

Vittorio Gallese

The Inner Sense of Action

Agency and Motor Representations

I: Introduction

We live in a meaningful world. Our capacity to deal with the ‘external world’ is constituted by the possibility of modifying the world by means of our actions; by the possibility of *representing* the world as an objective reality; and by the possibility of experiencing phenomenally this same objective reality, from a situated, self-conscious perspective. It is tempting to address these different articulations of the sense of ‘being related to the world’, of our *intentional relation* to the world, by using different languages, different methods of investigations, perhaps even different ontologies.

In the present paper I will start to explore the possibility of reconciling some of these different articulations of intentionality from a neurobiological perspective. I will confine my analysis to the relationship between agency and representation and I will show how representation is intrinsically related to action control. To that purpose, I will present a new account of action, arguing against what is still commonly held as its proper definition, namely the final outcome of a cascade-like process that starts from the analysis of sensory data, incorporates the result of decision processes, and ends up with responses (actions) to externally- or internally-generated stimuli. I will argue against this account of action by presenting and discussing recent findings from the investigation of neural mechanisms that are at the basis of sensorimotor integration. It will become clear that the so-called ‘motor functions’ of the nervous system not only provide the means to control and execute action but also to *represent* it. Actually, following this view, action control and action representation become two sides of the same coin.

One of the advantages of this approach is its empirical testability. If the results of this analysis turn out to be correct, we have, as philosophical pay-off, the possibility of overcoming the dichotomy between functional and semantic properties, by showing that causal properties *are* content properties. The neurophysiological

Correspondence: Vittorio Gallese, Istituto di Fisiologia Umana, Università di Parma, Via Volturno 39, I-43100 Parma, Italy.

and neuropsychological data that will be presented in the following sections provide preliminary, but very encouraging, evidence pointing in this direction. In sum, the overall goal of the present paper will be to convince the reader that — using the results of neuroscientific investigation as tools — representation can be naturalized.

II: Hierarchies: From Neurons to Behaviour

Neuroscience, being part of the biological sciences, should share with them an evolutionary and monist¹ foundational perspective. There is only one reality, organized — and therefore describable — along multiple levels of complexity. Each of these different levels of organization can be decomposed into other levels that are simpler, or at the very least endowed with different functional characterizations. The pay-off of this deconstructivist approach, when guiding empirical research, is that we can more easily trace an evolutionary path connecting all levels. Filling in the evolutionary mosaic of brain functions increases the heuristic power of our strategy of enquiry. Each of these different organizational levels requires an appropriate descriptive language. Which is therefore the appropriate level of enquiry of neuroscience?²

It is commonly held that we can devise a hierarchical model of the brain/mind in which we distinguish levels characterized by increasing complexity. Each of these levels of description (for example, ions, ion channel–receptor, neuron–synapse, neural assemblies–cortical circuits, brain–behaviour, cognition–consciousness) presupposes the possibility to delimit a *system*, characterized by principles of organization and functional mechanisms that can be defined as proper to that system. Each level of this hierarchy contains elements that we find also at the next ‘higher’ level, but framed within a more complex architectural and functional design. Our hierarchy is supposed to be topped by a level describing what is considered the distinctive hallmark of mankind: culture.

There is indeed an intuitive incommensurability when comparing the intrinsic features of levels of description such as neuron, behaviour, and experience. Neuroscience so far has almost exclusively explored the relation between neurons and behaviour, leaving any account of experience in the hands of philosophers. Let us focus first on the first of these levels, that of the neuron.

According to the classic neuroscientific model, each living organism is located within fields of energy — electromagnetic, mechanical and chemical energy. All these different forms of energy constitute the stimuli (visual, auditory, somatosensory, and the like) to which every organism is exposed. Energetic ‘stimuli’ are translated, or better, transduced into a *common code*; the action potentials of nerve cells. Action potentials express the electrochemical excitability of cells and

[1] With remarkable exceptions. See, for example, Eccles, 1989.

[2] Neuroscience comprises many different approaches addressing the functional properties of different levels of organization of the brain. In the present paper I’ll deliberately confine my discussion to the neurophysiological approach, which investigates the correlation between neural functions and behaviour.

constitute the ‘language’ spoken by the billions of neurons that make up the central nervous system. The receptors of the different sensory modalities are the agents of the transduction process: they translate the different types of energy into the common code of action potentials. This code is *truly common*; sensory pathways that, through receptors, receive information from the so-called ‘external world’, share the same code with motor pathways that, through different intermediate steps, activate neuro-muscular junctions that, in turn, *determine*³ our movements. (Following this line of argument one could object that nothing really prevents this common code from also ‘activating’ our consciousness,⁴ but we’ll come back to this point later on.)

So far, I have indicated that sensory and motor pathways share the same ‘language’ (incidentally, they also share it with the beating of our heart). This should warn us not to fall prey too easily to modularity mania, as when proposing strict dichotomies such as that between action and perception. Dichotomies have a very long historical record, and neuroscience in this respect is no exception. Today we are constantly exposed to the so popular *mantric* succession of dichotomies proposed as the state-of-the-art account of vision: where/what, how/what, pragmatic/semantic, egocentric/allocentric, and so on.

I hope to convince the reader with the arguments that will be presented in the following sections that these accounts leave many questions unanswered. From our neurobiological standpoint we then feel the urge to ask ourselves new questions, in part because we are not fully satisfied with many of the answers provided so far. A good starting point could be to explore the relationship that, at first sight, looks less problematic: that between neurons and behaviour.

III: How to Study Neurons: A Semantic Approach to Neural Computation

Each neuron elaborates and integrates excitatory and inhibitory influences that shape and define its final output. This process of integration that translates different inputs into a common output endows neurons with the properties of a ‘categorizing machine’. We can ask ourselves what are the mechanisms and the rules presiding over this particular level of integration. To this purpose we can study and model, for example, the dendritic spine as an integrative processing unit, like a little computer in the brain. In my opinion this amounts to an *analogue approach* to neural computation. But we should also ask ourselves *what kind of content* does the neuron transport? As soon as we have content, we have representation, and also the possibility of *misrepresentation*. This would then constitute a *semantic approach* to neural computation. The task of neurophysiology is to

[3] Without innervation there cannot be any movement. Thus neurons, in a sense, determine or cause movements. However, there are other ‘systems’ (muscles, blood vessels, the heart, etc.) that are concurrently necessary for movements to occur. It could be objected that we are facing a problem of causal overdetermination (see Kim, 1998 for a discussion of this point in relation to mental causation). My answer is that this problem perhaps stems from a misconception about the ontological status of these so-called ‘systems’. We should better speak of sub-components, the whole organism being the system.

