

THE BRAIN'S CONCEPTS: THE ROLE OF THE SENSORY-MOTOR SYSTEM IN CONCEPTUAL KNOWLEDGE

Vittorio Gallese

Università di Parma, Italy

George Lakoff

University of California, Berkeley, USA

Concepts are the elementary units of reason and linguistic meaning. They are conventional and relatively stable. As such, they must somehow be the result of neural activity in the brain. The questions are: Where? and How? A common philosophical position is that all concepts—even concepts about action and perception—are symbolic and abstract, and therefore must be implemented *outside* the brain's sensory-motor system. We will argue against this position using (1) neuroscientific evidence; (2) results from neural computation; and (3) results about the nature of concepts from cognitive linguistics. We will propose that the sensory-motor system has the right kind of structure to characterise both sensory-motor and more abstract concepts. Central to this picture are the neural theory of language and the theory of cogs, according to which, brain structures in the sensory-motor regions are exploited to characterise the so-called “abstract” concepts that constitute the meanings of grammatical constructions and general inference patterns.

INTRODUCTION

Concepts are the elementary units of reason and linguistic meaning. They are conventional and relatively stable. As such, they must somehow be the result of neural activity in the brain. The questions are: Where? and How?

First-generation cognitive science was strongly influenced by the analytic tradition of philosophy of language, from which it inherited the propensity to analyse concepts on the basis of formal abstract models, totally unrelated to the life of the body, and of the brain regions governing the body's functioning in the world.

Correspondence should be addressed to Vittorio Gallese, Dipartimento di Neuroscienze, Sezione di Fisiologia, Università di Parma, Via Volturno 39, 43100 Parma, Italy/(E-mail: vittorio.gallese@unipr.it).

For many of the ideas put forward in the present paper the authors are indebted to, and would like to thank, the following members of the NTL Project at the International Computer Science Institute of the University of California at Berkeley: Jerome Feldman, Srin Narayanan, Lokendra Shastri, Eve Sweetser, Nancy Chang, Shweta Narayan, Ellen Dodge, Keith Sanders, and Ben Bergen. Many thanks also to Giacomo Rizzolatti and Maria Alessandra Umiltà, for their most valuable comments on an earlier version of this paper.

Vittorio Gallese, who was George Miller Visiting Professor at the University of California at Berkeley, would like to thank the Miller Institute for awarding him the fellowship that made this work possible. He was also supported by the EU Mirrorbot Project, and by the Eurocores Project “The Origin of Man, Language, and Languages.”

Q1 Concepts, from this perspective, were conceived of as abstract, amodal, and arbitrary, represented in some “language of thought” (Fodor, 1975, 1987), made up of symbols and having the properties of productivity and compositionality, among others. In Fodor’s theory (see Fodor, 1975), the purported amodal nature of concepts draws a sharp dividing line between the modular input/output brain structures and a generalised cognitive system (unanalysed at the level of the brain), whose functioning rules are totally independent from those governing the input/output modules. The propositional picture of the mind conveyed by early cognitivism is that of a functional system whose processes can be described in terms of manipulations of abstract symbols according to a set of formal syntactic rules (see Fodor, 1981; Pylyshyn, 1984). Knowledge is therefore represented in amodal symbolic form. Meaning is referential, in that it derives from a posited correspondence between the system of abstract symbols and their corresponding extensions, the objects and events in the world. Thus, following the line of arguments of early cognitivism, concepts are symbolic representations by nature, and as thinking, they can be reduced to symbolic (not neural) computation.

Q1 We will propose a radically different view. We will argue that conceptual knowledge is embodied, that is, it is mapped within our sensory-motor system. We will argue that the sensory-motor system not only provides structure to conceptual content, but also characterises the semantic content of concepts in terms of the way that we function with our bodies in the world. Before delving deeply into the argument, we should discuss a major finding in neuroscience that we will be assuming throughout: *Imagining and doing use a shared neural substrate.*

When one imagines seeing something, some of the same part of the brain is used as when one actually sees. When we imagine moving, some of the same part of the brain is used as when we actually move. Note that these facts undermine the traditional rationale given above. We can imagine grasping an object without actually grasping it. From this, it does not follow that actual grasping and imaginary grasping do not use a common

neural substrate. One can reason about grasping without grasping; yet one may still use the same neural substrate in the sensory-motor system. Indeed, that is just what we shall argue. In doing so, we will extend what we know about doing and imagining sharing a common substrate via the following hypothesis: *The same neural substrate used in imagining is used in understanding.*

Consider a simple sentence, like “Harry picked up the glass.” If you can’t imagine picking up a glass or seeing someone picking up a glass, then you can’t understand that sentence. Our hypothesis develops this fact one step further. It says that *understanding is imagination*, and that *what you understand of a sentence in a context is the meaning of that sentence in that context.*

Our proposal is not an internalist theory of meaning. The reason is that imagination, like perceiving and doing, is *embodied*, that is, structured by our constant encounter and interaction with the world via our bodies and brains. The result is an *interactionist* theory of meaning.

Accordingly, we will argue that a key aspect of human cognition is *neural exploitation*—the adaptation of sensory-motor brain mechanisms to serve new roles in reason and language, while retaining their original functions as well. We will discuss two cases: conceptual metaphor and cogs.

As we shall see, circuitry across brain regions links modalities, infusing each with properties of others. The sensory-motor system of the brain is thus “multimodal” rather than modular. Accordingly, language is inherently multimodal in this sense, that is, it uses many modalities linked together—sight, hearing, touch, motor actions, and so on. Language exploits the pre-existing multimodal character of the sensory-motor system. If this is true, it follows that there is no single “module” for language—and that human language makes use of mechanisms also present in nonhuman primates.

According to our proposal, the concept *grasp*, from which we will start, gets its meaning via our ability to imagine, perform, and perceive *grasping*. Our ability to imagine grasping makes use of the same neural substrate as performing and perceiving *grasping*. According to our proposal, imagining is a form of *simulation*—a mental simulation of

action or perception, using many of the same neurons as actually acting or perceiving (Gallese, 2003a).

Before developing our arguments, we would like to conclude this introductory part by framing our proposal within the extant and copious literature on the neural underpinnings of conceptual knowledge. Any serious attempt to provide a neuroscientific account of conceptual content as nested in the activity of the brain faces the challenge of explaining how the localised patterns of activation of different neural cortical networks can enable the capacity to distinguish, recognise, categorise, and ultimately conceptualise objects, events, and the state of affairs in the world.

Two main approaches have so far challenged the view held by first-generation cognitive science on the nature of conceptual knowledge: clinical neuropsychology and cognitive neuroscience. Clinical neuropsychology has established a relationship between given patterns of localised brain damage and corresponding deficits in conceptual knowledge. Cognitive neuroscience has more recently tried to establish, mainly by means of brain imaging experiments, which brain regions are activated by different conceptual categories.

During the last two decades, an impressive amount of clinical data has accumulated, describing patients whose peculiarly localised brain lesions have determined selective impairments of their conceptual knowledge. Basically, most of these deficits encompass the loss of *some specific categories of conceptual knowledge*, such as living things, or nonliving objects (mainly tools and artifacts). An excellent and thorough survey of this literature can be found in the recent special issue of *Cognitive Neuropsychology* (Vol. 20, no. 3–6, 2003).

A parallel conspicuous brain imaging literature has accumulated on the neural correlates of distinct conceptual categories (for recent review, see Gainotti, 2004; Malach, Levy, & Hasson, 2002). The discussion of this vast literature is beyond the scope of the present article. What we would like to highlight here is that both the clinical and brain imaging literature have not to date provided a unified explanatory framework. Basically, three theories of conceptual deficit dominate the literature:

sensory-functional theory, correlated structure theory, and domain-specific theory (see Martin & Caramazza, 2003; Simmons & Barsalou, 2003). These theories differ along many dimensions, the principal one being the extent to which conceptual knowledge is structured—and henceforth selectively affected by localised brain damage, by property or by category.

Similarly, it has recently been argued (Gainotti, 2004) that the overall picture provided by functional brain imaging studies is by no means consistent and clear-cut (see also Gallese, 2003c). It should be added that even if one could provide unambiguous evidence for the neural correlates of specific object concepts, the general principle defining the topology of such neural representation would need to be convincingly demonstrated. Unfortunately, as argued by Malach et al. (2002), to date no such convincing solution has been proposed.

Our present proposal does not aim to settle the clinical issues, and so we will not deal with clinical cases. Our goal is to follow an alternative route, that is, to provide a testable embodied theory of concepts, based on the results of research in neuroscience, neural computation, and cognitive linguistics, capable of reconciling both concrete and abstract concepts within a unified framework. The structure of our argument is delineated in the next section.

THE STRUCTURE OF THE ARGUMENT

We will begin with the action concept *grasp*. The argument will take the following form.

1. Information structure: We will show that the information structure needed to characterise the conceptual structure of *grasp* is available at the neural level in the sensory-motor system. That includes the semantic role structure, the aspectual structure, and certain hierarchical category structures.

2. *Multimodality*. Mirror neurons and other classes of premotor and parietal neurons are

Q2

Q3

inherently “multimodal” in that they respond to more than one modality. Thus, the firing of a single neuron may correlate with both *seeing* and *performing* grasping. Such multimodality, we will argue, meets the condition that an action-concept must fit both the performance and perception of the action.

3. *Functional clusters.* Multimodality is realised in the brain through *functional clusters*, that is, among others, parallel parietal-premotor networks. These functional clusters form *high-level units*—characterising the discreteness, high-level structure, and internal relational structure required by concepts¹.

4. *Simulation.* To understand the meaning of the concept *grasp*, one must at least be able to imagine oneself or someone else grasping an object. *Imagination is mental simulation* (see Gallese, 2003a), carried out by the same functional clusters used in acting and perceiving. Any conceptualisation of grasping via simulation therefore requires the use of the same functional clusters used in the action and perception of grasping.

5. *Parameters.* All actions, perceptions, and simulations make use of neural parameters and their values. For example, the action of *reaching* for an object makes use of the neural parameter of direction; the action of *grasping* an object makes use of the neural parameter of force. So do the concepts of *reaching* and *grasping*. Such neural parameterisation is pervasive and imposes a *hierarchical structure* on the brain: The same parameter values that characterise the *internal structure* of actions and simulations of actions also characterise the *internal structure* of action concepts.

6. *Structured neural computation.* The neural theory of language (see Feldman & Narayanan, in press; Lakoff & Johnson, 1999) provides a theory of neural computation in which the same neural structures that allow for movement and perception in real time and in real sensory-motor contexts also permit real-time context-based inferences in reasoning. The same neural structures that carry out action and perception carry out inference.

These six points will allow us to characterise an *embodied* theory of concepts, grounded in the sensory-motor system. At first we will limit ourselves to the case of action-concepts like *grasp*. After that, we will suggest how this theory, with a couple of additions, will extend to concepts more generally.

