the external world in pure consciousness (denoted state VIII). This difference implies that state IX may be a continuation of state VIII achieved without complete destruction of the external input on WMsens but with continued parallel activation of WMcd. This is the parallel (continued) ‘PCE + Normal’ processing expected from the extended goal. Final awareness of agency would then arise solely from the motor IMC signal amplifying, by attention, the motor control signal accessing its buffer working memory site by peripheral feedback. Note that the PCE state, together with those developing from it, mentioned earlier, will still be expected to be unitary, since these states are constructed by reduction of all sensory and motor feedback from peripheral cortical sites.

7.3. Conclusions

The CODAM model is consistent with the existence of the Eastern meditation state of Pure Consciousness. The model is not in accord with the Western Cognitive Science approach to the mind, in which there is nothing but content. For CODAM, there is the crucial brief period of ownership of the amplified, and about to be experienced, content of the input. Which is correct: CODAM or Western Cognitive

Science? We turn in the next section to analysis of possible signatures for CODAM.

8. Signatures for CODAM

8.1. General signatures

We have considered, in earlier sections, the spatial aspects of the attention control system in the brain and its extension to incorporate conscious experience. There is now a wealth of brain imaging, single cell and deficit data, to support the general framework presented so far in this paper (some of which has been presented earlier): attention can effectively be viewed by means of an engineering control approach, with associated inverse, goal, monitor, and forward models (although the evidence for the existence of the latter is not as strong as it might be). The various networks of brain modules observed active in different attention paradigms can be assigned to the various functions expected.

However, there is a crucial aspect of the control system we have not yet properly explored: the temporal domain. Any set of forward and inverse models, implemented in a set of neural modules, will have a well-defined temporal flow of activity. Furthermore the CODAM model extension, designed to create the ownership of conscious experience, will have even more specific predictions. What is the situation with respect to such an extension, both from the point of view of predictions from the models and the experimental situation? We turn to develop answers to these questions in the following parts of this section. We expect to find specific timing patterns of the flow of activity which should be crucial in developing the CODAM model further, as well as helping understand, in a more detailed manner, the set of event related potentials (ERPs) obtained by averaging EEG signals over many trials (to remove noise), and related MEG signals seen in many studies. At the same time, we should be able to begin to derive explanations of important temporal attention paradigms, such as the AB, the perceptual refractory period (PRP), and the nature of simultaneous extinction. This will be natural to relate to the more traditional serial search results we already discussed briefly at the end of Section 4. Moreover, we should be able to cast some light on new results on awareness in motor responses (Johnson et al., 2002), mentioned earlier. We have already discussed that paradigm under aspects of the self in Section 6.5. Here, we will only consider sensory attention paradigms with especially challenging temporal dynamic features.

8.2. The temporal flow of activity in cortex: the Ni/Pi sequence

There are early electrical signals in cortex before the first 100 ms after stimulus onset, such as the C1. However, these activations are now known to be attention independent, corresponding to the forward flow of activity from input sites

in the retina to the thalamus and thence to the early visual cortices (V1, V2, V3, V4, . . .). The later sequence of distinct ERP signals, N1/P1, N2/P2, P3 and N4, are also well explored through EEG methods over several decades.

To begin with, the P1/N1 are detected at various brain sites, both posterior and anterior. Prefrontal activity, for example, is observed in various paradigms, at about 130–150 ms post-stimulus. In some paradigms, it is observed before posterior activations.

The distribution of the N2/P2 is also now becoming clearer, especially with the use of MEG to localize the source distribution (Hopf et al., 2000); the N2pc was found to have at least two components: the first was at 180–200 ms, and sited in the SPL, the second was in the temporal lobe, at about 220–240. The former signal was proposed (Hopf et al., 2000) as arising from a control signal for attention movement (from the IMC of Fig. 4). The second was regarded as due to the filtering process needed to remove distracters from the input in a visual search task.

It is now well accepted that the P3 arises from access to appropriate buffer working memory sites in parietal lobes, observed by numerous brain imaging methods. Thus, access to the WM Sensory site is expected to occur by about 450–600 ms after stimulus onset.

Finally, the N4 is considered to be due to activation of semantic-level representations of the input. This is now regarded as outside attention, as we will discuss in the next sub-section concerned with the AB. It is thought that the N4 arises by automatic access; the P3 is slightly later, and is now accepted as the period when consciousness of the input arises.

The ERPs, in general, have sources distributed in various places in cortex; any attempt to pin each component ERP down to a single site in the brain is inappropriate. However, the modules associated with the CODAM model can themselves be networks of brain areas. Thus, we propose that the ERPs are explicable in the CODAM model as arising from sequential activation of the various CODAM control modules, as information processing stages.

- P1/N1: goal and sensory cortex.
- P2/N2: IMC, WMcd and MON.
- N4: object module.
- P3: WM sensory.

These identifications are reasonably well supported except for the association of the WMcd and monitor modules with the P2/N2 periods. To explore that further we turn to the AB, a psychological paradigm that probes the order of temporal processing in attention with considerable sensitivity.