[4] A caveat: the term consciousness here has a double meaning: ‘being aware *that* P’ and ‘being aware *of being aware that* P’.

search for a reliable correlation between the number of action potentials produced by neurons of a given brain area (sampled either sequentially with the single neuron recording approach, or simultaneously by means of multiple recordings) within a given time unit, and a particular behaviour or chunk of behaviour of the organism whose brain is under investigation. A critical ‘off-line’ process of validation or confutation of the significance of the recorded correlation follows data acquisition. When describing correlations between neurons and behaviour we are forced to select a foundational perspective defining the broader context in which our investigation is supposed to be framed.

My personal view of this ‘broader context’ is that brain functions can be accounted for only by considering the dynamic interplay that occurs between the biological agent *as a whole*, and the ‘external world’ (see also Jarvilehto, 1998). Any attempt to characterize brain functions as the outcome of encoding devices whose final product is a symbolic ‘language’ totally remote from the acting body is bound to fail. The interdependence between brain and ‘world’ perhaps indicates, as will become clearer in the next sections, the necessity to question an absolute defining limit between these two realms.⁵

Furthermore, a critical aspect of neuroscientific investigation resides in the fact that the ‘answers’ we are getting from neuronal activity are strongly influenced by the way in which we pose our ‘questions’. For example, we can study brain functions by using a given behavioural paradigm and systematically apply it in different parts of the brain. The brain, however, isn’t equipotential, unless for its functional code that, as we have seen before, is ubiquitous. To say ‘different parts of the brain’ means distinct cortical areas and subcortical structures that can differ along many dimensions such as their cytoarchitectural features and/or the pattern of their anatomical connections. Unfortunately, when anatomical specificity is neglected, the simple correlation between a stereotyped behavioural paradigm and the neural activity of *different* cortical areas may induce very frustrating conclusions about the heuristic value of this experimental approach, leading one to conclusions that are, at best, too precipitous, if biased by such methodological flaws.

The ‘everything-is-distributed’ approach should be contrasted with a ‘naturalistic’ approach. This approach, when applied to neurophysiology, consists in choosing the most appropriate way of testing neural activity, by figuring out what are the stimuli or the behavioural situations that more closely approximate what the organism we are studying would experience in a natural environment. A second important tenet of the approach I am advocating consists in the systematic and preliminary functional characterization of the properties of a given, neural population. Only after this preliminary goal has been achieved⁶ can the most suitable paradigm be introduced.

[5] It has been proposed by several authors that the representational content of intentional states depends on external reality (see, for example, Bilgrami, 1992; Clark, 1997; even more radical are the theses of Jarvilehto, 1998, and of O’Regan and Noë, 2000).

[6] Certainly with some risk of incompleteness.

During the last two decades, the combination of this experimental type of investigation with the increased attention paid to determine *where* in the brain neural recordings were carried out has led to a series of important discoveries. The discovery of motor properties within the once-called parietal ‘association’ areas (among others, see Mountcastle *et al.*, 1975; Hyvärinen, 1981; Andersen, 1987; Sakata *et al.*, 1995); the discovery of the neural mechanisms at the basis of the visuo-motor transformation processes for the guide of eye and hand movements (see Andersen, 1987; Andersen *et al.*, 1993; Colby and Goldberg, 1999; Jeannerod *et al.*, 1995; Rizzolatti *et al.*, 2000b); and the functional parcellation of the premotor cortex (see Matelli *et al.*, 1985; Wise, 1985; Gentilucci *et al.*, 1988; Rizzolatti *et al.*, 1988; He *et al.*, 1993; Tanji, 1994; Matelli and Luppino, 1997; Rizzolatti *et al.*, 1998) have paved the road to a new way of conceiving motor systems as *representational* systems (see below).

If acting and representing are two sides of the same coin — as I’ll try to demonstrate — a large portion of the reality we live in and represent must inevitably be conceived of as *literally* being constituted by *actual* and potential (*simulated*)⁷ bodily motions, behaviours, and actions.

Another crucial problem posed by considering the relationship between neurons and the ‘world outside’ consists in defining how perception, action and cognitive processes interrelate with each other. Historically, neuroscience has tended to privilege, on one side, the study of sensory processes — with a strong emphasis put on the study of vision — and, on the other, the study of motor processes. A big void is left in between, occupied by cognitive processes, or, to use a more comprehensive definition, by what constitutes mental events: intentionality, desires and beliefs. These issues have represented the traditional playground of philosophy and, more recently, of the so-called ‘cognitive sciences’.

According to classic cognitivist models of mental events, perception and action not only play distinct, modular, and peripheral roles, but in virtue of their elementary role as input and output interfaces, respectively, they have also been kept rigidly separated from cognitive processes, which in turn were characterized as central and holistic (see Fodor, 1983; see Hurley, 1998 for an acute critique of the cognitivist model; see also Freeman and Núñez, 1999).

Has neuroscience today anything new to say about these issues? In recent years the relation between action and perception and the relation between pragmatic control and semantics have been addressed experimentally (see Jeannerod, 1994; 1997; Milner and Goodale, 1995 for thorough and comprehensive accounts of these issues). The processing of visual information, analyzed at an elementary level within the primary visual cortex, is considered to be carried out along two main visual pathways: a dorsal one supposed to terminate in the posterior parietal

[7] Simulation routines — as contrasted to theory–theory approaches — have been proposed as the most likely candidates to account for the capacity to attribute intentions, desires and beliefs to others (Goldman 1989; 1992; 1993; 2000; Gordon, 1986; Harris, 1989. See Gallese and Goldman, 1998 for a discussion of possible neural mechanisms at the basis of simulation routines). Here I refer to simulation, in a broader context, as a crucial mechanism to make the ‘external world’ phenomenally intelligible.

lobe, and a ventral one reaching the inferior temporal cortex. Ungerleider and Mishkin (1982) have influentially proposed that the dorsal pathway should function to analyze the spatial relationships among objects, while the ventral pathway should code their identity. This model posits that vision is ‘implemented’ along two parallel routes: the *where* and *what* pathways.