There are several points to be borne in mind: First, the neuroscientific research we will cite is partly done on monkeys and partly on humans. We will use the results on monkeys as applying to humans for the simple reason that there is enough evidence to support the notion of an analogy—when not a homology—between the monkey and human brain regions we will be discussing (see Rizzolatti, Fogassi, & Gallese, 2002).

Second, there is far more to the sensory-motor system than we will be discussing, and much of it is relevant. For example, we will not be discussing the roles of basal ganglia, cerebellum, thalamus, and somato-sensory cortices. Though they would add to the argument, they would also add greatly to the length of this study, and we believe we can make our point without them.

Third, as we stated at the outset, any theory of concepts must account for how concepts are implemented in the brain and must provide empirical evidence for such a theory. Let us now turn to the results from neuroscience.

MULTIMODAL FUNCTIONAL CLUSTERS AND EMBODIED SIMULATION

Before we look at the multimodality of action concepts, we need to look at the multimodality of actions themselves. The action of *grasping* has both a motor component (what you do in grasping) and various perceptual components (what it looks like for someone to grasp and what a graspable object looks like). Although we won't discuss them here, there are other modalities involved as

¹ Our theory assumes that concepts do have internal structure. This notion, however, is controversial. For an alternative account, see Fodor (1998).

well, such as the somato-sensory component (what it feels like to grasp something).

It is important to distinguish multimodality from what has been called “supramodality.” The term “supramodality” is generally (though not always) used in the following way: It is assumed that there are distinct modalities characterised separately in different parts of the brain and that these can only be brought together via “association areas” that somehow integrate the information from the distinct modalities. To claim that an action like grasping is “supramodal” is to say that it is characterised in an association area, distinct and different from the sensory-motor system, which integrates information from the motor system with information from sensory modalities. The point is that anything supramodal uses information coming from areas specialised for individual distinct modalities, but is not itself involved in the individual distinct modalities.

To claim, as we do, that an action like *grasping* is multimodal is to say that (1) it is neurally enacted using neural substrates used for both action and perception, and (2) that the modalities of action and perception are integrated at the level of the sensory-motor system itself and not via higher association areas.

To see the difference, consider the following example. Premotor area F4 (a sector of area 6 in the macaque monkey brain) was once conceived of as a relatively uninteresting extension of the primary motor cortex, whose only role was to control axial and proximal movements of the upper limbs. However, it has been shown that F4 contains neurons that integrate motor, visual, and somato-sensory modalities for the purpose of controlling actions in space and perceiving peri-personal space, that is, the area of space reachable by body parts (Fogassi et al., 1992, 1996a; Gentilucci et al., 1988; Gentilucci, Scandolaro, Pigarev, & Rizzolatti, 1983; Rizzolatti, Camarda, Fogassi, Gentilucci, Luppino, & Matelli, 1988; Rizzolatti, Fogassi, & Gallese, 2000b; Rizzolatti, Fadiga, Fogassi, & Gallese, 1997; Rizzolatti & Gallese, 2004; Rizzolatti, Matelli, & Pavesi, 1983; Rizzolatti, Scandolaro, Matelli, & Gentilucci, 1981b). Similar results about multimodal integration in area F4 were independently

obtained by Michael Graziano, Charlie Gross, and their co-workers (Graziano, Hu, & Gross, 1997a, b; Graziano, Yap, & Gross, 1994; Gross & Graziano, 1995). More recently, Graziano, Reiss, & Gross (1999) showed that F4 neurons integrate not only visual but also auditory information about the location of objects within peripersonal space.

The point here is that the very same neurons that control purposeful actions also respond to visual, auditory, and somato-sensory information about the objects the actions are directed to. They do so because they are part of a parietal-premotor circuit (F4-VIP, see below) in charge of overall control of purposeful bodily actions in peri-personal space. This contrasts with the old notion that sensory-motor integration is achieved at a “higher” level at which separate neural systems for motor control and sensory processing are brought together in a putative “association area.”

This is important theoretically because supramodality is consistent with the idea of strict modularity, while multimodality is not. Supramodality accords with a picture of the brain containing separate modules for action and for perception that need to be somehow “associated.” *Multimodality denies the existence of such separate modules.*

Multimodality does everything that supramodality has been hypothesised to do, and more. Multimodal integration has been found in many different locations in the brain, and we believe that it is the norm (for a review, see Fogassi & Gallese, 2004). That is, sensory modalities like vision, touch, hearing, and so on are actually integrated with each other *and* with motor control and planning. This suggests that there are no pure “association areas” whose only job is to link supposedly separate brain areas (or “modules”) for distinct sensory modalities.

The neuroscientific evidence accumulated during the last two decades shows the following. Cortical premotor areas are endowed with sensory properties. They contain neurons that respond to visual, somatosensory, and auditory stimuli. Posterior parietal areas, traditionally considered to process and associate purely sensory information, in fact play a major role in motor control. The premotor and parietal areas, rather than having separate and

independent functions, are neurally integrated not only to control action, but also to serve the function of *constructing* an integrated representation of (1) actions together with (2) objects acted on and (3) locations toward which actions are directed.

In particular, these multimodal functions have been described within three parallel parietal-premotor cortical networks: F4-VIP, F5ab-AIP, and F5c-PF, which we will characterise as “functional clusters.” By a “cluster” we do not just mean a bunch of individual neurons in the same place. A *functional cluster* is a cortical network that functions as a unit with respect to relevant neural computations.

1. The F4-VIP cluster functions to transform the spatial position of objects in peri-personal space into the most suitable motor programmes for successfully interacting with the objects in those spatial positions—reaching for them or moving away from them with various parts of your body such as the arm or head. The properties of the object are far less important than their spatial position. Damage to this cluster will result in the inability to be consciously aware of, and interact with, objects within the contralateral peri-personal space (see Rizzolatti, Berti, & Gallese, 2000a).

2. The F5ab-AIP cluster contains “canonical neurons,” which transform the intrinsic physical features of objects (e.g., shape, size) into the most suitable hand motor programmes required to act on them—manipulate them, grasp them, hold them, tear them apart. In this cluster, the properties of the objects are far more important than their spatial location. Accordingly, damage to this functional cluster will induce visuo-motor grasping deficits, that is, the inability to grasp an object, despite having the motor capacity for grasping (see Fogassi et al., 2001; Gallese, Murata, Kaseda, Niki, & Sakata, 1994).

3. The F5c-PF cluster contains mirror neurons that discharge when the subject (a monkey in the classical experiments) performs various types of hand actions that are goal-related and also when the subject observes another individual performing similar kinds of actions (see Rizzolatti, Fogassi, & Gallese, 2001).

In the next three sections we will review some of the crucial properties of these three functional clusters, and discuss their functional mechanisms in terms of simulation.

Actions and their locations

The F4-VIP cluster: Simulation in action-location neurons

Actions occur at locations. Natural language codes the location where a given action occurs via locative adverbs, as in *He grasped the cup in front of him*. The semantic relation between an action and its location is part of conceptual structure. We submit that this semantic relation be characterised neurally by the following means.

Within the F4-VIP cluster, there are neurons that discharge when a subject (a monkey) turns its head toward a given location in peri-personal space. The same neurons also discharge when an object is presented, or a sound occurs, at the very same location toward which the head would be turned, if it were actually turned. Peri-personal space is by definition a *motor space*, its outer limits defined by the action space of the various body effectors—hands and arms, feet, head. In these cases, a position in peri-personal space can be specified in a number of ways: sound, sight, and touch (Duhamel, Colby, & Goldberg, 1998; Fogassi et al., 1996a; Gentilucci et al., 1988; Graziano & Gross, 1995; Graziano et al., 1994; Rizzolatti et al., 1997).

We maintain that what integrates these sensory modalities is *action simulation*. Because sound and action are parts of an integrated system, the sight of an object at a given location, or the sound it produces, automatically triggers a “plan” for a specific action directed toward that location. What is a “plan” to act? We claim that it is a simulated potential action.

These neurons control the execution of a specific real action (turning the head, say, 15 degrees to the right). When they fire without any action in presence of a possible target of action seen or heard at the same location (say, 15 degrees to the right), we hypothesise that they are *simulating* the action. This is explanatory for the following reason. In

simulation the same neural substrate is used as in action. If simulation is being carried out here, this would explain why just those neurons are firing that otherwise could act on the same object in the same location.

Action, patients, and purposes

The F5ab-AIP cluster: Simulation in canonical neurons

Let us now turn to how various conceptual relations are realised neurally: The relation between an action and something acted on (a patient), as well as relations like manner and purpose. The same multimodal cluster characterises the relation between a general action and a specific type of action, as well as aspectual phases of an action. For the sake of brevity, we will focus only on the properties of the premotor pole of this cluster, that is, area F5 (for a description of AIP, see Murata, Gallese, Luppino, Kaseda, & Sakata, 2000; Rizzolatti, Fogassi, & Gallese, 2000a; Sakata et al., 2000).

Q5

In premotor area F5 (Matelli, Luppino, & Rizzolatti, 1985), there are action-only neurons, so-called because they only fire during real actions (Gentilucci et al., 1988; Hepp-Reymond, Hüsler, Maier, & Qi, 1994; Kurata & Tanji, 1986; Rizzolatti et al., 1988; Rizzolatti, Scandolara, Gentilucci, & Camarda, 1981a). These neurons discharge any time the subject (a monkey) performs hand or mouth movements directed to an object (Rizzolatti et al., 1988). Several aspects of these neurons are important. First, what correlates to their discharge is not simply a *movement* (e.g., flexing the fingers, or opening the mouth), but an *action*, that is, a *movement executed to achieve a purpose* (grasp, hold, tear apart an object, bring it to the mouth). Second, what matters is the *purpose of the action*, and not some dynamic details defining it, like force, or movement direction (see Rizzolatti et al., 2000).

Q6

Q13

For any particular type of purposeful action, there are a number of kinds of subclusters.

1. *The general-purpose subclusters*: The neurons of these subclusters indicate the general goal of the action (e.g., grasp, hold, tear an object). They are

not concerned with either the details of how the action is carried out, nor the effector used (e.g., hand, mouth), nor how the effector achieves the purpose of the action (e.g., grasping with the index and the thumb, or with the whole hand).

2. *The manner subclusters*: The neurons of these subclusters concern the various ways in which a particular action can be executed (e.g., grasping an object with the index finger and the thumb, but not with the whole hand).

3. *The phase subclusters*: The neurons of these subclusters deal with the temporal phases purposeful actions are segmented (e.g., hand/mouth opening phase, or hand/mouth closure phase).

Thus, there is a general grasping-purpose subcluster that is active whenever grasping of any kind is carried out. Consider a particular case: What is firing during the closure phase of a precision-grip grasp? Three subclusters. (1) The subcluster for general-purpose grasping. (2) The subcluster for precision-grip grasping (a particular manner). (3) The subcluster for closure phase grasping.