8.3. The attentional blink (AB)

In order to test the CODAM model, the AB is very appropriate (Duncan et al., 1994; Shapiro et al., 1997). The AB leads to the notion of ‘attention dwell time’, which I suggest to arise from inhibition from the WM Sensory activity,

brought about by the previous input, onto the WMcd (and possibly IMC), during the attempt of an input to access conscious awareness. The paradigm itself involves rapid serial visual presentation (RSVP), in which various visual images (such as letters, digits, words or pictures) are presented at a fast rate (at about 10 per second). The main task is to recognize or identify two targets, termed T1 and T2. The second target can be present in one of several lagged positions with respect to the first target. When there is a gap of about 250–300 ms between T1 and T2, there is a maximum level of error in reporting T2 at a later time. An important feature about the AB is the need for masking of each of the inputs; without such masking the AB effect is considerably reduced, if it does not disappear completely.

Various models of the AB have been suggested:

- The interference model, in which interference of T1 on T2, arising possibly through a competitive process, is supposed to be the main cause of the AB.
- The response competition model, in which there is competition between the representations for T1 and T2 when response has to be made.
- The two-stage model, in which the blink represents a period during which the T2 cannot achieve a reportable representation due to the available attention resource being used by T1.

There is also a hybrid model that combines the two-stage and the interference models, and which appears to be more flexible than either separately. Furthermore, recent experimental evidence shows that removal of the mask for T2 reduces the AB effect, so that simple detection was unaffected, although more complex identification was still not possible. This is important in relating the AB to serial search. This is because, in the latter, search times are only of the order of 50 ms per item, to be compared with a AB attention dwell time of up to 500 ms. This gap is now seen to have been opened up due to the difficulty, caused by its masking, of the T2 task. Both types of task (AB and serial search) may, thus, be very similar, and two distinct forms of attention would not need to be introduced.

The control framework for attention (Fig. 7) can be seen to be an extension of the two-stage model (stage 1 being the object module, stage 2 the WMsens module) that also has competition between activations at various levels (on IMC, on WMcd and on other modules, as well as between various of the modules, such as WMsens and WMcd). It is, thus, a version of the hybrid model, extended by the addition of the modules: goal, monitor, IMC and WMcd. Let us consider how these are relevant to the AB, and various experimental variations:

- The goal module is needed to specify T1 and T2, and the detailed nature of the AB task, such as detection or identification.
- The IMC/WMcd are needed to create a feedback signal to amplify the sensory object representation ac-

tivation, to achieve access to its representation on WMsens.

- The monitor module is needed to ensure accurate detection of T1, and then T2.
- The input buffer WMsens, whose activation is needed to attain awareness of T1.

Certain of these modules appear to be at the core of the AB effect, as running the interference process that prevents access by T2 to its representation on the WMsens module. To understand such competition better, let us take note of experimental data that indicates that the N4 of semantic access of T2 to its object module representation is not affected by the AB (Vogel et al., 1998). Thus, there is not expected to be interference from T1 representations onto the object module.

Nor would delayed timing be appropriate to explain the 250 ms maximum blink (Fell et al., 2002). A crucial aspect of processing to achieve the AB is that of awareness of T1. This is thought to occur through the P3 for T1, thus, at about 450 ms after T1 onset. Given that the maximum AB is about 250 ms after T1 onset, we conclude that the processing of T2 will be at about 200–250 ms after its onset, when maximum interference from T1 occurs. This is at about the P2 of T2 (not the P3 of T2) so at about the time that the IMC and WMcd system of T2 are being activated. Since the complexity of T2 is crucial here, and not just its presence or absence, we conclude that it is most likely interference affects from the P3 of T1 to the WMcd of T2 that cause the AB. This interference will depend on the level of complexity of T2, as well as being at the correct time. It may involve binding effects, as suggested by Fell et al. (2002), but these will need to be occurring as part of the control circuitry for T2 and not the access to its WMsens representation. This latter will be somewhat later, as corresponds to the P3 of T2 (if it occurs). Thus, the hybrid interference/2-stage model does not seem to be able to produce the correct timing for the AB, on the basis of ERP evidence, without introduction of the further internal control structure of Fig. 7. It is still necessary, in order to build an effective simulation of the AB, to model the slow growth of posterior activity, under attention, as allows the object activity to build up and access its WMsens representation. Such slow build-up has been reported in TE and in V1, V2, V4 (Mehta et al., 2000); it also is seen in the slow growth of P3 (Vogel et al., 1998).

In conclusion, we see that the AB requires much of the detailed control structure of the CODAM model, especially the WMcd. To probe this further, it is necessary to build a detailed simulation of the AB using the present ideas; that has now been achieved (Kockelkoren et al., 2003), with the WMcd functioning so as to hold the CD signal from the IMC long enough, in the difficult AB conditions (especially of masking of T2) to give the input suitable amplification to attain the WMsens. We note that the recent claim (Dehaene et al., 2003), that AB can be explained by a recurrent model of the cortico-thalamic system, neglects to show,

among other things, the importance of masks (they were not included in the simulation). Nor does it produce the correct timing in association with the P2/N2 or lag of the V1/V2/V4 re-activations associated with attention feedback; nor is this feedback given by any endogenous goal bias, as should be the case to detect T1 and T2 in the RSVP process.