From the early nineties this model has been questioned by an equally influential — and partly alternative — one (Milner and Goodale 1995; see also Gallese *et al.*, 1999 for a critical discussion of it). In Milner and Goodale’s view (1995) the dorsal pathway is involved in the sensorimotor ‘on-line’ control of action (the *where* becomes *how*), while the ventral pathway is maintained (pretty much in accord with Ungerleider and Mishkin) to be the privileged site for the semantic description of objects.

Both models, although with substantial differences, posit a strict dichotomy between regions of the brain supposed to control the *doing* of things, and other ones supposed to *know* what things really are. Apparently, even for modern neuroscience, action and perception have to be considered as distinct and segregated entities.

I think that it should be sufficiently clear, even from this very concise and perhaps oversimplified account that this interpretation of how the brain works is too simplistic.

It is no coincidence that the debate on the relation between action and perception, so richly intertwined with philosophical implications, has chosen the psychophysiology of vision as its empirical playground. In spite of the signs of snobbish detachment sometimes displayed by part of the neuroscientific community with respect to philosophical and foundational issues, the privileged attention devoted along the years to the empirical study of the psychophysiology of vision is deeply rooted in the representational and symbolic account of the mind, which has a long philosophical tradition. In fact, vision, beside having been considered as the paradigmatic model for the study of perceptual processes, has historically also provided the vast majority of metaphors to explain mental phenomena.

My point is that this strategy is profoundly wrong. Happily enough, in the present days, we see a revitalization of an old tradition of thought — obscured for nearly half-a-century, transversal to scholars of both phenomenology (Husserl, Edith Stein, Merleau-Ponty) and pragmatism (James, Dewey, Mead), that has emphasized — although with different qualifications — the *embodied* roots of mental processes. In the next sections I will address the relationship between action and perception quite differently from the tenets of classical cognitivism and neuroscience. This perspective will show the impossibility of drawing a sharp line between *acting* and *perceiving*. Furthermore, this account of sensorimotor processes will enable us to formulate some new hypotheses about how our brain is capable of *re-presenting* the world as phenomenally experienced.

IV: Actions and Their Representations

A series of neuroanatomical and neurophysiological data accumulated during the last twenty years conveys a dramatically different picture of how the brain treats sensorimotor information.

The frontal agranular cortex (constituted by the primary motor cortex, the supplementary motor area and the premotor cortex), as well as parietal cortices, are constituted by a mosaic of areas, endowed with peculiar anatomo-functional properties, which interact by means of reciprocal connections⁸ within distinct cortico-cortical networks. Each of these networks integrates sensory and motor information relative to a given body part in order to *control* that particular body part within distinct spatial reference frames. In other words, the picture emerges of a multiplicity of ‘cortical representations’ of distinct effectors, each of which is related to specific and distinct functions. When I am speaking of ‘cortical representations’ I do not refer, however, to a mere symbolic, abstract equivalence between a *real* entity in the world and a computational code, which in principle can be multiply instantiated in whatever substrate. The expression ‘the cortical representation of . . . ’ needs to be qualified according to its aboriginal meaning, that of *control*.

To make myself clearer I will use an example. Let us imagine for a while, through the power of mental concentration, trying to reduce the rate of our heart-beat by about 10 per cent. A pretty tough enterprise at least for the layman not accustomed to meditative practices. Let us imagine now the less exotic situation of being able to visualize our heart rate on a monitor plugged to an electrocardiograph that records our heartbeats. Surprisingly enough, after some practice we would realize that indeed we are capable of reducing our heart rate to the desired level. There is no magic in it; we simply use a well-known practice named bio-feedback. What does this trivial example tell us? It tells us that a very efficient way to control a given variable (the heart rate in our example) is to produce a copy, a *re-presentation* of that variable. The notion of representation needs to be freed from its abstract connotation — typical of the representational-computational account of the mind — and has to be relocated within a naturalistic perspective. This new account of representation stresses its pre-conceptual and pre-linguistic roots.

What does it precisely mean to define representation in control terms? It means to underline its relational — and therefore intentional — character. Jaegwon Kim in his beautiful *Mind in a Physical World* (1998) has written:

Thus, that a given intentional state of an organism instantiates a certain semantic property is a relational fact, a fact that essentially involves the organism’s relationship to various external environmental and historical factors. This makes semantic properties relational, or extrinsic, whereas we expect causative properties involved in behaviour production to be non-relational, or intrinsic, properties of the organism.

[8] Edelman (Edelman and Tononi, 2000) and Zeki (1993) have both recently stressed in different contexts and at different scales the importance of re-entrant signalling.

My objection to Kim is that the causative properties in a sensorimotor loop, in active vision, those involved in behaviour production, actually *also* belong to the target object. I want to question the whole conceptual distinction between extrinsic/ semantic properties and intrinsic/causative properties. I think we now have a way of describing the overall dynamics of behavioural and perceptual events in a new conceptual framework, which is not only conceptually convincing, but also empirically increasingly plausible. We can transcend the old dichotomy in a way that still generates testable predictions, but at the same time has a deeper philosophical flavour. The philosophical intuition behind my thesis is that causative properties *are* content properties, that — at least for *some* forms of mental content — their meaning is literally *constituted* by the way they are ‘enacted’ by a situated and functionally grounded organism.

This point will be clearer once we have specified the double executive/representational character of the motor system. To do that we should provide a *naturalistic* account of action.

Every action is characterized — by definition — by the presence of a goal. The same *movements* (for example, flexing the fingers of our hand) can be executed to achieve different purposes, such as grasping a mug, scratching our head or simply playing with our fingers. The achievement of different goals turns those very same movements into different *actions*. What relationship exists between the motor system, movements and actions? Until not so many years ago the motor system was conceived as a mere movement controller. However, recent neurophysiological findings convey a totally different picture: the motor system controls *actions*.

A series of experiments of single-neuron recordings in the premotor cortex of behaving monkeys carried out during the 1980s led to the discovery of neurons that became activated not during the execution of simple movements, but during the execution of motor acts aimed at the achievement of a specific purpose (Rizzolatti *et al.*, 1988; Gentilucci *et al.*, 1988). More specifically, these neurons discharged any time the monkey was grasping a given object, regardless of which was the effector employed (the right hand, the left hand, or the mouth). It is clear that the movement of each of these different effectors is controlled by different groups of muscles. Neither muscle activation, nor movements can therefore account for the activation of this class of neurons. What drives these neurons is the *goal* of the action. What does that imply? Premotor neurons, traditionally considered part of the output pathway by means of which agents respond to external or internally generated stimuli, do indeed, as low-level physical events, actually correlate with events that we can only understand if we move to the most abstract level of description of action: its purpose.