Of course, the general-purpose subcluster for grasping can never function alone in action, since all actions are carried out in some manner and are in one phase or another at some time. However, it is at least in principle possible for the general-purpose subcluster for grasping to fire without a manner subcluster firing, *in simulation*. That is, you should be able to simulate something in imagination that you cannot do—carry out a *general* action without specifying manner. This is important for the theory of concepts. We can conceptualise a generalised grasping without any particular manner being specified.

The action-only neurons fire only when actions are carried out. But premotor area F5 also contains what are called “canonical neurons”—grasping-related neurons that fire not only when a grasping action is carried out, but also when the subject (a monkey) sees an object that it could grasp, but doesn't (see Rizzolatti et al., 2000b). These canonical neurons have both a general-purpose subcluster and a manner subcluster for cases where the grasping action is carried out. No experiments have yet been done to determine in detail the phases of

firing in such subclusters, though it is surmised that they will have phase subclusters as well.

There is a simulation explanation for the behavior of canonical neurons: If the sight of a graspable object triggers the simulation of grasping, we would expect there to be firing by at least some of the neurons that fire during actual grasping. This indeed is what happens with canonical neurons.

Strong evidence for the simulation hypothesis comes from the following data. In most canonical grasping-manner neurons, there is a strict correlation: The same neurons fire for a given manner of grasping as for merely observing an object that, if grasped, would require the same manner of grasping. For example, if a small object is presented, no matter what its shape is, then the same neurons fire as would fire if that small object were being picked up with a precision grip (as afforded by a small object of any shape). This is strong *prima facie* evidence that simulation is taking place: When you observe a graspable object, only the neurons with the right manner of grasping for that object fire (see Gallese, 2003b).

Observing the actions of others

The F5c-PF cluster: Simulation in mirror neurons

Within the F5c-PF cluster, there are individual neurons that are activated both during the execution of purposeful, goal-related hand actions, such as grasping, holding, or manipulating objects, and during the observation of similar actions performed by another individual. These neurons are called “mirror neurons” (Gallese, 1999, 2000a, 2001, 2003a, b; Gallese, Fadiga, Fogassi, & Rizzolatti, 1996; Gallese, Fogassi, Fadiga, & Rizzolatti, 2002; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996; Rizzolatti et al., 2000b, 2001). Mirror neurons, unlike canonical neurons, do not fire when just presented with an object one can act upon. They also do not fire when the observed action is performed with a tool, such as pliers or pincers.

Some mirror neurons (roughly 30%) are “strictly congruent.” They fire when the action seen is exactly the same as the action performed.

Others (roughly 70%) show hierarchical relations: They fire when either (1) the monkey grasps with a pincer grip, or (2) the monkey sees someone perform *any* type of grasping. (1) is a special case of (2), and hence the relationship between the special case and the general case is there in the neural structure.

Certain mirror neurons show aspectual phases such as the aspectual phases of an action, e.g., they fire during the central part of this action and the concluding part.

Here too, there is a general explanation in terms of simulation: When the subject (a monkey) observes another individual (monkey or human) doing an action, the subject is automatically simulating the same action. Since action and simulation use some of the same neural substrate, that would explain why the same neurons are firing during action-observation as during action-execution.

An even stronger argument in favour of the simulation interpretation comes from the following experiments. In the first series of experiments, F5 mirror neurons were tested in two conditions: (1) a condition in which the subject (a monkey) could see the entire action (e.g., a grasping-action with the hand), and (2) a condition in which the same action was presented, but its final critical part—that is, the hand-object interaction—was hidden. In the hidden condition the monkey only “knew” that the target object was present behind the occluder. The results showed that more than half of the recorded neurons responded in the hidden condition (Umiltà et al., 2001). These data indicate that, like humans, monkeys can also infer the goal of an action, even when the visual information about it is incomplete. This inference can be explained as the result of a simulation of that action by a group of mirror neurons (see also Gallese, 2003a).

A second series of experiments investigated what could possibly be the neural mechanism underpinning the capacity to understand the meaning of an action on the basis of its sound alone. F5 mirror neurons were tested in four conditions: When the monkey (1) executed noisy actions (e.g., breaking peanuts, tearing sheets of paper apart, and the like); and (2) just saw, (3) saw and heard, and (4) just heard the same actions performed by another

Q8

individual. The results showed that a consistent percentage of the tested mirror neurons fired under all four conditions (Keysers, 2003; Kohler et al., 2001, 2002). These neurons not only responded to the sound of actions, but also discriminated between the sounds of different actions: Each sound matched the appropriate action, whether observed or executed.

The hypothesis again is simulation: When the subject (a monkey) hears another individual performing an action with a distinctive sound, the subject is simulating the same action. Since action and simulation use some of the same neural substrate, that would explain why the same neurons are firing during observing, hearing, and executing the same action.

EVIDENCE FAVOURING THE HYPOTHESIS OF EMBODIED SIMULATION IN HUMANS

All of the cases cited above come from studies of monkeys. There are also correlates of the same results for humans. To the extent that the monkey studies constitute evidence for simulation, so do the studies on humans. This evidence for simulation makes even stronger the case for simulation made by the evidence given from the studies of visual and motor imagery (see below).

First, the action-location neurons: Recent brain-imaging experiments probed a cluster in humans located in the ventral premotor cortex and in the depth of the intraparietal sulcus, homologous to F4-VIP in monkeys. Neurons in this cluster were activated when subjects heard or saw stimuli being moved in their peri-personal space (Bremner et al., 2001). The significance of this is that one of the areas activated during such perception is a premotor area, the area that would most likely control movements aimed at objects in peri-personal space.

Second, the canonical neurons: In several recent brain-imaging experiments, subjects were asked to (1) observe, (2) name silently, and (3) imagine using various man-made objects (e.g., hammers, screwdrivers, and so on). In all these cases, there was activation of the ventral premotor cortex, that is,

the brain region activated when using those same tools to perform actions (Chao & Martin, 2000; Grafton, Arbib, Fadiga, & Rizzolatti, 1996; Martin et al., 1996; Perani et al., 1995).

Third, the mirror neurons: Several studies using different experimental methodologies and techniques have demonstrated in humans the existence of a mirror system, similar to that observed in monkeys, matching action observation and execution (see Buccino et al., 2001; Cochin, Barthelemy, Lejeune, Roux, & Martineau, 1998; Decety et al., 1997; Fadiga, Fogassi, Pavesi, & Rizzolatti, 1995; Grafton et al., 1996; Hari et al., 1998; Iacoboni et al., 1999; Rizzolatti et al., 1996). In particular, brain-imaging experiments in humans have shown that, during action observation, there is a strong activation of premotor and parietal areas, which are very likely to be the human homologue of the monkey areas in which mirror neurons were found (Buccino et al., 2001; Decety & Grèzes, 1999; Decety et al., 1997; Grafton et al., 1996; Iacoboni et al., 1999; Rizzolatti et al., 1996).

MENTAL IMAGERY: EMBODIED SIMULATION

All human beings entertain the capacity to imagine worlds that they have or have not seen before, to imagine doing things that they have or have not done before. The power of our imagination is seemingly infinite. Indeed, mental imagery has been considered for ages as one of the most characteristic aspects of the human mind, as it has been taken to epitomise its disembodied nature. Mental imagery used to be thought of as "abstract" and "fanciful", far from, and independent of, the perception of real objects and actions. In the light of neuroscientific research, though, things look quite different: We now know that visual and motor imagery *are embodied*.

1. Embodied visual imagery: *Some of the same parts of the brain used in seeing are used in visual imagination* (imagining that you are seeing). (For a comprehensive review, see Farah, 2000; Kosslyn & Thompson, 2000.)

Q9

2. Embodied motor imagery: *Some of the same parts of the brain used in action are used in motor imagination* (imagining that you are acting). Thus, imagination is not separate in the brain from perception and action (see Jeannerod, 1994).

The evidence comes from a variety of studies. For example, the time it takes to scan a visual scene is virtually identical to the time employed to scan the same scene when it is only imagined (Kosslyn, Ball, & Reiser, 1978). Furthermore, and more importantly, brain-imaging studies show that when we engage in imagining a visual scene, we activate regions in the brain that are normally active when we actually perceive the same visual scene (Farah, 1989; Kosslyn, 1994; Kosslyn et al., 1993). This includes areas, such as the primary visual cortex, involved in mapping low-level visual features (Le Bihan et al., 1993).

Motor imagery shows the same embodied nature as visual imagery. Mentally rehearsing a physical exercise has been shown to induce an increase of muscle strength comparable to that attained by a real exercise (Yue & Cole, 1992). When we engage in imagining the performance of a given action, several bodily parameters behave similarly to when we actually carry out the same actions. Decety (1991) has shown that heartbeat and breathing frequency increase during motor imagery of physical exercise. As in real physical exercise, they increase linearly with the increase of the imagined effort.

Finally, brain-imaging experiments have shown that motor imagery and real action both activate a common network of brain motor centres, such as the premotor cortex, the supplementary motor area (SMA), the basal ganglia, and the cerebellum (Decety, Sjöholm, Ryding, Stenberg, & Ingvar, 1990; Fox, Pardo, Petersen, & Raichle, 1987; Parsons et al., 1995; Roland, Larsen, Lassen, & Skinhoj, 1980; see also Jeannerod, 1994).

These data all together show that typical human cognitive activities such as visual and motor imagery, far from being of a disembodied, modality-free, and symbolic nature, make use of the activation of sensory-motor brain regions.

Let us conclude this section with a note on multimodality. In natural language, the same verbs

refer to actions performed and the corresponding actions seen or heard. The verbs are not limited by a particular modality. There are two possible accounts of the neural underpinnings of this aspect of language. (1) Modality-neutrality—a neural system outside the sensory-motor system altogether, with verbs expressing modality-neutral concepts. (2) Multimodality of the sort we have just seen: Connections across brain areas result in coordinated multimodal neural firing, and verbs express such multimodal concepts. We have just seen evidence for the existence of the appropriate multi-modality. Before we go on to the implications of all this for concepts, we will take up the topic of parameters.

PARAMETERS

A cat has three gaits—strutting, trotting, and galloping. Each gait requires a distinct motor programme. In galloping, for example, the front legs move together and the back legs move together. Strutting and trotting involve very different motor control of the legs. In short, the cat has three very different motor circuits to control its gait.

What is remarkable is that it has been discovered that there is a single cluster of neurons, a central pattern generator, that controls which gait is chosen. When those neurons are firing at low frequency, the cat struts; when the firing is at intermediate frequency, the cat trots; and at high frequency, the cat gallops (for review, see Grillner & Wallen, 2002; Yamaguchi, 2004). In other words, there are three values of firing frequency over a single collection of neurons—low, medium, and high—that result in the activation of either the strutting, trotting, or galloping gait. The firing frequency over that collection of neurons is a *neural parameter* and the mutually exclusive low, medium, and high firing frequencies are *values* of that neural parameter.