8.4. Conclusions

Further tests are needed of the CODAM model, especially of the temporal dynamics it leads to from the efferent copy signal of the attention movement. This signal is expected to arise at its buffer within 10 or 20 ms of the activation of the attention movement control signal, thus, at about 150–250 ms after stimulus input. Evidence of the N2/P2, as a complex set of signals produced at about 200–250 ms post-stimulus, is now being searched for in a GO/NOGO paradigm by MEG (Ioannides and Taylor, 2003). Similar temporal dynamics should be clear in the AB, as discussed above. This and further paradigms need to be explored with MEG to track down the CODAM signals more fully, so as to validate the model.

9. Conclusions

9.1. Summary

This review has developed an approach to the movement of attention by regarding it as an action, and using the engineering control approaches for the latter to apply to the former. The resulting control framework automatically brings with it various components of the overall control system, especially inverse and forward models. There are also monitor and goal modules. At the same time, attention must be paid to features of brain processing known from psychological paradigms, such as the existence of suitable buffers for holding activity in parietal sites, as well as those for transforming it in frontal sites. The resulting buffers were incorporated into the observer and goal modules in an appropriate manner.

Next, we took account that there are found to be 2 sorts of attention: (1) sensory and (2) motor.

These were observed as lateralized to the right hemisphere (sensory) and the left (motor). We then described a range of simulations, based on the engineering control approach, which enabled some quantitative justification to be given of the presence of various of the modules introduced: the IMC, the goals module, the monitor module.

Attention is increasingly accepted as being necessary for consciousness; however, it is not sufficient. The engineering control approach to the movement of attention allowed a further exploration of what further components would be relevant to begin to understand the creation of consciousness in the brain. To embark on that ambitious trip, it was noted that consciousness had several important features that were only recently being incorporated into modern research:

- The ‘gappy’/discontinuous nature of conscious experience.
- The presence of an ‘inner’ or pre-reflective self (PRS) as a separate component of consciousness.
- There are also altered states of consciousness brought about by intense meditation, such as the pure conscious experience (PCE).

It is these features of consciousness that have been used recently to help guide the extension of the engineering control model of attention towards an explanation of consciousness itself. This was through the CODAM model, in which the plant working memory site was divided into two buffers: one for the efferent copy of the attention movement signal, and the other that to receive the amplified sensory input. An explanation was given of how the pre-reflective self (PRS) can be regarded as an experience of the ‘ownership’ of the about-to-be-experienced input of the attended external event, brought about by the efferent copy of the attention movement control signal being sent to an appropriate buffer (possibly in parietal lobe). Thus, the gappy nature of conscious experience arises from the temporal flow of neural activity:

IMC → WMcd → WMsens → inhibit WMcd,

where the amplified input arrives with enough strength, to gain access to the correct site on WMsens, and then inhibits the WMcd, leading to content-full experience. The ‘gap’ in experience is, thus, that of the activation of the WMcd; there is no loss of conscious experience, per se, only a change of its nature from the ‘blankness’ of the ownership experience, in the PRS, to the content-fullness of the later experience.

The mystical state of PCE was then explained as being achieved, through meditation learning, by inhibition of sensory input so that attention is only attending to itself. This means that in PCE activity on the WMcd is temporally extended, so leading to ‘stillness’ as an extended content-free experience. Thus, PCE is the ‘inner eye’ of the WMcd activity ‘looking at itself’ by recurrence.

We have explained the PRS as arising from the efferent copy of the attention movement signal. Is this a forward model (a predictor of future input for state updating) so that the signal is truly of sensory form? If so, could it be that it is just in a different layer of the same area as the buffer for the sensory input? Or is it in a different region completely?

9.2. Unanswered questions for CODAM

We have only presented a sketch of what is a large program to explore CODAM and states of consciousness. We need to:

- Explore mental states by analogy with their known features (as done by James Clerk Maxwell to prove that light is a form of electromagnetic radiation, and especially that they had the same speed). Look at the qualitative features only (transparency, closeness, temporality, and so on); there has been an extended discussion of this elsewhere

(Taylor, 1999, 2001a), but it now needs to be completed using CODAM.

- Test CODAM by comparison with the observed temporal flow of information in the brain, for a variety of paradigms. This has already been discussed in a preliminary manner for the AB and extinction paradigms in Section 8; there are many more paradigms of value (GO/NOGO, various working memory paradigms, such as the n-back, extinction, neglect) that deserve careful study.
- Explore the temporal and topographic brain activations of subjects entering and leaving the PCE, and in particular consider the prefrontal and parietal dynamics.
- Explore the range of mental states of normal subjects (awake versus REM sleep versus SWS sleep). There are increasing numbers of experimental results on brain patterns of activity in these different states; these need close analysis from the CODAM point of view.
- Explore the mental states of disturbed patients, especially for those with disturbances of self (schizophrenia, AD, ADHD, OCD). There is already a large literature on the approach to schizophrenia as a breakdown of ipseity (Sass, 1999; Parnas, 1999, 2000). CODAM needs to be developed more closely in relation to modified brain activations in such people.
- Explore the underpinning modifications of neuro-modulators in the various paradigms discussed on attention. The manner acetylcholine, noradrenaline, serotonin and dopamine function is highly relevant to understanding the molecular basis of attention and of CODAM, and hence of consciousness.
- Construct increasingly detailed simulations of the various regions of the brain observed active in the paradigms, using the CODAM model to guide the simulations along the lines of those reported in Sections 4.3 and 5.2. This will produce simulations of increasingly many regions of the brain (including the need to incorporate various sub-cortical regions: basal ganglia, cerebellum, NRT, amygdala, etc). Such a program has already been started as part of the Lobal technologies attack on language processing through the LAD Brain Project (Taylor et al., 2003; Taylor and Taylor, 2003). Further results will be reported on this approach elsewhere.