On the basis of the results reported above, in a new series of experiments carried out in our laboratory, we wanted to investigate, by means of a specifically devised behavioural paradigm, the role of the monkey ventral premotor cortex in the mechanisms of visuo-motor control of goal-related hand movements. We discovered that a particular class of premotor neurons — ‘canonical neurons’ — are selectively activated when the monkey, in the absence of any active movement,

observes objects whose intrinsic features such as size and shape are strictly related to the type of action that the very same neurons motorically code (see Murata *et al.*, 1997; Rizzolatti *et al.*, 2000; Gallese, 2000a). The most interesting feature of these neurons is that the vast majority of them show a strict congruence between their high motor selectivity for a particular type of prehension and the visual selectivity for objects that, although differing in shape (e.g. cube, cone and sphere), nevertheless require the same type of prehension in order to be grasped. It has to be stressed that these ‘visual’ responses do not precede any impending movement, because the monkey is simply required to fixate the objects.

It is hard to conceptualize these responses in purely sensory or motor terms. It is more plausible to postulate that the objects whose observation triggers the neurons’ response are analyzed in *relational* terms. Object observation, even within a behavioural context not specifically requiring an active interaction on the side of the observer, determines the activation of the motor program that *would be* required were the observer actively interacting with the object. To observe objects is therefore equivalent to automatically evoking the most suitable motor program required to interact with them. Looking at objects means to unconsciously ‘simulate’ a potential action. In other words, the *object-representation* is *transiently integrated* with the *action-simulation* (the ongoing simulation of the *potential action*).⁹

If this interpretation is correct, objects are not merely identified and recognized by virtue of their physical ‘appearance’, but in relation to the effects of the interaction with an agent. In such a context, the object acquires a meaningful value by means of its dynamic relation with the agent of this relation. This dynamic relation is multiple, as multiple are the ways in which we can interact with the world by acting within it. The *object-representation* ceases to exist by itself. The object *phenomenally exists* to the extent it represents the target of an action.

The ecological approach to perception, influentially heralded by Gibson (1979) has contributed to a great extent to corroborate a notion of the subject ever less *other* with respect to the ‘outside world’. The subject — an acting subject — is defined by her/his reciprocal dynamic relation with the world, that world whose unstable and changeable borders are unceasingly set by acting on it. With respect to Gibson, however, who assigns to active but also to passive *movement* a purely instrumental role in defining the invariant features already present in sensory data, I think we should stress the positive role of *action* in providing meaning to the overall world-model or the world *as represented*. Objects’ invariance shouldn’t be considered an intrinsic feature of the physical world, but rather the result of the peculiar interactions with the acting organism (see also Merleau-Ponty, 1962).

[9] A similar approach can be found in the work of Wolpert and Kawato (see Wolpert *et al.*, 1995; Wolpert, 1997; Kawato, 1997; 1999). These authors propose — very convincingly in my view — that motor *internal models* may constitute not only the best way to control motor learning, but may also provide insights for cognitive domains such as communication, thinking and consciousness. Also relevant is the Emulation Theory of Representation (ETR) proposed by Grush (1997). My own perspective, as presented in the present paper, owes much to the exciting contributions of all these authors.

This is not meant to deny obviously that a dynamic ‘catalogue’ of different object descriptions is located in the infero-temporal cortex. What I want to stress is that each of these descriptions becomes *fully* meaningful only once the individual, by acting (looking at the object by moving the eyes, moving around the object, interacting with the object), has had experience with the object.

The German philosopher Thomas Metzinger has introduced a technical term for this: the *phenomenal model of the intentionality relation* (PMIR) (Metzinger, 2000). The PMIR is a consciously available, dynamic representation of a *subject-object relation*, of the organism as a whole as currently interacting with a perceptual or behavioural target-object. An object-representation in the infero-temporal cortex becomes *more* meaningful for the organism, if it is integrated into (something equivalent to) Metzinger’s PMIR, because it has an expanded *proper function* for the organism in the sense of Millikan (1984; 1993). What I am proposing here is that to be phenomenally *conscious* of the meaning of a given object depends also on the *unconscious* simulation of actions directed to that object.

In humans, the development of language allows a new way of categorizing objects by means of their naming. By receiving a verbal description of an object one can infer its category without the need of acting on it. However, ‘to receive a verbal description of an object’, if one looks closer at it, could still be a way of experiencing this object, by involving the *internal* simulation of an action directed to that object. This prediction could be easily tested empirically.

In addition, it should be stressed that, beside the mechanism discussed above, it is likely that the evolutionarily most ancient systems linked to emotional life may also provide a further, and possibly even more basic, description of objects such as ‘edible’, ‘not edible’, ‘dangerous’, ‘sweet’, etc. Thus, contrary to ‘classic’ cognitivism, the data here reviewed suggest that object meaning derives from the biology of the organism and not from some *a priori* equivalence between objects in the world and internal descriptions.

Let us imagine an alien creature who, suddenly materialized on our planet, and being endowed with a physical body totally different from ours, interacted with terrestrial objects by employing its peculiar — and alien — way of acting, constrained by its peculiar anatomy. Are we sure that the alien’s description of the same objects we are dealing with in our daily life would be the same as ours? If we allow ourselves a negative answer to this question, the possibility opens up to define vision (and, translationally, also the other different sensory modalities) in a completely different way. Perceptual processes *are part* of action.

According to this view, a definition of action and perception as distinct and separated entities becomes highly problematic. Furthermore, the empirical data here briefly summarized constitute a very strong argument against the dichotomous — and too simplistic — account of visual processing we were discussing earlier on. Canonical neurons are located in a sector of the ventral premotor cortex (area F5) reciprocally connected with the posterior parietal cortex, the terminal of the dorsal visual pathway, characterized as the *where* or *how* pathway according to the models of Ungerleider and Mishkin (1982) and Milner and

Goodale (1995), respectively. According to these models, canonical neurons should therefore be the recipient of visual information fed by a pathway — the dorsal one — totally unrelated with the semantic coding of objects, considered to be a prerogative of the ventral pathway, the *what* one.