Parameters can be seen as “higher-level” features of neural organisation, while the neural firings in particular motor circuits for various gaits can be seen as being at a “lower level” of organisation. Given the higher-level firing, all the lower-level

Q10

firings are automatically driven as part of an encapsulated routine. To the higher-level parameters, the lower-level structure is “invisible.” Parameterisation thus imposes a hierarchical structure on the neural system.

Parameterisation is a pervasive feature of the brain. Here are some further examples:

1. In any given motor task, a certain level of force is appropriate. Level of force is a parameter for each motor task, and degrees of force are its values. The degree of force is controlled in the brain in one of two ways: either the *level* of activation of some cluster of motor neurons, or the *number* of motor neurons activated (see Porter & Lemon, 1993).

2. Direction of motion is also a parameter for actions. Two mechanisms have been proposed for determining values of the direction of movement parameter: (a) groups of neurons are selectively tuned to control a movement in a particular direction (see Gentilucci et al., 1988); (b) direction is determined by a “vector sum” over a whole population of neurons, each of which is only broadly tuned, that is, tuned to a range of directions (see Georgeopoulos, Schwartz, & Kettner, 1986). In either case, there is a direction parameter and a neural mechanism for determining specific values of that parameter.

3. In any given action description, role parameters play a major role. Mirror neurons map different actions (e.g., grasping, holding, tearing, placing, kicking an object) by specifying the agentive relation, while being neutral about the specific quality or identity of the agentive/subjective parameter. Other clusters map this information. For example, the activation of pre-SMA or the primary motor cortex is present only when one executes the action, but not when one is observing it being performed by someone else (Buccino et al., 2001; Ehrsson et al., 2000; for a recent review of the neural correlates of the *who* parameter, see Jackson & Decety, 2004).

The parameter–simulation link

In the enactment of any particular movement, say, pushing an object in a direction with a given force, the parameter values chosen determine where and

how hard one pushes. Moreover, if the force required is very high, what is required is *showing* rather than mere *pushing*. Shoving requires a different motor programme: setting the weight on the back foot, and so on. Thus, the choice of parameter values also determines motor programmes for humans as well as for cats. Moreover, parameter values govern simulations as well. Imagining pushing is different from imagining shoving.

The parameterisation hierarchy and the capacity to set parameter values are basic features of the brain. The parameters used in everyday perception and action are stable—built into our neural structure. In order to carry out any action or simulation, suitable parameter values must be activated. But there is a difference between parameter structures, on the one hand, and the actions and simulations they control. Both simulations and actions are *dynamic* and *contextually adapted*. Parameters are *fixed*. Whenever you act, there is always a neurally determined action, force, direction, amplitude, and so on. But the situation you are in affects the ultimate *values* of the parameters—exactly when, where, and how the action is carried out. Similarly, all simulation occurs via choice of the values of fixed parameters, which are determined dynamically in the context of the simulation.

The accessibility of parameters

Parameters and their values impose a hierarchical structure on the brain in the following sense. Once a value for a parameter is chosen, lower-level automatic neural mechanisms take over, say, to apply a force of a given magnitude or move in a given direction. Parameters and the kinds of values they have may be brought to consciousness. For example, all of us know we can press our palms forward with high degree of force. But we do not know how that is carried out neurally. Thus, parameters and their values are accessible to consciousness, while anything below the parameter value level is inaccessible.

Similarly, language may express parameters and their values, but language cannot express anything below the level of parameter values. Parameters and their values are thus also accessible to language, while lower-level neural structures are not.

After having discussed the relevance of parameters, there are more results from cognitive science that need to be mentioned before we move on. They concern basic-level categories.

BASIC-LEVEL CATEGORIES

The classic, and long taken for granted theory of categorisation assumed that categories formed a hierarchy—bottom to top—and that there was nothing special about those categories in the middle. This view was challenged by the research by Berlin, Rosch, and their co-workers in the 1970s (Berlin, Breedlove, & Raven, 1974; Berlin & Kay, 1969; Rosch, 1977, 1978, 1981, 1994; Rosch & Lloyd, 1978; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Take hierarchies like *furniture/chair/rocking chair* or *vehicle/car/sports car*. The categories in the middle—*chair* and *car*—are special; what Rosch called “basic-level” categories. One can get a mental image of a chair or a car, but not of a piece of furniture in general or a vehicle in general. We have motor programmes for interacting with chairs and cars, but not with furniture in general or vehicles in general. The basic level is the highest level at which this is true. Moreover, words for basic-level categories tend to be recognisable via gestalt perception, be learned earlier, to be shorter (e.g., *car* vs. *vehicle*), to be more frequent, to be remembered more easily, and so on.

Rosch observed that the basic level is the level at which we interact optimally in the world with our bodies. The consequence is that categorisation is embodied—given by our interactions, not just by objective properties of objects in the world, as a long philosophical tradition had assumed. Without us—without the way we sit and the way we form images—the wide range of objects we have called “chairs” do not form a category. A simple sentence like *Some chairs are green* is not true of the world independent of us, since there are neither chairs nor green things independent of us. It is, of course, true relative to *our body-based understanding* of the world. Our concepts must also be characterised relative to such a body-based understanding.

AN EMBODIED NEURAL THEORY OF CONCEPTS

We are now in a position to propose how these basic results from neuroscience and cognitive science allow us to characterise in neural terms not just actions, but *action concepts*. We have chosen to start with the concept of *grasping* for two reasons. First, we know quite a lot about the neuroscience of grasping; enough, we believe, to get fairly far. Second, the traditional theory requires *all* concepts to be disembodied, that is, above the level of everything we have talked about so far and not making use of any of it. This includes action concepts like *grasping*. The usual assumption is that a *concept* is disembodied, that is, modality-neutral and symbolic, while the *action that the concept designates* is of course embodied. Proponents of the traditional view would therefore maintain that any attempt to say that concepts are embodied would amount to confusing the *concept* with *what the concept designates*.

Our response will have two parts: First, we will argue that parameters and simulations can do the jobs that everyone agrees that concepts must do. Second, we will argue that the traditional theory does not accord with the results of neuroscience that we have just given.

What is an embodied concept?

Here is our central claim on embodied concepts: *The job done by what have been called “concepts” can be accomplished by schemas characterised by parameters and their values.* Such a schema, from a neural perspective, consists of a network of functional clusters. The network constituting a schema contains:

1. One cluster for each *parameter*—a cluster that characterises that parameter.
2. One cluster for each *parameter value*, or range of values.
3. One “controller” cluster, whose activation is linked to the activation of the parameters and their values in the following way: If the controller is active, each of its parameters and the accompanying

values are active. If a sufficient number of parameters and their values are active (this may be as few as one), the controller is active.

These are neural computational conditions on the networks we call “schemas.”

We have hesitated to call schemas “concepts,” simply because concepts have long been traditionally thought of as being direct reflections or representations of external reality. Schemas are clearly not that at all. Schemas are *interactional*, arising from (1) the nature of our bodies, (2) the nature of our brains, and (3) the nature of our social and physical *interactions* in the world. Schemas are therefore not purely internal, nor are they purely representations of external reality. We will, for the moment, think of concepts as schemas, though that idea will be extended below when we discuss abstractions.

The example of grasp

The example we have been using all the way through is the concept *grasp*. Here is what a schema for *grasp* might look like in this theory. The parameters divide up in the following ways:

The grasp schema.

1. The role parameters: agent, object, object location, and the action itself.
2. The phase parameters: initial condition, starting phase, central phase, purpose condition, ending phase, final state.
3. The manner parameter.
4. The parameter values (and constraints on them).

The various parameters can be described as follows.

Agent: An individual.

Object: A physical entity with parameters: size, shape, mass, degree of fragility, and so on.

Initial condition:² Object Location: Within peri-personal space.

Starting phase: Reaching, with direction: Toward object location; opening effector.

Central phase: Closing effector, with force: A function of fragility and mass.

Purpose condition: Effector encloses object, with manner (a grip determined by parameter values and situational conditions).

Final state: Agent in-control-of object.

This should give the reader a pretty clear idea of how a *grasp* schema is structured in terms of neural parameters and values of the sort we described in the sections above on neuroscience. Note that we have written down symbols (e.g., *final state*) as our notation for functional clusters. This does *not* mean that we take functional clusters themselves to be symbolic. The symbols are only *our names* for functional clusters, which, as we have seen, function from a computational point of view as neurally realised units.

A note about schemas

Traditionally, concepts were seen as a set of necessary and sufficient conditions operating in a system of logic. Indeed, for many philosophers, that was a defining characteristic of what a concept was to be.

It might look from the notation as though the *grasp* schema is indeed defined by such a set of necessary and sufficient conditions. This is not the case. First, the activation of functional clusters is not all-or-none; there are degrees of activation. Such gradations are not part of the traditional notion of necessary and sufficient conditions. Second, there are variations on schemas, as when certain phases are optionally left out. Third, there are extensions of schemas; for example, we will discuss metaphorical extensions below. Fourth, and perhaps most important, schemas combine and operate dynamically, in context, by neural optimisation—that is, via best fit principles. For example, imagine that you intend to grasp, pick up, and throw what appears to be a ball. But the ball turns out to be made of iron and to have a slippery surface. You

²The “:” notation indicates the content of a phase.

will grasp it as well as you can, though perhaps not exactly fitting the schema (tightening your grip might be difficult). You may manage to pick it up, but hardly in the normal way. And being slippery and very heavy, your attempt to throw it may result in something closer to a shot-put motion.

In short, schemas are not like logical conditions. They run bodies—as well as they can.

The theory we are outlining uses the computational modelling mechanisms of the neural theory of language (NTL) developed in Berkeley by the groups of Jerome Feldman and George Lakoff (see Feldman & Narayanan, in press; Lakoff & Johnson, 1999). NTL makes use of a *structured connectionist* version of neural computation (see Feldman, 1982), which, though “localist,” has units that are not just individual neurons, but rather *functional clusters* as we have discussed them throughout this paper. In such a structured connectionist model operating on functional clusters, the death or plasticity of individual neurons has virtually no effect, so long as the connectivity of the rest of the cluster remains intact.

NTL is therefore not subject to the “grandmother cell” objection, which assumes the following caricature of localist computation. In the caricature, each concept—say, the concept of your grandmother—is represented by one and only one neuron. If that neuron dies, then you lose the concept of your grandmother. No localist ever proposed such a theory, and nor do we.

From the structured connectionism perspective, *the inferential structure of concepts is a consequence of the network structure of the brain and its organisation in terms of functional clusters*. This brain organisation is, in turn, a consequence of our evolutionary history—of the way in which our brains, and the brains of our evolutionary ancestors, have been shaped by bodily interactions in the world.