The above bullet-points cover an enormous range of brain processing as observed experimentally and analyzed by simulation. It will be only by this joint experimental/simulation approach that we will finally come to understand that presently most scientifically elusive of all features of the Universe: human consciousness.

Acknowledgements

I would like to thank my colleagues N. Fragopanagos, S. Kockelkoren, S. Kasperides, M. Hartley, N. Taylor, at King's College London and Lobal Technologies, for constant stim-

ulation and support for these developing ideas on attention and consciousness, and developing ever-more complex brain simulations. I would also like to thank my colleagues from IME Juelich: N. Schmitz, J. Shah, H. Mueller-Gaertner and F. Binkofski, as well as A. Ioannides and L. Chan, at the Brain Sciences Institute, Tokyo, for my involvement in relevant brain imaging experiments. Thanks also go to the referees for numerous helpful comments and questions. Finally, I would like to thank the EC, through the ORESTEIA and ERMIS projects, for providing me with suitably challenging application arenas.

References

- Andersen, R.A., 2002. Private communication.
- Ariff, G., Donchin, O., Nanyakhara, T., Shadmehr, R., 2002. A real-time state predictor in motor control study of saccadic eye movements during unseen reaching movements. *J. Neurosci.* 22, 7721–7729.
- Austin, J., 1998. *Zen and the Brain*. Cambridge University Press, Cambridge, MA.
- Baerentsen, K.B., et al., 2001. Onset of meditation explored with fMRI. *NeuroImage* 13, S927.
- Berti, A., Ladavas, E., Della Corte, M., 1996. Anosognosia for hemiplegia, neglect dyslexia, and drawing neglect: clinical findings and theoretical considerations. *J. Int. Neuropsychol. Soc.* 2, 426–440.
- Binkofski, F., Fink, G.R., Geyer, S., Buccino, G., Gruber, O., Shah, N.J., Taylor, J.G., Seitz, R.J., Zilles, K., Freund, H.-J., 2002. Neural activity in human motor cortex areas 4a and 4p is modulated differentially by attention to action. *J. Neurophysiol.* 88, 514–519.
- Bisiach, E., 1988. The (haunted) brain. In: Marcel, A.J., Bisiach, E. (Eds.), *Consciousness in Contemporary Science*. Clarendon Press, Oxford, pp. 101–120 (Chapter 5).
- Botvinik, M.M., Braver, T.S., Barch, D.M., Carter, C.S., Cohen, J.D., 2001. Conflict monitoring and cognitive selection for control. *Psychol. Rev.* 108, 624–652.
- Blakemore, S.J., Wolpert, D.M., Frith, C.D., 1998. Central cancellation of self-produced tickle sensations. *Nat. Neurosci.* 1, 635–640.
- Buechel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI. *Cereb. Cortex* 7, 768–778.
- Chalmers, D., 1996. *The Conscious Mind: Towards a Fundamental Theory*. Oxford University Press, Oxford.
- Cole, J., Paillard, J., 1995. Living without touch and peripheral information about body position and movement: studies of deafferented subjects. In: Bermudez, J.L., Marcel, A., Eilam, N. (Eds.), *The Body and Self*. MIT Press, Cambridge, MA, pp. 245–266.
- Corbetta, M., 1998. Fronto-parietal cortical networks for directing attention and the eye to visual locations: identical, independent, or overlapping neural systems? *Proc. Natl. Acad. Sci. U.S.A.* 95, 831–838.
- Corbetta, M., Shulman, G.L., 2002. Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* 3, 201–215.
- Corchs, S., Deco, G., 2001. A neuro-dynamical model for the selective visual attention using oscillators. *Neural Netw.* 14, 981–990.
- Cotterill, R.J., 1989. *No Ghost in the Machine*. Heinemann, London.
- Cotterill, R.M.J., 2001. Cooperation of the basal ganglia, cerebellum, sensory cerebrum and hippocampus: possible implications for cognition, consciousness, intelligence and creativity. *Prog. Neurobiol.* 64, 1–33.
- Coull, J.T., 1998. Neural correlates of attention and arousal: insights from electrophysiology. *Prog. Neurobiol.* 55, 343–361.
- Coull, J.T., Nobre, A.C., 1998. Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed to both PET and fMRI. *J. Neurosci.* 2, 7426–7735.