What then really constitutes the *meaning* of an observed and internally represented object? A purely pictorial description of its shape, size and colour features, or rather *also* its intentional value? The pictorial description only gains its full, interesting meaning by being transiently bound to an individual first-person perspective on the level of conscious experience, by becoming the object-component of a much bigger, comprehensive picture. A possible account of this more comprehensive picture could be Metzinger's PMIR, the ongoing representation of an agent *as currently integrated with this object*, by *intending* it, for example in terms of intending to grasp it, intending to hold it in the focus of perceptual attention, in terms of intending to *think* about it. All these different situations involve an object and a subject. They generate an ongoing internal simulation of this object and this subject within a relation of dynamic complementarity. If both representations are *integrated*, then a full-blown holistic kind of content emerges. Maybe a good metaphor for this process is: the object becomes *intentionally charged*, because it is now not represented in isolation, but under a PMIR.

The *intentional* character of representation should now be clearer. We could hypothesize that re-presentations (see also Grush, 1997), as I have qualified them, do not primarily originate — either phylogenetically, or ontogenetically — with a specific *semantic* value. This feature is likely the later result of the functional reorganization of processes originally selected for a different purpose. I submit that this purpose is to achieve a better control of the dynamic relation between an open system — the living organism — and the environment.

This scenario can be framed within the concept of 'exaptation' introduced by Gould and Lewontin (1979). With the term exaptation these authors denote the emergence of new skills (behavioural, mental and the like) that developed by exploiting in a totally new way resources that had been selected for other purposes. An alternative approach has been proposed by Dennett. In *Darwin's Dangerous Idea* (1995), Dennett defends the thesis that a theory of evolutionary adaptation can ground a theory of natural teleology that, in turn, grounds mental representations and the intentionality of the mind. This creates a bridge between the intentionality of natural selection and the intentionality of the mind. As — quite correctly, in my view — pointed out by Fodor (1998), natural teleology implies a 'Designer', therefore '*... you can't explain intentionality by appealing to the notion of design because the notion of design presupposes intentionality*' (Fodor, 1998). I share with Fodor the idea that Mother Nature has no foresight. I also agree with Fodor (1998) that '*... the representational is a (possibly large) superset of the intentional*', although, as should be clear from what has been said so far, I am afraid that I do not share his enthusiasm for *informational* semantics. The evolutionary account I sketched above bypasses the question-begging arguments of Dennett and at the same time provides, I hope, a good starting point for a

naturalistic theory of representation. If causative properties are content properties representation can be naturalized.

The sensorimotor integration processes supported by the multiple and parallel fronto-parietal cortical networks instantiate ‘internal copies’, duplicates of actions utilized to generate and control goal-related behaviours (see footnote 9), *but also* to provide — at a pre-conceptual and pre-linguistic level — a meaningful account of things in the world. Objects acquire their full meaning only to the extent that they constitute one of the poles of the dyadic dynamic relation with the acting subject, who, in turn, constitutes the second pole of this relationship.

Our vision of the world can be characterized as dynamic, relational and intentional. As Merleau-Ponty writes in *The Phenomenology of Perception* (1962), ‘*The identity of the thing through perceptual experience is only another aspect of the identity of one’s body throughout exploratory movements; thus they are the same in kind as each other*’. I posit that this dynamic account of vision depends on agency. In this respect vision shares with other *dynamic* senses (for example, touch) the same logic of operation.

V: Neurons, Behaviour and Experience

We started our analysis by exploring the relationship between neurons and behaviour. This analysis, by assigning a pivotal role to agency, has brought us to the conclusion that by using a naturalistic approach we can provide new answers to the problems raised by this relationship that, to a careful scrutiny, appeared more complex than one should have expected. A second important result of this analysis consists in the possibility of empirically addressing what also looked like *the most problematic* relationship: that between neurons and experience. When we speak of subjective experience we refer to a set of properties that we assign to an entity that we call *the self*.

The self can be conceived as a *theoretical construct* (see Baars, 1997; Metzinger, 1993; 2000) of the phenomenological experience of being a well-identified ‘I’. This construct is a model — the Phenomenal Self Model (PSM) — (see Metzinger, 1993; 2000) that is articulated into three main phenomenal properties: *mineness* (I consciously experience *my* body as having always phenomenally belonged to me); *selfhood* (I am directly acquainted with the contents of my self-consciousness); and *perspectivalness* (phenomenal space is organized around a centre, a supramodal point of view). The idea that the self can be described — and de-constructed — as a *model* becomes immediately appealing because it provides a functional architecture that can be easily matched with the mechanisms investigated by neuroscience at the neural level.

Throughout this paper I have emphasized the relevance of *agency*¹⁰ in modelling the reality we’re living in, and showed its neurobiological roots. By doing so, the gap between the neural and the mental levels of description can be greatly narrowed, especially if we characterize mental phenomenal properties as Metzinger

[10] See Proust (1998; 2000) for an insightful and stimulating account of how agency may support the attribution of beliefs and intentions to other agents.

(1993; 2000) does. I will now briefly focus on *perspectivalness* from a neurobiological standpoint, and I'll try to show how this higher-order component constituted by the PSM¹¹ is also intimately related to agency.

Our body is made of different parts, and indeed the neurophysiological literature on non-human primates has shown that different parts of the body (for example, the eyes, the head, the hands) rely on different spatial frames of reference supported by specific, distinct neural circuits (see for review Colby and Goldberg, 1999). For example, multimodal neurons discovered in area VIP and in sectors of area PF in the posterior parietal cortex, and in area F4 in the ventral premotor cortex, respond to the presence of objects only within a limited space sector, peri-personal space (Gentilucci *et al.*, 1988; Graziano *et al.*, 1994; Fogassi *et al.*, 1996; Rizzolatti *et al.*, 1997; Graziano and Gross 1995; Duhamel *et al.*, 1998).

Peri-personal space is by definition a motor space, its outer limits being defined by the working space of different body effectors such as the head or the arms. In fact, what is relevant to the neurons of these brain sectors is the location, with respect to the body, of 'something' that will become the target of a purposeful action. Again, we see that even space is inherently, intrinsically dependent on the dynamic relationship between agent and environment. Even more suggestive that this perspective is right are the data by Iriki and co-workers (1996). These authors discovered in a sector of the monkey posterior parietal cortex 'bimodal' (visual and somatosensory) neurons that fired when a tactile stimulus was delivered to the monkey's hand and when visual objects were presented near the hand's tactile receptive field. The striking characteristic of these neurons was that, during a reaching movement performed with a rake, their visual receptive field expanded to include the entire length of the rake. In other words, in that experiment, the monkey's body schema changed when the tool was used: the tool was assimilated to the hand and became part of the hand representation. As a consequence, the far space reached by the rake was re-mapped as 'near' space and the neurons which fired for near space also fired for the remapped far space.