WHAT WE HAVE SHOWN SO FAR

We have provided a reasonably detailed neural theory for one action concept—*grasping*. We have shown how the *sensory-motor* system can characterise a sensory-motor *concept*, not just an action or a perception, but a *concept* with all that that requires.

But we think we have shown something more powerful than that. According to our hypothesis, understanding requires simulation. The understanding of concrete concepts—physical actions, physical objects, and so on—requires sensory-motor simulation. But sensory-motor simulation, as suggested by contemporary neuroscience, is carried out by the sensory-motor system of the brain. It follows that the sensory-motor system *is required for understanding at least concrete concepts*. We see this as an insurmountable difficulty for any traditional theory that claims that concrete concepts are modality-neutral and disembodied.

There is a further argument against the traditional modality-neutral, disembodied account of concepts. In order to have a neural account of such a theory of action concepts, action *concepts*, like all other concepts, would have to be represented neurally *outside the sensory-motor system altogether*. That, in turn, would require complete duplication of the structure-characterising concepts that neuroscience has found *in the sensory-motor system*, namely, all the structure we have just outlined: *the manner subcases, the agent-object-location structure, the purpose structure, and the phase structure*.

The reason is this: All of that sensory-motor structure (agent-object-location, manner, purpose, and phases) has to be there, in anyone’s account. Any neural theory of modality-neutral concepts must claim that such structure is located in the brain *outside the sensory-motor system*. But we know, from independent evidence that we have cited, that all that structure is indeed *inside the sensory-motor system*. The only way it could also be outside is if it were duplicated. Not just for one concept, but for *every action concept*. And it would not just have to be duplicated. There would have to be one-to-one connections to just the right parts of the premotor-parietal system in order for the concept to apply to real cases that are performed, observed, and simulated.

In short, we think there is an Occam’s Razor argument here. The modality-neutral structure is just not needed.

We think we can safely say that something like the proposal we have made would have to work for the action concept of grasping, and probably for

every other action concept as well. We believe that the same basic structures—schemas structuring sensory-motor parameterisations—can be used to characterise all concrete concepts. Take, for example, the basic-level concepts that we described above—chair, car, etc. As we saw, basic-level concepts are defined by the convergence of (1) gestalt object perception (observed or imaged) and (2) motor programmes that define the prototypical interaction with the object (again, performed or imaged). Thus, a chair looks a certain way (and we can imagine how it looks) and it is used for sitting in a certain position (and we can imagine sitting that way). What brings together the perceptual and motor properties are, we believe, functional neural clusters showing the analogous characteristics in humans that have been found for *canonical neurons* thus far in the brains of monkeys. Indeed, evidence suggesting the presence of the equivalent of canonical neurons in humans does exist, as we mentioned above (Chao & Martin, 2000; Grafton et al., 1996; Martin et al., 1996; Perani et al., 1995). The existence of canonical neurons and their putative equivalent in humans could underpin basic-level categories of objects. Subordinate-level categories, which *have more perceptual and motor details filled in*, would be accounted for by functional clusters of the type described before, with the values of more parameters specified. Schemas structuring the parameters over functional clusters would have the properties of basic-level and subordinate concrete object concepts.

We believe that *all* concrete concepts—concepts of things we can see, touch, and manipulate—can be addressed by the strategy outlined so far. Indeed, our proposal for object schemas is very close to the “sensory/motor model of semantic representations of objects” of Martin, Ungerleider, and Haxby (2000), which is based on a thorough survey of neuroscientific evidence.

THE LARGER CONTEXT

Neuroscience does not exist in a vacuum; it is part of the larger field of cognitive science, which has branches like neural computation and cognitive

linguistics. The results we have just summarised fit into the larger context. Within the field of structured connectionist neural modelling, Srinivas Narayanan (1997, 1999) has constructed computational neural models of motor actions, including the tripartite breakdown: premotor, motor, and premotor-motor connections. The premotor model functioned dynamically to “choreograph” and carry out in proper sequence the simple movements of the motor cortex.

These premotor models turned out to have a uniform structure: (1) initial state, (2) starting phase transition, (3) precentral state, (4) central phase transition (either instantaneous, prolonged, or ongoing), (5) postcentral state, (6) ending phase transition, (7) final state. At the postcentral state, there are the following options: (a) a check to see if a goal state has been achieved, (b) an option to iterate or continue the main process, (c) an option to stop, and (d) an option to resume. Each complex motor programme is a complex combination of structures of this form, either in sequence, in parallel, or embedded one in another. What distinguishes actions from one another is (1) the version of this premotor structure and (2) bindings to the motor cortex and other sensory areas (for perceptual and somatosensory feedback). These premotor structures are called “executing schemas,” or X-schemas for short. Narayanan (1997) noted that premotor structures also fit the perceptual structure of the motor actions modelled. In short, he modelled the structures described above for mirror neurons, canonical neurons, and action-location neurons.

Dynamic X-schemas can be seen as linking the functional clusters we have just discussed to characterise their temporal activation in either action or perception—or in imagination. In short, they are capable of carrying out imaginative simulations. Furthermore, those imaginative simulations can carry out *abstract conceptual reasoning* as well as actions and perceptions. The result is a neural theory of conceptual metaphor. We know from cognitive semantics that conceptual metaphors are one of the basic mechanisms of mind. Each conceptual metaphor is a mapping across conceptual domains, from a (typically) sensory-motor source domain to a (typically) non-sensory-motor target domain.

Q11

For example, the conceptual metaphor *love is a journey* maps travelers to lovers, vehicles to relationships, destinations to common life goals, and impediments to relationship difficulties, as shown by English expressions about love like *It's been a long bumpy road*, *The marriage is on the rocks*, *We're spinning our wheels*, *We're going in different directions*, *We're at a crossroads in the relationship*, and so on. The concept of love has a minimal nonmetaphorical structure with a lover, a beloved, a love relationship, and not much more. More than a dozen conceptual metaphors of this sort add to that minimal structure a very rich conceptual structure. Henceforth love can be conceptualised and reasoned about in terms not only of a journey, but also of a partnership, a joining-together, magic, heat, and so on. Most (not all) of the richness of the concept comes from these metaphors. (For detailed discussions of such metaphors and the evidence for them, see Kövecses 2002; Lakoff & Johnson, 1980, 1999; and the references discussed in these works.)

What Narayanan (1997) did was to construct a computational neural model of such metaphorical mappings, in which each mapping is carried out by neural circuitry of certain regular structures. He then chose an abstract domain—international economics—and worked out the conceptual metaphors mapping physical actions to economics. He constructed computational neural models of both target and source domains and took sentences from such sources as the *NY Times Business Section* and the *Wall Street Journal*—sentences like *France fell into a recession; Pulled out by Germany*, and *India is stumbling toward economic liberalisation*, in which there are physical sensory-motor expressions like *fall into*, *pull out*, and *stumble toward*. He then showed that by using the mappings to combine source (sensory-motor) and target (economic) inferences, he could get the correct inferences in a neural computational simulation.

The same computational models of neural circuitry that can direct action and perceptions of actions can also simulate actions with the right structures to get all the conceptual inferences right. The sensory-motor system can characterise action concepts and, in simulation, characterise

conceptual inferences. And the concepts characterised in the sensory-motor system are of the right form to characterise the source domains of conceptual metaphors.

THE THEORY OF COGS

Narayanan (1997) made another discovery. His premotor X-schemas have exactly the right structure to characterise the collection of concepts that linguists refer to as “aspect”—concepts that characterise the structure of events and our reasoning about events. Every language has a way of indicating aspect. English has terms like *about to*, *start to*, *be + Verb + ing*, and *have + Verb + Past Participle*. Thus, *he is about to run*, *he is starting to run*, *he is running*, *he has run*.

In actions, the premotor cortex is neurally connected to the motor cortex, choreographing simple movements into complex actions. But those premotor-to-motor connections can be inhibited, and the X-schemas of the premotor system can function independently, characterising the logic of aspect in the abstract. Thus, a sentence like *he is doing something stupid* doesn't tell what action he is carrying out, but does specify the ongoing aspectual structure, with the inferences that *he has already started doing something stupid* and *he hasn't finished doing something stupid*. These inferences are being computed via neural simulation by X-schema structure circuitry in the premotor cortex—with no active connections to the motor cortex. In short, a portion of the sensory-motor system (the premotor cortex) is being used to do abstract reasoning, reasoning that is not about any particular sensory-motor activity. Indeed, *the stupid thing he is doing* need not be physical at all.

Let us call the premotor cortex a “secondary” area—an area not directly connected to sensors or effectors, but which provides structure to information going to effectors or coming from sensors. On Narayanan's hypothesis, abstract aspectual concepts have the following properties:

- They are neurally simulated in a secondary area with no active connections to a primary

area. For example, all such connections may be inhibited.

- Their inferences are computed via that simulation.
- They characterise concepts in the grammar of a natural language.
- As such, these concepts are general, and can apply to any special-case concepts when there are active connections to primary areas.
- When there are such active connections, the general concepts are an inseparable part of the structure of the special case concepts.

Lakoff (personal communication) has proposed a generalisation of Narayanan's account of aspect to include all concepts with such properties. Any concept that has the properties given in the above list is called a *cog*. In the theory of cogs, all concepts in the grammars (as opposed to the lexicons) of natural languages should have the properties given in the list above. That is, they should be computed in secondary areas. This could include all the primitive image-schemas, such as containment, source-path-goal, force dynamics, orientation schemas, etc. (see Casad & Langacker, 1985; Johnson, 1987; Lakoff, 1987; Langacker, 1986, 1990, 1991; Lindner, 1981; Talmy, 1983, 1988, 1996, 1999).

For example, according to Regier's (1996) neural computational theory of container schemas, containment is computed via a complex of topographical maps of the visual field. The computation is general; it can take any size or shape of container—a bottle, a cup, a room, a castle, a triangle—as input and characterise its interior, boundary, and exterior. But that input information (including shape) is not computed in topographic map areas. It is computed elsewhere, with input presumably coming via the parietal cortex. The complex of topographic maps computing containment constitutes a “secondary” area, while the input to it would come via an area that is more “primary.”

What we know about image-schemas is that they characterise a wide range of general inference patterns, characterising all forms of causal, spatial, and event-based reasoning. They are universal, general, and appear in the semantics of grammar

around the world. If they are all computed in secondary regions, that would explain why there is a limited range of them (there is a relatively small number of such regions), why they are universal (we all have the same basic brain structure), and why they are general (they provide structure to primary regions with specific information).