- Crick, F.H.C., 1994. *The Astonishing Hypothesis*. Simon & Schuster, London.
- Crick, F.H.C., Koch, C., 1990. Towards a neurobiological theory of consciousness. *Semin. Neurosci.* 2, 263–275.
- Crick, F.H.C., Koch, C., 1998. Consciousness and neuroscience. *Cereb. Cortex* 8, 97–107.
- Crick, F.H.C., Koch, C., 2003. A framework for consciousness. *Nat. Neurosci.* 6, 119–126.
- Damasio, A., 1994. *Descartes Error*. Picador, London.
- Deco, G., 1999. Biased competition mechanisms for visual attention in a multimodal neurodynamics system. In: Wermter, S., Austin, J., Willshaw, D. (Eds.), *Emergent Neural Computational Architectures Based on Neuroscience*. Springer, Berlin, pp. 114–126.
- Dehaene, S., Sergent, C., Changeux, J.-P., 2003. A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proc. Natl. Acad. Sci. U.S.A.* 100, 8520–8525.
- Dennett, D., 1991. *Consciousness Explained*. Allen Lane, London.
- Desimone, R., Duncan, J., 1995. Neural mechanics of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- Desmurget, M., Grafton, S., 2000. Forward modeling allows feedback control for fast reaching movements. *Trends Cogn. Sci.* 4, 423–431.
- Downar, J., Crawley, A.P., Mikulis, D.J., Davis, K., 2000. A multimodal cortical network for the detection of changes in the environment. *Nat. Neurosci.* 3, 277–283.
- Doya, K., et al., 2002. Computational models of neuromodulation. *Neural Netw.* 15 (Special Issue 7).
- Duggan, A., Frith, C., 2000. Workshop on Abnormalities of Consciousness with Special Reference to Schizophrenia. Towards a Science of Consciousness Conference, Tucson IV, April 2000.
- Duhamel, J.-R., Colby, C.L., Goldberg, M.E., 1992. The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255, 90–92.
- Duncan, J., Ward, R., Shapiro, K., 1994. Direct measurement of attentional dwell time in human vision. *Nature* 369, 313–315.
- Edelman, G., 1992. *Bright Air, Brilliant Fire*. Basic Books, New York.
- Eliassen, J.C., Souza, T., Sanes, J.N., 2001. Human brain activation accompanying explicitly directed movement sequence learning. *Exp. Brain Res.* 141, 269–280.
- Feinberg, I., 1978. Efference copy and corollary discharge: implications for thinking and its disorders. *Scizophrenia Bull.* 4, 636–640.
- Fell, J., Klaver, P., Elger, C.E., Fernandez, G., 2002. Suppression of EEG gamma activity may cause the attentional blink. *Conscious. Cogn.* 11, 114–122.
- Forman, R.K.C., 1990. *The Problem of Pure Consciousness*. Oxford University Press, New York.
- Fourneret, P., Jeannerod, M., 1998. Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia* 36, 1133–1140.
- Franklin, G.F., Powell, J.D., Workman, M.L., 1989. *Digital Control of Dynamic Systems*. Addison-Wesley, Reading, MA.
- Frith, C., 1992. *The Cognitive Neuropsychology of Schizophrenia*. Lawrence Erlbaum, Hillsdale, NJ.
- Gallagher, S., 1999. Self-reference and schizophrenia: a cognitive model of immunity to error through misidentification. In: Zahavi, D. (Ed.), *Exploring the Self*. Benjamin, Amsterdam.
- Gallagher, S., 2000. Philosophical conceptions of the self: implications for cognitive science. *Trends Cogn. Sci.* 4, 14–21.
- Gallagher, S., Shear, J. (Eds.), 1999. *Models of the Self*. Imprint Academic, Thorverton, UK.
- Giteleman, D.R., Nobre, A.C., Parrish, T.B., LaBar, K.S., Kim, Y.-H., Meyer, J.R., Mesulam, M.-M., 1999. A large-scale distributed network for covert spatial attention. *Brain* 122, 1093–1106.
- Hayward, J., 1999. A rDzogs-chen Buddhist interpretation of the sense of self. In: Gallagher, S., Shear, J. (Eds.), *Models of the Self*. Imprint Academic, Thorverton, UK, pp. 379–395.
- Heilman, K.M., 1991. Anosognosia: possible neuropsychological mechanisms. In: Prigatano, G.P., Schachter, D.L. (Eds.), *Awareness of Deficit After Brain Injury*. Oxford University Press, Oxford, pp. 53–61 (Chapter 4).
- Henry, M., 1963. *L'essence de la manifestation*. PUF, Paris.
- Herzog, H., Lele, V.R., Kuwert, T., Langen, K.-J., Kops, E.R., Feiendengen, L.E., 1990. Changed pattern of regional glucose metabolism during Yoga meditative relaxation. *Neurosychobiology* 23, 182–187.
- Hillyard, S.A., Anllo-Vento, L., 1998. Event-related brain potentials in the study of visual selective attention. *Proc. Natl. Acad. Sci. U.S.A.* 95, 781–787.
- Honey, G.D., Bullmore, E.T., Sharma, T., 2000. Prolonged reaction time to a verbal working memory task predicts increased power of posterior parietal cortical activation. *NeuroImage* 12, 495–503.
- Honey, G.D., Fu, C.H.Y., Kim, J., Brammer, M.J., Croudace, T.J., Suckling, J., Pich, E.M., Williams, S.C.R., Bullmore, E.T., 2002. Effects of verbal working memory load on corticocortical connectivity modelled by analysis of functional magnetic resonance imaging data. *NeuroImage* 17, 573–582.
- Hopf, J.-M., Luck, S.J., Girelli, M., Hagner, T., Mangun, G.R., Scheich, H., Heinze, H.-J., 2000. Neural sources of focussed attention in visual search. *Cereb. Cortex* 10, 1233–1241.
- Hopfinger, J.B., Buonocore, M.H., Mangun, G.R., 2000. The neural mechanisms of top-down attentional control. *Nat. Neurosci.* 3, 284–291.
- Hopfinger, J.B., Woldorff, M.G., Fletcher, E.M., Mangun, G.R., 2001. Dissociating top-down attentional control from selective perception and action. *Neuropsychologia* 39, 1277–1291.
- Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Putz, B., Yoshioko, T., Kawatao, M., 2000. Human cerebellar activity reflecting acquired internal model of a new tool. *Nature* 403, 192–195.
- Ioannides, A.A., Taylor, J.G., 2003. Testing models of attention with MEG. In: *Proceedings of the IJCNN'03. IJCNN2003. IEEE Catalogue Number 03CH37464C, ISBN# 0-7803-7899-7*, pp. 287–297.
- Jacobs, O.L.R., 1993. *An Introduction to Control Theory*. Oxford University Press, Oxford.
- Jackson, S.R., Marocco, R., Posner, M.I., 1994. Networks of anatomical areas controlling visuospatial attention. *Neural Netw.* 7, 925–944.
- James, W., 1890. *Principles of Psychology*. Macmillan, London.
- Johansen-Berg, H., Matthehs, P.M., 2002. Attention to movement modulates activity in sensory-motor areas, including primary motor cortex. *Exp. Brain Res.* 142, 13–24.
- Johnson, H., van Beers, R.J., Haggard, P., 2002. Action and awareness in pointing tasks. *Exp. Brain Res.* 146, 451–459.
- Kastner, S., Ungerleider, L.G., 2000. Mechanisms of visual attention in the human cortex. *Annu. Rev. Neurosci.* 23, 315–341.
- Kawato, M., 1999. Internal models for motor control and trajectory processing. *Curr. Opin. Neurobiol.* 9, 718–727.
- Kentridge, R.W., Heywood, C.A., Weiskrantz, L., 1999. Attention without awareness in blindsight. *Proc. R. Soc. Lond. B* 266, 1805–1811.
- Kim, Y.-H., Gitelman, D.R., Nobre, A.C., Parrish, T.B., LaBar, K.S., Mesulam, M.-M., 2000. The large-scale neural network for spatial attention displays. *NeuroImage* 9, 269–277.
- Koch, C., Ullman, S., 1985. Shifts in selective visual attention: towards the underlying neural circuitry. *Hum. Neurobiol.* 4, 219–227.
- Kockelkoren S., Fragonanagos N., Taylor, J.G., 2003. Early Temporal Signals in the Attentional Blink, KCL, in preparation.
- Konishi, S., Nakajima, K., Uchida, I., Kameyama, M., Nakahara, K., Sekihara, K., Miyashita, Y., 1998. Transient activation of inferior prefrontal cortex during cognitive set shifting. *Nat. Neurosci.* 1, 80–84.
- LaBerge, D., 1995. *Attentional Processing*. Harvard University Press, Cambridge.
- Lamme, V.A.F., 2003. Why visual attention and awareness are different. *Trends Cogn. Sci.* 7, 12–18.
- Levine, J., 1983. Materialism and qualia: the explanatory gap. *Pacific Phil. Q.* 64, 354–361.
- Libet, B., 2001. Consciousness, free action and the brain. *J. Conscious. Stud.* 8 (8), 59–65.
- Lumer, E.D., Friston, K.J., Rees, G., 1998. Neural correlates of perceptual rivalry in the human brain. *Science* 280, 1930–1934.