This action-driven modulation of spatial maps has also been shown in humans. Berti and Frassinetti (2000) showed that when the cerebral representation of body space is extended to include tools that the subject uses, space previously mapped as far is treated as near. They described a patient (P.P.), who, after a right hemisphere stroke, exhibited a clear neglect in near space in many different tasks such as line cancellation, reading, and line bisection. When the lines were positioned far from the body (approximately 1 m) and the patient bisected them using a laser pointer, the neglect tended to disappear almost completely. However, when line bisection in far space was tested using a stick by means of which the patient could *reach* the line, neglect reappeared and was as severe as neglect in near space. As in monkeys, the use of a tool extended the body space of the patient and enlarged his peri-personal space to include all the space between the body and the stimulus.

[11] According to Metzinger (2000), the self-model is a component of the PMIR: Object-model plus 'dynamically complementarized' self-model = PMIR = a consciously experienced first-person perspective = perspectivalness (as a property of the current world-model of the system).

As a consequence far space was remapped as near. Because near space was affected by the lesion, neglect reappeared.

In conclusion, what these studies suggest is that — both in monkeys and humans — the neural sensorimotor systems supporting spatial cognition have a composite architecture. The lesion of a specific sector of this architecture impairs spatial awareness of the corresponding space sector, but leaves intact other parts of space representation.

According to Metzinger (1993; 2000) *perspectivalness* anchors the viewpoint from which ‘the world outside’ is experienced to a temporally extended, non-conceptual, supramodal, body-centred (and therefore *embodied*) frame of reference. Our body, however, as we have seen is made of different parts. By exploring the functional properties of specific and different neural circuits, we have learned that the different parts of the body (for example, the eyes, the head and the hands) rely on different spatial frames of reference. Furthermore, as we have seen, hemispatial neglect, a disturbance of spatial *awareness*, can be exclusively confined either to the peri-personal or to the extrapersonal space (for review, see Rizzolatti *et al.*, 2000a).

These results tell us that even one of the single features composing the self-model can be further decomposed, at the neuronal as well as at the neuropsychological level. After reaching these conclusions we need a ‘glue’ capable of giving coherence to the cubist multiplicity of perspectives and levels of descriptions produced by this de-constructivist account of the self and its features. My suggestion is that a good candidate for ‘gluing the pieces together’ might be the motor system, in its double aspect of action generation and action representation.

VI: Conclusions

At the beginning of the paper I emphasized the intrinsic ambiguity of the neural code. This ambiguity shouldn’t necessarily make us feel too embarrassed. The price we are traditionally supposed to pay for not feeling too embarrassed, however, is that we must stick to a representational–computational account of the mind. The central idea at the core of this account consists in positing that what the brain does can be described in terms of algorithms manipulating symbols. In other words, input information is symbolically represented and progressively transformed for perception and motor output. These internal symbolic representations of the external world have been referred to as a ‘language of thought’ or ‘mentalese’ (Fodor, 1975). Symbol manipulation obeys a set of rules — a universal syntax — considered by some to be innate. Knowledge is represented in a symbolic form. Meaning is referential since it derives from the posited correspondence between the system of symbols and the objects in the world. These are classified in terms of classical categories, whose membership is defined in terms of singly necessary and jointly sufficient conditions.

In this approach to the mind it is possible to recognize, on one side, the powerful influence of mathematical logic and artificial intelligence, on the other, that of

the Chomskian ‘universal grammar’ theory (Chomsky, 1986). This extension of some principles of logic and linguistics to biology was prompted also by the felt necessity to reject behaviourism which, in its most radical formulation, restricted the inquiry about the biological world to a too narrow and simplistic relationship between stimuli and responses. Unfortunately this rejection ended up by neglecting the whole biological substrate of the mind, the nervous system being reduced to one of the infinite possible machines that may manipulate symbols (see Edelman, 1992 for a strong defence of a ‘biological counter-revolution’).

On metaphysical ground, on the other hand, this representational–computational account of the mind forces us on to a track leading to dualism or to modified versions of it. We shouldn’t be discouraged, however, by these difficulties. The approach that I proposed in the present paper can lead us back to the safe bastions of physicalism, while at the same time preventing us from being trapped within the claustrophobic cabin of radical eliminativism (life is full of dangers . . .). Keeping neurons as a crucial level of description doesn’t necessarily impoverish our account of the world of subjective experience.

It is clear from this brief and incomplete survey of part of the recent neurophysiological literature that there are apparently conflicting representational rules in need of reconciliation, if they are supposed to ground the notion of consciousness. Our apparently unitary model of the world relies on the possibility of switching from one frame of *representational* reference to the other: first- and third-person perspectives (see Barresi and Moore, 1996) coexist within consciousness and at the same time define it.

It is possible to be even more explicit by saying that, in a broad sense, being *conscious* depends on the acquired capacity to recognize the existence of multiple frames of reference, and to put them in dynamic relation with each other by a continuous process of analogy and differentiation. This capacity likely developed as the best adaptive solution to the powerful pressure exerted by the presence of multiple and, in principle, conflicting frames of reference. Consciousness can be seen as the adaptive tool able to give coherence to these interacting levels of representation and to define *one* of these levels as the intranscendable background of reality. The adaptive logic of consciousness, once rooted in this evolutionary perspective, can be traced back to more ancient mechanisms of which it may represent the human homologue. The investigation and the understanding of these mechanisms is in my view inescapable for any serious account of consciousness. Consciousness, in this light, becomes the decisive step making the switch from biological to cultural evolution possible.

Furthermore, by stressing the relational nature of consciousness and the behavioural levels at which it operates, the links between consciousness and self appear to be stricter than one would assume they are. If consciousness is a model of the world, it must incorporate also the possibility of demarcating the world as *that something* that sets the limits of the self. In this perspective self and world depend on each other. Self-consciousness, as suggested by Bermudez (1998), can in fact be seen as a *contrastive* notion (see also Gallese, 2000b).

Is it possible to define the constituent parts of consciousness in terms of the functional properties of a series of neural circuits? I think that the answer is: yes and no. The clarification of the functional mechanisms and behavioural relevance of neural circuits, as I said before, is crucial to understanding how consciousness *has evolved* and how consciousness *operates*. However, this sub-personal level of analysis needs to be framed within a more comprehensive view that *must* incorporate the personal level of description, viewed as the dynamic interplay between an open system — our living body — and the environment.