The theory of cogs is at present still vague. It does not specify exactly which areas compute which cogs using which circuitry. Its utility, however, lies in that: (1) it gives neuroscientists something to look for that they might not have thought to look for; (2) it provides a possible explanation for why image-schemas exist and why there is such a thing as the semantics of grammar; (3) it posits an interesting theory of the learning of general linguistic structures. Briefly, here is what the theory says:

- Because neural structures in secondary areas are inseparable in behaviour from the primary structures that they are connected to, they characterise generalisations that are inherent in and inseparable from special cases.
- The “learning” of general cases is not the acquisition of new structures, but rather the inhibition of the connections between secondary and primary areas.
- In other words, the generalisations are inherent in the special cases that are learned first. What is learned is the control of inhibitory connections.

EMPIRICAL VALIDATION OF OUR THEORY: FUTURE EXPERIMENTS

The theory of concepts that we have outlined here has empirical consequences that can be tested. In this theory, the same circuitry that can move the body and structure perceptions, also structures abstract thought. Of course, this does not imply that a paresis or sensory loss due to a localised brain damage affecting selected regions within the sensory-motor system should *necessarily* produce a deficit in the capacity to entertain abstract conceptual knowledge. This possibility is not, in principle, precluded, but is unlikely to arise. The reason

is this: The distributed nature of the sensory-motor neural clusters responsible for structuring conceptual knowledge spans from frontal to parieto-temporal cortices. Because the operative neural clusters are distributed in this way, it is unlikely that a restricted brain lesion would be sufficient to impair their ability to function.

It should be added that standard neuropsychological testing for conceptual knowledge does not routinely investigate metaphorical conceptualisation in brain-lesioned patients. A systematic study of patients along these lines should be encouraged. If, however, a hypothetical brain lesion would prevent a patient from grasping objects, recognising other people doing it, or even imagining himself or others grasping objects, then our theory predicts that the same patient should also be impaired in metaphorically conceptualising grasping as understanding. We are not aware of any clinical report of specific and localised brain lesions (within or outside the sensory-motor system) causing the selective loss of *specific abstract concepts* such as causation, identity, love, and the like. In fact, at the core of our theory is the claim that there are no dedicated and specialised brain circuits for concepts in general, or for abstract concepts in particular.

The difficulty in relying on clinical cases, though, does not mean that our theory cannot generate empirically testable predictions. We know, for example, that the motor cortex is topographically organised, with the movements of effectors like the mouth, hands, and feet controlled in different sectors. These sectors are separated far enough from each other in the motor cortex so that they are clearly distinguishable in fMRI studies. The same holds for the premotor cortex, though with more overlap between regions. Moreover, premotor activations are clearly distinguishable from motor activations and from activations in other regions. We know that both premotor and motor cortices are both activated during mental motor imagery. According to our hypothesis, mental motor imagery is voluntary simulation of motor action (see Gallese, 2003 a, b).

In the theory we have outlined, the properties of action concepts are specified by parameter values and the simulations they govern. Such

action simulations should be detectable through fMRI studies in the sensory-motor cortices that are responsible for controlling the corresponding actions. For example, in the case of the concept of *grasping*, one would expect the parietal-premotor circuits that form functional clusters for grasping to be active not only when actually grasping, but also when understanding sentences involving the concept of *grasping*. Moreover, when processing sentences describing actions performed by different effectors—the mouth in biting, the hands in holding, the feet in kicking—one would expect parietal-premotor regions for action by the mouth vs. hands vs. feet to be active not only when actually acting, but also when understanding the corresponding sentences. Preliminary evidence seems to confirm that this is the case (Tettamanti et al., in press)

A further prediction of our theory of concepts is that such results should be obtained in fMRI studies, not only with literal sentences, but also with the corresponding metaphorical sentences. Thus, the sentence *He grasped the idea* should activate the sensory-motor *grasping-related regions* of the brain. Similarly, a metaphorical sentence like *They kicked him out of class* should activate the sensory-motor *kicking-related regions* of the brain. A series of brain-imaging experiments are currently being carried out to test this prediction.

Another series of experiments is being designed to test the parameters of force and direction in concepts. On the grounds of our proposal, we would predict that asking subjects to exert a force while reading sentences about force, or to move their hand along a given direction while reading sentences about directional actions, would either facilitate or interfere with the comprehension of such sentences. Indeed, Glenberg and Kaschak (2002) recently showed an action-sentence compatibility effect in normal subjects required to determine the correctness of read sentences describing concrete or abstract actions towards or away from the body, by moving their hand either towards or away from the body.

Such an experimental design could also be used to study metaphorical sentences, in which there is a metaphorical use of force or direction predicates (e.g., throw, bring, knock). The same logic should

apply here as in the literal sentences. The application of force, or movement exerted along a given direction, should facilitate or interfere with the understanding of metaphorical sentences, depending on whether the force or direction in the source domain of the metaphor is or is not consistent with magnitude of the force applied or with the direction of movement.

CONCLUSIONS

We have argued that contemporary neuroscience seems to suggest that concepts of a wide variety make direct use of the sensory-motor circuitry of the brain. We began the argument with action concepts and with four central ideas triggered by neuroscience: multimodality, functional clusters, simulation, and parameters. We then turned to neuroscientific results: (1) visual and motor mental imagery, which, according to our hypothesis, imply sensory-motor simulation using the same brain resources as in observation and action; (2) detailed results concerning mirror neurons, canonical neurons, and action-location neurons. By applying the four ideas to these results, we proposed, for the action concept of *grasping*, that a directly embodied schema for grasping satisfies all principal criteria for concepts. We argued that a disembodied, symbolic account of the concept of *grasping* would have to duplicate elsewhere in the brain the complex neural machinery in three parietal-premotor circuits, which is implausible to say the least. We concluded that the action *concept* of *grasping* is embodied in the sensory-motor system.

We then went on to argue that arguments of the same form may apply to all other action concepts, to object concepts, and to abstract concepts with conceptual content that is metaphorical. Finally, we considered cogs, which we posit to be structuring circuits in the sensory-motor system, which normally function as part of sensory-motor operations, but whose neural connections to specific details can be inhibited, allowing them to provide inferential structure to "abstract" concepts. If all this is correct, then abstract reasoning in general *exploits* the sensory-motor system.

What is the import of our proposal on the nature of human cognition? We believe that it suggests that rational thought is not entirely separate from what animals can do, because it directly uses sensory-motor bodily mechanisms—the same ones used by nonhuman primates to function in their everyday environments. According to our hypothesis, rational thought is an exploitation of the normal operations of our bodies. As such, it is also largely unconscious.

Another major consequence concerns language. Language makes use of concepts. Concepts are what words, morphemes, and grammatical constructions express. Indeed, the expression of concepts is primarily what language is about. If we are right, then:

1. Language makes direct use of the same brain structures used in perception and action.
2. Language is not completely a human innovation.
3. There is no such thing as a "language module."
4. Grammar resides in the neural connections between concepts and their expression via phonology. That is, grammar is constituted by the connections between conceptual schemas and phonological schemas. Hierarchical grammatical structure is conceptual structure. Linear grammatical structure is phonological.
5. The semantics of grammar is constituted by cogs—structuring circuits used in the sensory motor system.
6. Neither semantics nor grammar is modality-neutral.
7. Neither semantics nor grammar is symbolic, in the sense of the theory of formal systems, which consists of rules for manipulating disembodied meaningless symbols.

Future empirical research will show whether and to what extent our hypotheses are correct.

REFERENCES

- Albritton, D. (1992). *The use of metaphor to structure text representations: evidence for metaphor-based schemas*. PhD dissertation, Yale University.

- Bailey, D. (1997). *A computational model of embodiment in the acquisition of action verbs*. PhD dissertation. Computer Science Division, EECS Department, University of California, Berkeley.
- Berlin, B., Breedlove, D., & Raven, P. (1974). *Principles of Tzeltal plant classification*. New York: Academic Press.
- Berlin, B., & Kay, P. (1969). *Basic color terms: Their universality and evolution*. Los Angeles: University of California Press.
- Boroditsky, L. (1997). Evidence for metaphoric representation: Perspective in space and time. In M. G. Shafto and P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates Inc.
- Bremner, F., Schlack, A., Jon Shah, N., Zafiris, O., Kubischik, M., Hoffmann, K. P., Zilles, K., & Fink, G. R. (2001). Polymodal motion processing in posterior parietal and premotor cortex: A human fMRI study strongly implies equivalences between humans and monkeys. *Neuron*, *29*, 287–296.
- Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., Gallese, V., Seitz, R. J., Zilles, K., Rizzolatti, G., & Freund, H.-J. (2001). Action observation activates premotor and parietal areas in a somatotopic manner: An fMRI study. *European Journal of Neuroscience*, *13*, 400–404.
- Carpenter, P. A., & Eisenberg, P. (1978). Mental rotation and the frame of reference in blind and sighted individuals. *Perception and Psychophysics*, *23*, 117–124.
- Casad, E., & Langacker, R. W. (1985). “Inside” and “outside” in Cora grammar. *International Journal of American Linguistics*, *51*, 247–281.
- Chao, L. L., & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *Neuroimage*, *12*, 478–484.
- Cochin, S., Barthelemy, C., Lejeune, B., Roux, S., & Martineau, J. (1998). Perception of motion and qEEG activity in human adults. *Electroencephalography and Clinical Neurophysiology*, *107*, 287–295.
- Colby, C. L., Duhamel, J.-R., & Goldberg M. E. (1993). Ventral intraparietal area of the macaque: anatomic location and visual response properties. *Journal of Neurophysiology*, *69*, 902–914.
- Decety, J., & Grèzes, J. (1999). Neural mechanisms subserving the perception of human actions. *Trends in Cognitive Sciences*, *3*, 172–178.
- Decety, J., Grezes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., Grassi, F., & Fazio, F. (1997). Brain activity during observation of actions. Influence of action content and subject’s strategy. *Brain*, *120*, 1763–1777.
- Decety, J., Jeannerod, M., Germain, M., & Pastene, J. (1991). Vegetative response during imagined movement is proportional to mental effort. *Behavioral Brain Research*, *34*, 35–42.
- Decety, J., Jeannerod, M., & Prablanc, C. (1989). The timing of mentally represented actions. *Behavioral Brain Research*, *34*, 35–42.
- Decety, J., Sjöholm, H., Ryding, E., Stenberg, G., & Ingvar, D. (1990). The cerebellum participates in cognitive activity: Tomographic measurements of regional cerebral blood flow. *Brain Research*, *535*, 313–317.
- DeValois, R. L., & DeValois, K. (1975). Neural coding of color. In E. C. Careterette & M. P. Friedman (Eds.), *Handbook of perception, Vol. V, Seeing*. New York: Academic Press.
- Duffy, J. R., & Watkins, L. B. (1984). The effect of response choice relatedness on pantomime and verbal recognition ability in aphasic patients. *Brain and Language*, *21*, 291–306.
- Duhamel, J.-R., Colby, C. L., & Goldberg, M. E. (1998). Ventral intraparietal area of the macaque: Congruent visual and somatic response properties. *Journal of Neurophysiology*, *79*, 126–136.
- Ehrsson, H. H., Fagergren, A., Jonsson, T., Westling, G., Johansson, R. S., & Forssberg H. (2000). Cortical activity in precision- versus power-grip tasks: An fMRI study. *Journal of Neurophysiology*, *83*, 528–536.
- Fadiga, L., Fogassi, L., Pavesi, G., & Rizzolatti, G. (1995). Motor facilitation during action observation: A magnetic stimulation study. *Journal of Neurophysiology*, *73*, 2608–2611.
- Farah, M. J. (1989). The neural basis of mental imagery. *Trends in Neuroscience*, *12*, 395–399.
- Farah, M. J. (2000). The neural bases of mental imagery. In m. s. Gazzaniga (Ed.), *The cognitive neurosciences (2nd ed.)*. Cambridge, MA: MIT Press.
- Feldman, J. (1982). Dynamic connections in neural networks. *Biological Cybernetics*, *46*, 27–39.
- Feldman, J., & Narayanan, S. (in press). Embodied meaning in a neural theory of language. *Brain and Language*.
- Fillmore, C. (1982). Frame semantics. In Linguistic Society of Korea (Ed.), *Linguistics in the morning calm* (pp. 111–138). Seoul: Hanshin.
- Fillmore, C. (1985). Frames and the semantics of understanding. *Quaderni di Semantica*, *6*, 222–253.
- Fogassi, L., & Gallese, V. (2004). Action as a binding key to multisensory integration. In G. Calvert, C. Spence, & B. E. Stein (Eds.), *Handbook of multisensory processes*. Cambridge: MIT Press.