- Mack, Rock, 1998. Inattention blindness: perception without attention. In: Wright, R.D. (Ed.), *Visual Attention*. Oxford University Press, Oxford, pp. 55–76.
- McAdams, C.J., Maunsell, J.H.R., 1999. Effects of attention on orienting-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* 19, 431–441.
- Mehta, A.D., Ulbert, I., Schroeder, C.E., 2000. Intermodal selective attention in monkeys I & II. *Cereb. Cortex* 10, 343–358, 359–370.
- Merleau-Ponty, M., 1962. In: Smith, C. (Ed.), *The Phenomenology of Perception*. Routledge, New York.
- Mesulam, M.-M., 1981. A cortical network for directed attention and unilateral neglect. *Ann. Neurol.* 10, 309–325.
- Miall, R.C., Wolpert, D., 1996. Forward models for physiological motor control. *Neural Netw.* 9, 1265–1279.
- Mozer, M.C., Sitton, M., 1999. Computational modeling of spatial attention. In: Pashler, H. (Ed.), *Attention*. Taylor & Francis, New York, pp. 341–393 (Chapter 9).
- Nagel, T., 1974. What is it like to be a bat? *Phil. Rev.* 83, 434–450.
- Noe, A., Thompson, E., 2003. Are there neural correlates of consciousness? *J. Conscious. Stud.*, in press.
- Olshausen, B.A., Anderson, C.H., Van Essen, D.C., 1993. A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* 13, 4700–4719.
- Parnas, J., 1999. The self and intentionality in the pre-psychotic stages of schizophrenia. In: Zahavi, D. (Ed.), *Exploring the Self*. Benjamin, Amsterdam, pp. 115–147.
- Parnas, J., 2000. The self and intentionality in schizophrenia. In: Zahavi, D. (Ed.), *Exploring the Self*. Benjamin, Amsterdam.
- Paulesu, E., Frith, C.D., Frackowiak, R.S.J., 1993. The neural correlates of the verbal component of working memory. *Nature* 362, 342–345.
- Perriot-Desilligny, C., Rivaud, S., Gaynard, B., Muri, R., Vermersch, A.-I., 1995. Cortical control of saccades. *Ann. Neurol.* 37, 557–567.
- Posner, M.I., 1980. Orienting of attention. *Q. J. Exp. Psychol.* 32, 3–25.
- Rensink, R.A., O'Regan, J.K., Clark, J.J., 1997. To see or not to see: the need for attention to perceive changes in scene. *Psychol. Sci.* 8, 368–373.
- Reynolds, J.H., Chelazzi, L., Desimone, R., 1999. Competitive mechanisms subserve attention in Macaque areas V2 and V4. *J. Neurosci.* 19, 1736–1753.
- Rizzolatti, G., 1983. Mechanisms of selective attention in mammals. In: Ewert, J.P., Capranica, R.R., Ingle, D.J. (Eds.), *Advances in Vertebrate Neuroethology*. Plenum Press, New York.
- Rizzolatti, G., Riggio, L., Sheliga, B.M., 1994. Space and selective attention. In: Umiltà, C., Moscovitch, M. (Eds.), *Attention and Performance XV*. MIT Press, Cambridge, MA, pp. 231–265 (Chapter 9).
- Robinson, D.L., Petersen, S.E., 1992. The pulvinar and visual salience. *Trends Neurosci.* 15, 127–132.
- Rolke, B., Heil, M., Hennighausen, E., 2001. Missed prime words within the attentional blink evoke an N400 semantic priming effect. *Psychophysiology* 38, 165–174.
- Rushworth, M.F., Nixon, P.D., Renowden, S., Wade, D.T., Passingham, R.E., 1997. The left parietal cortex and motor attention. *Neuropsychologia* 35, 1261–1273.
- Rushworth, M.F., Krams, M., Passingham, R.E., 2001a. The attentional role of the left parietal cortex: the distinct lateralization and localization of motor attention in the human brain. *J. Cogn. Neurosci.* 13, 698–710.
- Rushworth, M.F., Ellison, A., Walsh, V., 2001b. Complimentary localization and lateralisation of orienting and motor attention. *Nat. Neurosci.* 4, 656–661.
- Ryle, G., 1952. *The Concept of Mind*. Hutchinson, London.
- Sabes, 2000. The planning and control of reaching movements. *Curr. Opin. Neurobiol.* 10, 740–746.
- Sakai, K., Hikosaka, O., Takino, R., Miyauchi, S., Nielsen, M., Tamada, J., 2000. What and when: parallel and convergent processing in motor control. *J. Neurosci.* 20, 2691–2700.
- Sartre, J.-P., 2001. In: Barnes, H.E. (Ed.), *Being and Nothingness*. Routledge, London.