This is the reason why I think that *agency* provides the key approach to the fascinating enquiry about how we are capable of modelling ourselves and the ‘world outside’.

Acknowledgements

The author wishes to thank Thomas Metzinger and four anonymous referees for their most valuable comments on an earlier version of the paper. Supported by a Human Frontier Scientific Program grant.

References

- Andersen, R.A. (1987), ‘Inferior parietal lobule function in spatial perception and visuomotor integration’, in *Handbook of Physiology. The Nervous System. Higher Functions of the Brain. Section 1, Volume 5*. (Bethesda, MA: American Physiological Society).
- Andersen R.A., Snyder, L.H., Chiang-Shan, L. and Stricanne, B. (1993), ‘Coordinate transformations in the representation of spatial information’, *Current Opinion in Neurobiology*, **3**, pp.171–6.
- Baars, B. (1997), *In the Theatre of Consciousness. The Workspace of the Mind* (London: Oxford University Press).
- Barresi, J. and Moore, L. (1996), ‘Intentional relations and social understanding’, *Behavioral Brain Sciences*, **19**, pp. 107–54.
- Bermudez, J.L. (1998), *The Paradox of Self-Consciousness* (Cambridge, MA: MIT Press).
- Berti, A. and Frassinetti, F. (2000), ‘When far becomes near: Re-mapping of space by tool use’, *Journal of Cognitive Neuroscience*, **12**, pp. 415–20.
- Bilgrami, A. (1992), *Belief and Meaning* (Oxford: Blackwell).
- Colby, C.L. and Goldberg, M.E. (1999), ‘Space and attention in parietal cortex’, *Annual Review of Neuroscience*, **22**, pp. 319–49.
- Clark, A. (1997), *Being There* (Cambridge, MA: MIT Press).
- Chomsky, N. (1986), *Knowledge of Language: Its Nature, Origin and Use* (New York and Philadelphia: Praeger Special Studies).
- Dennett, D.C. (1995), *Darwin’s Dangerous Idea* (New York: Simon and Schuster).
- Duhamel, J-R., Colby, C.L. and Goldberg, M.E. (1998), ‘Ventral intraparietal area of the macaque: Congruent visual and somatic response properties’, *Journal of Neurophysiology*, **79**, pp. 126–36.
- Eccles, J. (1989), *Evolution of the Brain: Creation of the Self* (London: Routledge).
- Edelman, G.M. (1992), *Bright Air, Brilliant Fire* (New York: Basic Books).
- Edelman, G.M. and Tononi, G. (2000), *A Universe of Consciousness: How Matter Becomes Imagination* (New York: Basic Books).
- Fodor, J. (1975), *The Language of Thought* (New York: Thomas Y. Crowell).
- Fodor, J. (1983), *The Modularity of Mind* (Cambridge, MA: MIT Press).
- Fodor, J. (1998), *In Critical Condition* (Cambridge, MA: MIT Press).
- Fogassi, L., Gallese, V., Fadiga, L., Luppino, G., Matelli, M. and Rizzolatti, G. (1996), ‘Coding of peripersonal space in inferior premotor cortex (area F4)’, *Journal of Neurophysiology*, **76**, pp. 141–57.
- Freeman, W.J. and Núñez, R. (1999), ‘Restoring to cognition the forgotten primacy of action, intention and emotion’, *Journal of Consciousness Studies*, **6** (11–12), pp. 9–19.

- Gallese, V. (2000a), 'The acting subject: Towards the neural basis of social cognition', in *Neural Correlates of Consciousness. Empirical and Conceptual Questions*, ed. T. Metzinger (Cambridge, MA: MIT Press).
- Gallese V. (2000b), 'The brain and the self: Reviewing the neuroscientific evidence', *Psychology*, **11** (34) (<http://www.cogsci.soton.ac.uk/cgi/psyc/newpsy?11.034>).
- Gallese, V. and Goldman, A. (1998), 'Mirror neurons and the simulation theory of mind-reading', *Trends in Cognitive Sciences*, **12**, pp. 493–501.
- Gallese, V., Craighero, L., Fadiga, L. and Fogassi, L. (1999), 'Perception through action', *Psyche*, **5** (21) (<http://psyche.cs.monash.edu.au/v5/psyche-5-21-gallese.html>).
- Gentilucci, M., Fogassi, L., Luppino, G., Matelli, M., Camarda, R. and Rizzolatti, G. (1988), 'Functional organization of inferior area 6 in the macaque monkey: I. Somatotopy and the control of proximal movements', *Experimental Brain Research*, **71**, pp. 475–90.
- Gibson, J. (1979), *An Ecological Approach to Visual Perception* (Boston: Houghton Mifflin).
- Goldman, A. (1989), 'Interpretation psychologized', *Mind and Language*, **4**, pp. 161–85.
- Goldman, A. (1992), 'In defence of the simulation theory', *Mind and Language*, **7**, pp. 104–19.
- Goldman, A. (1993), 'The psychology of folk psychology', *Behavioral Brain Sciences*, **16**, pp. 15–28.
- Goldman, A. (2000), 'The mentalizing folk', in *Metarepresentation*, ed. D. Sperber (London: Oxford University Press).
- Gordon, R. (1986), 'Folk psychology as simulation', *Mind and Language*, **1**, pp. 158–71.
- Gould, S.J. and Lewontin, R.C. (1979), 'The spandrels of San Marco and the Panglossian paradigm. A critique of the adaptationist programme', *Proceedings of the Royal Society of London*, **205**, pp. 281–8.
- Graziano, M.S.A., Yap, G.S. and Gross, C.G. (1994), 'Coding of visual space by premotor neurons', *Science*, **266**, pp. 1054–7.
- Graziano, M.S.A. and Gross, C.G. (1995), 'The representation of extrapersonal space: A possible role for bimodal visual–tactile neurons', in *The Cognitive Neurosciences*, ed. M.S. Gazzaniga (Cambridge, MA: MIT Press).
- Grush, R. (1997), 'The architecture of representation', *Philosophical Psychology*, **10**, pp. 5–25.
- Harris, P. (1989), *Children and Emotion* (Oxford, Blackwell Publishers).
- He, S.Q., Dum, R.P. and Strick, P.L. (1993), 'Topographic organization of corticospinal projections from the frontal lobe–motor areas on the lateral surface of the hemisphere', *Journal of Neuroscience*, **13**, pp. 952–80.
- Hurlay, S.R. (1998), *Consciousness in Action* (Harvard, MA: Harvard University Press).
- Hyvärinen, J. (1981), 'Regional distribution of functions in parietal association area 7 of the monkey', *Brain Research*, **206**, pp. 287–303.
- Iriki, A., Tanaka, M. and Iwamura, Y. (1996), 'Coding of modified body schema during tool use by macaque postcentral neurones', *NeuroReport*, **7**, pp. 2325–30.
- Jarvilehto, T. (1998), 'Efferent influences on receptors in knowledge formation', *Psychology*, **9**, p. 41.
- Jeannerod, M. (1994), 'The representing brain: Neural correlates of motor intention and imagery', *Behavioral Brain Sciences*, **17**, pp. 187–245.
- Jeannerod, M. (1997), *The Cognitive Neuroscience of Action* (Oxford: Blackwell).
- Jeannerod, M., Arbib, M.A., Rizzolatti, G. and Sakata, H. (1995), 'Grasping objects: The cortical mechanisms of visuomotor transformation', *Trends in Neuroscience*, **18**, pp. 314–20.
- Kawato, M. (1997), 'Bidirectional theory approach to consciousness', in *Cognition, Computation and Consciousness*, ed. M. Ito, Y. Myashita and E.T. Rolls (Oxford: Oxford University Press).
- Kawato, M. (1999), 'Internal models for motor control and trajectory planning', *Neuroreport*, **9**, pp. 718–27.
- Kim, J. (1998), *Mind in a Physical World* (Cambridge, MA: Bradford Book-MIT Press).
- Matelli, M. and Luppino, G. (1997), 'Functional anatomy of human motor cortical areas', in *Handbook of Neuropsychology*, vol. 11, ed. F. Boller and J. Grafman (Elsevier Science B. V.).
- Matelli, M., Luppino, G. and Rizzolatti, G. (1985), 'Patterns of cytochrome oxidase activity in the frontal agranular cortex of the macaque monkey', *Behavioral Brain Research*, **18**, pp. 125–37.
- Merleau-Ponty, M. (1962), *The Phenomenology of Perception*, tr. C. Smith (London: Routledge).
- Metzinger, T. (1993), *Subjekt und Selbstmodell* (Paderborn: Schoeningh).
- Metzinger, T. (2000), 'The subjectivity of subjective experience: A representationalist analysis of the first-person perspective', in *Neural Correlates of Consciousness. Empirical and Conceptual Questions*, ed. T. Metzinger (Cambridge, MA: MIT Press).
- Millikan, R.G. (1984), *Language, Thought, and other Biological Categories* (Cambridge, MA: MIT Press).