- Fogassi, L., Gallese, V., Buccino, G., Craighero, L., Fadiga, L., & Rizzolatti, G. (2001). Cortical mechanism for the visual guidance of hand grasping movements in the monkey: A reversible inactivation study. *Brain*, *124*, 571–586.
- Fogassi, L., Gallese, V., di Pellegrino, G., Fadiga, L., Gentilucci, M., Luppino, G., Matelli, M., Pedotti, A., & Rizzolatti, G. (1992). Space coding by premotor cortex. *Experimental Brain Research*, *89*, 686–690.
- Fogassi, L., Gallese, V., Fadiga, L., Luppino, G., Matelli, M., & Rizzolatti, G. (1996a). Coding of peripersonal space in inferior premotor cortex (area F4). *Journal of Neurophysiology*, *76*, 141–157.
- Fogassi, L., Gallese, V., Fadiga, L., & Rizzolatti, G. (1996b). Space coding in inferior premotor cortex (area F4): Facts and speculations. In F. Laquaniti & P. Viviani (Eds.), *Neural basis of motor behavior* (pp. 99–120). *NATO ASI Series*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Fox, P., Pardo, J., Petersen, S., & Raichle, M. (1987). Supplementary motor and premotor responses to actual and imagined hand movements with Positron Emission Tomography. *Society for Neuroscience Abstracts*, *13*, 1433.
- Gainotti, G. (2004). A metanalysis of impaired and spared naming for different categories of knowledge in patients with a visuo-verbal disconnection. *Neuropsychologia*, *42*, 299–319.
- Gallese, V. (2000a). The acting subject: Towards the neural basis of social cognition. In T. Metzinger (Ed.), *Neural correlates of consciousness. Empirical and conceptual questions* (pp. 325–333). Cambridge, MA: MIT Press.
- Gallese, V. (2000b). The inner sense of action: Agency and motor representations. *Journal of Consciousness Studies*, *7*, 23–40.
- Gallese, V. (2001). The “shared manifold” hypothesis: From mirror neurons to empathy. *Journal of Consciousness Studies*, *8*, 33–50.
- Gallese, V. (2003a). The manifold nature of interpersonal relations: The quest for a common mechanism. *Philosophical Transactions of the Royal Society of London, B*, *358*, 517–528.
- Gallese, V. (2003b). A neuroscientific grasp of concepts: From control to representation. *Philosophical Transactions of the Royal Society of London, B*, *358*, 1231–1240.
- Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, *119*, 593–609.
- Gallese, V., Fogassi, L., Fadiga, L., & Rizzolatti, G. (2002). Action representation and the inferior parietal lobule. In W. Prinz & B. Hommel (Eds.), *Attention and performance XIX* (pp. 247–266). Oxford: Oxford University Press.
- Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, *12*, 493–501.
- Gallese, V., Murata, A., Kaseda, M., Niki, N., & Sakata, H. (1994). Deficit of hand preshaping after muscimol injection in monkey parietal cortex. *NeuroReport*, *5*, 1525–1529.
- Gentilucci, M., Fogassi, L., Luppino, G., Matelli, M., Camarda, R., & 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*, 475–490.
- Gentilucci, M., Scandolara, C., Pigarev, I. N., & Rizzolatti, G. (1983). Visual responses in the postarcuate cortex (area 6) of the monkey that are independent of eye position. *Experimental Brain Research*, *50*, 464–468.
- Gentner, D., & Gentner, D. R. (1982). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.), *Mental models*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*, 1416–1419.
- Gibbs, R. (1994). *The poetics of mind: Figurative thought, language, and understanding*. Cambridge: Cambridge University Press.
- Grady, J. (1997). *Foundations of meaning: Primary metaphors and primary scenes*. PhD dissertation, University of California at Berkeley.
- Grafton, S. T., Arbib, M. A., Fadiga, L., & Rizzolatti, G. (1996). Localization of grasp representations in humans by PET: 2. Observation compared with imagination. *Experimental Brain Research*, *112*, 103–111.
- Grafton, S. T., Fadiga, L., Arbib, M. A., & Rizzolatti, G. (1997). Premotor cortex activation during observation and naming of familiar tools. *Neuroimage*, *6*, 231–236.
- Graziano, M. S. A., & Gross, C. G. (1995). The representation of extrapersonal space: A possible role for bimodal visual-tactile neurons. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1021–1034). Cambridge, MA: MIT Press.
- Graziano, M. S. A., Hu, X., & Gross, C. G. (1997a). Visuo-spatial properties of ventral premotor cortex. *Journal of Neurophysiology*, *77*, 2268–2292.

- Graziano, M. S. A., Hu, X., & Gross, C. G. (1997b). Coding the locations of objects in the dark. *Science*, 277, 239–241.
- Graziano, M. S. A., Reiss, L. A. J., & Gross, C. G. (1999). A neuronal representation of the location of nearby sounds. *Nature*, 397, 428–430.
- Graziano, M. S. A., Yap, G. S., & Gross, C. G. (1994). Coding of visual space by premotor neurons. *Science*, 266, 1054–1057.
- Grillner, S., & Wallen, P. (2002). Cellular bases of a vertebrate locomotor system—steering, intersegmental and segmental co-ordination and sensory control. *Brain Research Review*, 40, 92–106.
- Hari, R., Forss, N., Avikainen, S., Kirveskari, S., Salenius, S., & Rizzolatti, G. (1998). Activation of human primary motor cortex during action observation: A neuromagnetic study. *Proceedings of the National Academy of Sciences USA*, 95, 15061–15065.
- Heilman, K. M., & Rothi, L. J. (1993). Apraxia. In K. M. Heilman and E. Valenstein (Eds.), *Clinical neuropsychology* (pp. 141–163, 3rd ed.). New York: Oxford University Press.
- Heilman, K. M., & Rothi, L. J., & Valenstein, E. (1982). Two forms of ideomotor apraxia. *Neurology*, 32, 342–346.
- Hepp-Reymond, M.-C., Hüsler, E. J., Maier, M. A., & Qi, H.-X. (1994). Force-related neuronal activity in two regions of the primate ventral premotor cortex. *Canadian Journal of Physiology and Pharmacology*, 72, 571–579.
- Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical mechanisms of human imitation. *Science*, 286, 2526–2528.
- Jackson, P. L., & Decety, J. (2004). Motor cognition: A new paradigm to study self-other interactions. *Current Opinions in Neurobiology*, 14, 259–263.
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral Brain Science*, 17, 187–245.
- Jeannerod, M., Arbib, M. A., Rizzolatti, G., & Sakata, H. (1995). Grasping objects: The cortical mechanisms of visuomotor transformation. *Trends in Neuroscience*, 18, 314–320.
- Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination and reason*. Chicago: University of Chicago Press.
- Kay, P., & McDaniel, C. (1978). The linguistic significance of the meanings of basic color terms. *Language*, 54, 610–646.
- Kemper, S. (1989). Priming the comprehension of metaphors. *Metaphor and Symbolic Activity*, 4, 1–18.
- Kerr, N. H. (1983). The role of vision in “visual imagery” experiments: Evidence from the congenitally blind. *Journal of Experimental Psychology: General*, 112, 265–277.
- Kohler, E., Keysers, C., Umiltà, M. A., Fogassi, L., Gallese, V., & Rizzolatti, G. (2002). Hearing sounds, understanding actions: Action representation in mirror neurons. *Science*, 297, 846–848.
- Kohler, E., Umiltà, M. A., Keysers, C., Gallese, V., Fogassi, L., & Rizzolatti, G. (2001). Auditory mirror neurons in the ventral premotor cortex of the monkey. *Society for Neuroscience Abstracts*, 27.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Kosslyn, S. M., Alpert, N. M., Thompson, W. L., Maljkovic, V., Weise, S., Chabris, C., Hamilton, S. E., Rauch, S. L., & Buonanno, F. S. (1993). Visual mental imagery activates topographically organized visual cortex: PET investigations. *Journal of Cognitive Neuroscience*, 5, 263–287.
- Kosslyn, S. M., Ball, T. M., & Reiser, B. J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 47–60.
- Kosslyn, S. M., & Thompson, W. L. (2000). Shared mechanisms in visual imagery and visual perception: Insights from cognitive science. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (2nd ed.). Cambridge, MA: MIT Press.
- Kövecses, Z. (2002). *Metaphor: A practical introduction*. Oxford: Oxford University Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh*. New York: Basic Books.
- Lakoff, G., & Turner, M. (1989). *More than cool reason: A field guide to poetic metaphor*. Chicago: University of Chicago Press.
- Langacker, R. (1983). Remarks on English aspect. In P. Hopper (Ed.), *Tense and aspect: Between semantics and pragmatics* (pp. 265–304). Amsterdam: John Benjamins.
- Langacker, R. W. (1986). *Foundations of cognitive grammar, vol. 1*. Stanford, CT: Stanford University Press.