- Sass, L.A., 1999. Schizophrenia, self-experience and the so-called 'negative symptoms'. In: Zahavi, D. (Ed.), *Exploring the Self*. Benjamin, Amsterdam.
- Schluter, N.D., Krams, M., Rushworth, M.F., Passingham, R.E., 2001. Cerebral dominance for action in the human brain: the selection of actions. *Neuropsychologia* 39, 103–113.
- Schweighofer, N., Spoeckstra, J., Arbib, M., Kawato, M., 1998. Role of the cerebellum in reaching movements in humans. II. A neural model of the intermediate cerebellum. *Eur. J. Neurosci.* 10, 95–105.
- Shadmehr, R., Brashers-Krug, T., 1997. Functional stages in the formation of human long-term motor memory. *J. Neurosci.* 17, 409–419.
- Shapiro, K.L., Arnell, K.M., Raymond, J.E., 1997. The attentional blink. *Trends Cogn. Sci.* 8, 291–296.
- Shibata, T., Ioannides, A.A., 2001. Contribution of the human superior parietal lobe to spatial selection process: an MEG study. *Brain Res.* 897, 164–168.
- Shoemaker, S., 1984. *Identity, Cause and Mind*. Cambridge University Press, Cambridge, MA.
- Sirigu, A., et al., 1996. The mental representation of hand movements after parietal cortex damage. *Science* 273, 1564–1568.
- Snyder, L.H., Batista, A.P., Andersen, R.A., 2000. Intention-related activity in the posterior parietal cortex. *Vision Brain Res.* 40, 1433–1441.
- Sperry, R.W., 1950. Neural basis of the spontaneous optokinetic response produced by visual inversion. *J. Compar. Physiol. Psychol.* 43, 482–489.
- Strawson, G., 1999. The self. In: Gallagher, S., Shear, J. (Eds.), *The Self*. Imprint Academic, Thorverton, Devon, pp. 1–24.
- Taylor, J.G., 1999. *The Race for Consciousness*. MIT Press, Cambridge, MA.
- Taylor, J.G., 2000a. A general framework for functions of the brain. In: *Proceedings of the International Joint Conference on Neural Networks 2000*. IEEE Press, ISBN 0-7695-0619-4/00.
- Taylor, J.G., 2000b. The central representation: the where, what & how of consciousness. In: *Proceedings of the Emergence of Mind*, Milan, April 2000. Fondazione Carlo Erba, Milan.
- Taylor, J.G., 2000c. Attentional movement: the control basis for consciousness. *Neurosci. Abstr.* 26 (Part 2, 839.3), 2231.
- Taylor, J.G., 2001a. The importance of the parietal lobe for consciousness. *Conscious. Cogn.* 10, 379–417.
- Taylor, J.G., 2001b. Attention as a neural control system. In: *Proceedings of the International Joint Conference on Neural Networks (IJCNN'01)*. IEE Cat. #01CH37222C, pp. 272–276, ISBN 0-7803-7046-5.
- Taylor, J.G., 2002a. Paying attention to consciousness. *Trends Cogn. Sci.* 6, 206–210.
- Taylor, J.G., 2002b. From matter to mind. *J. Conscious. Stud.* 9, 3–22.
- Taylor, J.G., 2002c. Toward global principles of brain processing. In: Hecht-Nielsen, R., McKenna, T. (Eds.), *Computational Models for Neuroscience*. Springer, London, pp. 221–244 (Chapter 8).
- Taylor, J.G., Rogers, M., 2002. A control model of the movement of attention. *Neural Netw.* 15, 309–326.
- Taylor, J.G., Fragopanagos, N., 2003. Simulations of attention control models in sensory and motor paradigms. In: *Proceedings of the IJCNN'03*. IJCNN2003, IEEE Catalogue Number 03CH37464C, ISBN# 0-7803-7899-7, pp. 298–303.
- Taylor, J.G., Kasderides, S., 2003. Living with the dangerous. In: Carbonell, J.G., Sickmann, J. (Eds.), *Proceedings of the KES'03 Conference*, Oxford. Springer, London, pp. 615–621.
- Taylor, N., Taylor, J.G., 2003. *Neural Networks/Biological Cybernetics/ACTION* net web page at <http://www.kcl.ac.uk>.
- Taylor, J.G., Schmitz, N., Ziemons, K., Grosse-Ruyken, M.-L., Gruber, O., Mueller-Gaertner, H.-W., Shah, N.J., 2000. The network of brain areas involved in the motion aftereffect. *NeuroImage* 11, 257–270.
- Taylor, J.G., Taylor, N., Hartley, M., 2003. *The LAD Brain Program*. Lobal Technologies Internal Report.
- Treisman, A., 1988. Features and objects. The fourteenth Bartlett memorial lecture. *Q. J. Exp. Psychol.* 40A, 201–237.