- Millikan, R.G. (1993), *White Queen Psychology and Other Essays for Alice* (Cambridge, MA: MIT Press).
- Milner, D. and Goodale, M.A. (1995), *The Visual Brain in Action* (Oxford: Oxford University Press).
- Mountcastle, V.B., Lynch, J.C., Georgopoulos, A., Sakata, H. and Acuna, C. (1975), 'Posterior parietal association cortex of the monkey: Command functions for operations within extrapersonal space', *Journal of Neurophysiology*, **38**, pp. 871–908.
- Murata, A., Fadiga, L., Fogassi, L., Gallese, V., Raos, V. and Rizzolatti, G. (1997), 'Object representation in the ventral premotor cortex (Area F5) of the monkey', *Journal of Neurophysiology*, **78**, pp. 2226–30.
- O'Regan, J.K. and Noë, A. (2000, in press), 'What it is like to see: A sensorimotor theory of perceptual experience', *Synthese*.
- Proust, J. (2000), 'Awareness of being the actor of one's actions: Three levels of analysis', in *Neural Correlates of Consciousness. Empirical and Conceptual Questions*, ed. T. Metzinger (Cambridge, MA: MIT Press).
- Proust, J. (1998), 'Presentation', in *Subjectivité et Conscience d'Agir*, ed. H. Grivois and J. Proust (Paris: Presses Universitaires de France).
- Rizzolatti, G., Camarda, R., Fogassi, M., Gentilucci, M., Luppino, G. and Matelli, M. (1988), 'Functional organization of inferior area 6 in the macaque monkey: II. Area F5 and the control of distal movements', *Experimental Brain Research*, **71**, pp. 491–507.
- Rizzolatti, G., Fadiga, L., Fogassi, L. and Gallese, V. (1997), 'The space around us', *Science*, **277**, pp. 190–91.
- Rizzolatti, G., Berti, A. and Gallese, V. (2000a), 'Spatial neglect: neurophysiological bases, cortical circuits and theories', in *Handbook of Neuropsychology* (second edition, Vol.1), ed. F. Boller, G. Grafman and G. Rizzolatti (Amsterdam: Elsevier), pp. 503–37.
- Rizzolatti, G., Fogassi, L. and Gallese, V. (2000b), 'Cortical mechanisms subserving object grasping and action recognition: A new view on the cortical motor functions', in *The Cognitive Neurosciences, Second Edition*, ed. M.S. Gazzaniga (Cambridge, MA: MIT Press).
- Rizzolatti, G., Luppino, G. and Matelli, M. (1998), 'The organization of the cortical motor system: New concepts', *Electroencephalography and Clinical Neurophysiology*, **106**, pp. 283–96.
- Sakata, H., Taira, M., Murata, A. and Mine, S. (1995), 'Neural mechanisms of visual guidance of hand action in the parietal cortex of the monkey', *Cerebral Cortex*, **5**, pp. 429–38.
- Tanji, J. (1994), 'The supplementary motor area in the cerebral cortex', *Neuroscience Research*, **19**, pp. 251–68.
- Ungerleider, L.G. and Mishkin, M. (1982), 'Two visual systems', in *Analysis of Visual Behavior*, ed. D.J. Ingle, M.A. Goodale and R.J.W. Mansfield (Cambridge, MA: MIT Press).
- Wise, S.P. (1985), 'The primate premotor cortex: Past, present, and preparatory', *Annual Review of Neuroscience*, **8**, pp. 1–19.
- Wolpert, D.M. (1997), 'Computational approaches to motor control', *Trends in Cognitive Sciences*, **1**, pp. 209–16.
- Wolpert, D.M., Ghahramani, Z. and Jordan, M.I. (1995), 'An internal model for sensorimotor integration', *Science*, **269**, pp. 1880–2.
- Zeki, S. (1993), *A Vision of the Brain* (Oxford: Blackwell).