- Langacker, R. W. (1990). *Concept, image, and symbol: The cognitive basis of grammar*. Berlin: Mouton de Gruyter.
- Langacker, R. W. (1991). *Foundations of cognitive grammar, vol. 2*. Stanford, CT: Stanford University Press.
- Le Bihan, D., Turner, R., Zeffiro, T. A., Cuenod, C. A., Jezzard, P., & Bonnerot, V. (1993). Activation of human primary visual cortex during visual recall: A magnetic resonance imaging study. *Proceedings of the National Academy of Sciences*, *90*, 11802–11805.
- Lindner, S. (1981). *A lexico-semantic analysis of verb-particle constructions with up and out*. PhD dissertation, University of California at San Diego.
- Luppino, G., Murata, A., Govoni, P., & Matelli, M. (1999). Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas AIP and VIP) and the ventral premotor cortex (areas F5 and F4). *Experimental Brain Research*, *128*, 181–187.
- Malach, R., Levy, I., and Hasson, U. (2002). The topography of high-order human object areas. *Trends in Cognitive Sciences*, *6*, 176–184.
- Marmor, G. S., & Zaback, L. A. (1976). Mental rotation by the blind: Does mental rotation depend on visual imagery? *Journal of Experimental Psychology: Human Perception and Performance*, *2*, 515–521.
- Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion in Neurobiology*, *11*, 194–201.
- Martin, A., Ungerleider, L. G., & Haxby, J. V. (2000). Category specificity and the brain: The sensory/motor model of semantic representations of objects. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 1023–1036, 2nd ed.). A Bradford book. Cambridge, MA: MIT Press.
- Matelli, M., Camarda, R., Glickstein, M., & Rizzolatti, G. (1986). Afferent and efferent projections of the inferior area 6 in the macaque monkey. *Journal of Comparative Neurology*, *251*, 281–298.
- Matelli, M., & Luppino, G. (1997). Functional anatomy of human motor cortical areas. In F. Boller, & J. Grafman (Eds.), *Handbook of neuropsychology, Vol. 11*. Amsterdam: Elsevier Science.
- Matelli, M., Luppino, G., & Rizzolatti, G. (1985). Patterns of cytochrome oxidase activity in the frontal agranular cortex of the macaque monkey. *Behavioral Brain Research*, *18*, 125–137.
- Mervis, C. (1986). Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 201–233). Cambridge: Cambridge University Press.
- Mervis, C., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, *32*, 89–115.
- Mountcastle, V. B. (1995). The parietal system and some higher brain functions. *Cerebral Cortex*, *5*, 377–390.
- Murata, A., Fadiga, L., Fogassi, L., Gallese, V., Raos, V., & Rizzolatti, G. (1997). Object representation in the ventral premotor cortex (Area F5) of the monkey. *Journal of Neurophysiology*, *78*, 2226–2230.
- Murata, A., Gallese, V., Luppino, G., Kaseda, M., & Sakata, H. (2000). Selectivity for the shape, size and orientation of objects in the hand-manipulation-related neurons in the anterior intraparietal (AIP) area of the macaque. *Journal of Neurophysiology*, *83*, 2580–2601.
- Narayanan, S. (1997). *KARMA: Knowledge-based Active Representations for Metaphor and Aspect*. PhD dissertation, Computer Science Division, University of California, Berkeley (<http://www.icsi.berkeley.edu/~snarayan/thesis.pdf>).
- Narayanan, S. (1999). Moving right along: A computational model of metaphoric reasoning about events. *Proceedings of the National Conference on Artificial Intelligence AAAI-99*, Orlando, FL (<http://www.icsi.berkeley.edu/~snarayan/met.ps>).
- Nishitani, N., & Hari, R. (2000). Temporal dynamics of cortical representation for action. *Proceedings of the National Academy of Sciences*, *97*, 913–918.
- Parsons, L., Fox, P., Downs, J., Glass, T., Hirsch, T., Martin, C., Jerabek, J., & Lancaster, J. (1995). Use of implicit motor imagery for visual shape discrimination as revealed by PET. *Nature*, *375*, 54–58.
- Perani, D., Schnur, T., Tettamanti, M., Gorno-Tempini, M., Cappa, S. F., & Fazio, F. (1999). Word and picture matching: A PET study of semantic category effects. *Neuropsychologia*, *37*, 293–306.
- Porro, C. A., Francescato, M. P., Cettolo, V., Diamond, M. E., Baraldi, P., Zuiani, C., Bazzocchi, M., & di Prampero, P. E. (1996). Primary motor and sensory cortex activation during motor performance and motor imagery. A functional magnetic resonance study. *Journal of Neuroscience*, *16*, 7688–7698.
- Porter, R., & Lemon, R. N. (1993). *Corticospinal function and voluntary movement*. Oxford: Clarendon Press.
- Regier, T. (1995). A model of the human capacity for categorizing spatial relations. *Cognitive Linguistics*, *6*–1, 63–88.
- Regier, T. (1996). *The human semantic potential: Spatial language and constrained connectionism*. Cambridge, MA: MIT Press.

- Rizzolatti, G., Berti, A., & Gallese, V. (2000a). Spatial neglect: Neurophysiological bases, cortical circuits and theories. In F. Boller, J. Grafman, & G. Rizzolatti (Eds.), *Handbook of neuropsychology, Vol. I* (pp. 503–537, 2nd ed.). Amsterdam: Elsevier Science.
- Rizzolatti, G., Camarda, R., Fogassi, M., Gentilucci, M., Luppino, G., & 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*, 491–507.
- Rizzolatti, G., Fadiga, L., Fogassi, L., & Gallese, V. (1997). The space around us. *Science*, *277*, 190–191.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, *3*, 131–141.
- Rizzolatti, G., Fogassi, L., & Gallese, V. (2000b). Cortical mechanisms subserving object grasping and action recognition: A new view on the cortical motor functions. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 539–552, 2nd ed.). Cambridge, MA: MIT Press.
- Rizzolatti, G., Fogassi, L., & Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Neuroscience Reviews*, *2*, 661–670.
- Rizzolatti, G., Fogassi, L., & Gallese, V. (2002). Motor and cognitive functions of the ventral premotor cortex. *Current Opinion in Neurobiology*, *12*, 149–154.
- Rizzolatti, G., & Gallese, V. (2004). Do perception and action result from different brain circuits? The three visual systems hypothesis. In L. van Hemmen & T. Sejnowski (Eds.), *Problems in systems neuroscience*. Oxford: Oxford University Press.
- Rizzolatti, G., & Gentilucci, M. (1988). Motor and visuo-motor functions of the premotor cortex. In P. Rakic & W. Singer (Eds.), *Neurobiology of neocortex* (pp. 269–284). Chichester, UK: Wiley.
- Rizzolatti, G., Luppino, G., & Matelli, M. (1998). The organization of the cortical motor system: New concepts. *Electroencephalography and Clinical Neurophysiology*, *106*, 283–296.
- Rizzolatti, G., Matelli, M., & Pavesi, G. (1983). Deficits in attention and movement following the removal of postarcuate (area 6) and prearcuate (area 8) cortex in macaque monkeys. *Brain*, *106*, 655–673.
- Rizzolatti, G., Scandolara, C., Gentilucci, M., & Camarda, R. (1981a). Response properties and behavioral modulation of “mouth” neurons of the postarcuate cortex (area 6) in macaque monkeys. *Brain Research*, *255*, 421–424.
- Rizzolatti, G., Scandolara, C., Matelli, M., & Gentilucci, M. (1981b). Afferent properties of periarculate neurons in macaque monkeys. II. Visual responses. *Behavioral Brain Research*, *2*, 147–163.
- Roland, P., Larsen, B., Lassen, N., & Skinhoj, E. (1980). Supplementary motor area and other cortical areas in organization of voluntary movements in man. *Journal of Neurophysiology*, *43*, 118–136.
- Rosch, E. (1977). Human categorization. In N. Warren (Ed.), *Studies in cross-cultural psychology*. London: Academic Press.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Rosch, E. (1981). Prototype classification and logical classification: The two systems. In E. Scholnick (Ed.), *New trends in cognitive representation: Challenges to Piaget's theory* (pp. 73–86). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Rosch, E. (1994). Categorization. In V. S. Ramachandran (Ed.), *The encyclopedia of human behavior*. San Diego, CA: Academic Press.
- Rosch, E., & Lloyd, B. B. (1978). *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
- Sakata, H., & Taira, M. (1994). Parietal control of hand action. *Current Opinion in Neurobiology*, *4*, 847–856.
- Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning. *Behavioral and Brain Sciences*, *16*, 417–494.
- Talmy, L. (1983). How language structures space. In H. L. Pick & L. P. Acredolo (Eds.), *Spatial orientation: Theory, research, and application*. New York: Plenum Press.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, *12*, 49–100.
- Talmy, L. (1996). Fictive motion in language and “ception.” In P. Bloom, M. Peterson, L. Nadel, & M. Garrett (Eds.), *Language and space* (pp. 211–275). Cambridge, MA: MIT Press.
- Talmy, L. (1999). *Toward a cognitive linguistics*. Cambridge, MA: MIT Press.
- Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Scifo, P., Fazio, F., Rizzolatti, G., Cappa, S. F., & Perani, D. (in press). Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience*.

- Thompson, E. (1995). *Colour vision. A study in cognitive science and the philosophy of perception*. London: Routledge.
- Umiltà, M. A., Kohler, E., Gallese, V., Fogassi, L., Fadiga, L., Keysers, C., & Rizzolatti, G. (2001). "I know what you are doing": A neurophysiological study. *Neuron*, 32, 91–101.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Yamaguchi, T. (2004). The central pattern generator for forelimb locomotion in the cat. *Progress in Brain Research*, 143, 115–122.
- Yue, G., & Cole, K. (1992). Strength increases from the motor program: Comparison of training with maximal voluntary and imagined muscle contractions. *Journal of Neurophysiology*, 67, 1114–1123.
- Zeki, S. M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *Journal of Physiology*, 236, 549–573.
- Zimler, J., & Keenan, J. M. (1983). Imagery in the congenitally blind: How visual are visual images? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 269–282.

CN SI OBA 9

Queries

Gallese & Lakoff

The following text citations of articles, etc. do not appear in the reference list. Please provide full reference or delete the text citation.

- Q1 Fodor 1975, 1987, 1981
Pylysyn 1981
- Q2 Martin & Caramazza 2003
Simmons & Barsalou 2003
- Q3 Gallese 2003c
- Q4 Fodor 1998
- Q5 Sakata et al 2000—all author names need to be inserted in the text unless there are more than 6.
- Q6 Kurata & Tanji 1986
- Q7 Gallese 1999
- Q8 Keyser 2003
- Q9 Martin et al 1996—all author names need to be inserted in the text unless there are more than 6.
Perani et al 1995—all author names need to be inserted in the text unless there are more than 6.
- Q10 Decety 1991
- Q11 Martin et al 1996
Perani et al 1995
- Q12 Gleason & Kaschak 2002

Other queries:

- Q13 Rizzolatti et al 2000—do you mean 2000a or 2000b?