- Treue, S., Maunsell, J.H.H., 1999. Effects of attention processing in macaque middle temporal and medial superior temporal visual cortical areas. *J. Neurosci.* 19, 7591–7603.
- Umeno, M.M., Goldberg, M.E., 1997. Spatial processing in the monkey frontal eye field. I. Predictive visual responses. *J. Neurophysiol.* 78, 1373–1383.
- van Veen, V., Carter, C.S., 2002. The timing of action-monitoring processes in the anterior cingulate cortex. *J. Cogn. Sci.* 14, 591–602.
- Vidyasagar, T.R., 1999. A neuronal model of attentional spotlight: parietal guiding the temporal. *Brain Res. Rev.* 30, 66–76.
- Vogel, E.K., Luck, S.J., Shapiro, K., 1998. Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *J. Exp. Psychol.* 241, 656–1674.
- von Holst, E., Mittelstaedt, H., 1950. Das Reafferenzprinzip: Wechselwirkungen zwischen Zentralnervensystem und Peripherie. *Naturwissenschaften* 37, 464–476.
- Vuilleumeir, P.O., Rafal, R.D., 2000. A systematic study of visual extinction. *Brain* 123, 1263–1279.
- Vuilleumeir, P.O., Armony, J.L., Clarke, K., Husain, M., Driver, J., Dolan, R.J., 2002. Neural responses to emotional faces with and without awareness: event-related fMRI in a parietal patient with visual extinction and spatial neglect.
- Willingham, D.B., 1998. A neuropsychological theory of motor skill learning. *Psychol. Rev.* 105, 558–584.
- Wolfe, J.M., 1997. In a blink of the mind's eye. *Nature* 387, 736–737.
- Wolpert, D.M., Gharharmini, Z., 2000. Computational principles of movement neuroscience. *Nat. Neurosci. Suppl.* 3, 1212–1217.
- Wright, R.D., Ward, L.M., 1998. The control of visual attention. In: Wright, R.D. (Ed.), *Visual Attention*. Oxford University Press, Oxford, pp. 132–186 (Chapter 6).
- Zahavi, D., 1999. *Self-Awareness and Alterity*. Northwestern University Press, Evanston, IL.