



PERGAMON

Neural Networks 13 (2000) 975–997

Neural
Networks

www.elsevier.com/locate/neunet

2000 Special Issue

Synthetic brain imaging: grasping, mirror neurons and imitation

M.A. Arbib^{a,b,*}, A. Billard^b, M. Iaconi^c, E. Oztop^{a,b}

^aUSC Brain Project, University of Southern California, Los Angeles, CA 90089-2520, USA

^bComputer Science Department, University of Southern California, Los Angeles, CA 90089-2520, USA

^cDivision of Brain Mapping, Neuropsychiatric Institute, UCLA School of Medicine, Los Angeles, CA 90095-7085, USA

Received 4 July 2000; accepted 5 July 2000

Abstract

The article contributes to the quest to relate global data on brain and behavior (e.g. from PET, Positron Emission Tomography, and fMRI, functional Magnetic Resonance Imaging) to the underpinning neural networks. Models tied to human brain imaging data often focus on a few “boxes” based on brain regions associated with exceptionally high blood flow, rather than analyzing the cooperative computation of multiple brain regions. For analysis directly at the level of such data, a schema-based model may be most appropriate. To further address neurophysiological data, the Synthetic PET imaging method uses computational models of biological neural circuitry based on animal data to predict and analyze the results of human PET studies. This technique makes use of the hypothesis that rCBF (regional cerebral blood flow) is correlated with the integrated synaptic activity in a localized brain region. We also describe the possible extension of the Synthetic PET method to fMRI. The second half of the paper then exemplifies this general research program with two case studies, one on visuo-motor processing for control of grasping (Section 3 in which the focus is on Synthetic PET) and the imitation of motor skills (Sections 4 and 5, with a focus on Synthetic fMRI). Our discussion of imitation pays particular attention to data on the mirror system in monkey (neural circuitry which allows the brain to recognize actions as well as execute them). Finally, Section 6 outlines the immense challenges in integrating models of different portions of the nervous system which address detailed neurophysiological data from studies of primates and other species; summarizes key issues for developing the methodology of Synthetic Brain Imaging; and shows how comparative neuroscience and evolutionary arguments will allow us to extend Synthetic Brain Imaging even to language and other cognitive functions for which few or no animal data are available. © 2000 Published by Elsevier Science Ltd.

Keywords: Affordances; Brain imaging; Computational model; fMRI; Grasping; Imitation; Mirror neurons; Neural networks; Parietal cortex; PET; Premotor cortex; Schemas; Synthetic PET

1. Global brain function

This paper is part of a Special Issue of *Neural Networks* on the subject “The Global Brain: Imaging and Neural Modeling”, with an emphasis on relating brain imaging data to “the activities of the underpinning neural networks, considered at the most appropriate level”. The Call for Papers states that “The hope of ... deriving a simple theoretical understanding of the action of the brain drives the research at both theoretical and experimental levels”. This raises many challenging issues:

1.1. Which “global brain”?

Which “global brain”? A human brain, a monkey brain

* Corresponding author. Address: USC Brain Project, University of Southern California, Los Angeles, CA, USA. Tel.: +1-213-740-9220; fax: +1-213-740-5687.

E-mail addresses: arbib@pollux.usc.edu (M.A. Arbib), billard@pollux.usc.edu (A. Billard), erhan@java.usc.edu (E. Oztop).

(and which species?) or some other brain? If a human brain, is it a mature brain, a developing brain, or a diseased brain? A male brain or a female brain? And what of individual differences? What of brain–body interactions? What of the effects of nature and nurture and social circumstance? We suspect that this very diversity makes “a simple theoretical understanding of the action of the brain” impossible. However, we believe that continuing study can develop a large store of models and methods which will let us integrate our insights into diverse brain mechanisms and behaviors, helping us to see order in complexity even if we cannot achieve simplicity. In particular, this paper will introduce our insight that data on the anatomy and physiology of the macaque monkey and the imaging of adult humans can contribute together to an integrated view of what might otherwise seem quite disparate behaviors: visually guided grasping, imitation, and language (Arbib, 2000; Iaconi et al., 1999; Rizzolatti & Arbib, 1998). Data on the mirror system (neural circuitry which allows

the brain to recognize actions as well as execute them) in monkey provide the integrating factor.

1.2. The discrepancy between brain imaging and neurophysiology

We note the discrepancy between the methodologies employed for gaining data on brain imaging and on the underpinning biological networks:

1. A brain imaging experiment typically operates at the task level: given a pair of tasks and some level of significance, find all areas of the brain as indicated by “blobs” which signal that part of the brain is significantly more active in task A rather than task B. It is common to interpret such a “blob” as indicating that a specific brain region X is crucial to task A but not to task B. A key challenge is then to go beyond the “blob” that indicates a statistical confidence measure for differences in rCBF to an understanding of actual “information processing” in neural networks. Brain imaging presents a problem of “too much” and “too little” — “too much” because its global view of brain activity may draw our attention to many brain regions whose causal role in the given task is unclear; “too little” because, in emphasizing regions “especially active” in the given task, it may downplay the role of regions whose activity, while not yielding a large rCBF signal, is nonetheless crucial to the neural implementation of the task.
2. Neurophysiological data are normally gathered in experiments on animals rather than humans. They typically focus on just one, or very few, brain regions considered relevant to execution of a given task. They thus lack the “global assay” of brain imaging, but offer the advantages of combining greater attention to parametric variation of the task with detailed insights into the variation of activity in neurons in such a way as to constrain hypotheses on actual circuitry. The disadvantage is that since each neurophysiological experiment is “local” as to task and circuitry, and since protocols for even apparently similar experiments differ greatly, there may be major interpretive challenges in integrating data from different experiments.

1.3. “Underpinning” neural networks

Much work in this journal has concerned *artificial* neural networks — a parallel, often adaptive, computing structure, usually inspired only loosely by neurobiology, which performs satisfactorily some given function. But the word “underpinning” implies some measure of biological constraint. The loosest constraint is to have a modular neural network in which each subnetwork is taken to represent the function of a particular brain region, but there is no constraint on the internal workings of the network. Some of the modeling described below is of this category, but we often add (as in the

FARS model below) the constraint that the components of the subnetwork match the anatomical characterization of cell types in the brain region being modeled and explain the neurophysiological data gathered from such neurons when the animal performs the task of interest. Let us reserve the term “*biological* neural network” for a neural network designed to model a specific brain region or set of brain regions, and whose unit neurons are constrained to match some neurophysiological data. More detailed models may attend to the dendritic structure of individual neurons, or the neurochemistry of, e.g. cellular activity.

While there are many ways to represent neural networks for computer simulation, the models described (with more or less detail) in this paper (Sections 3.1, 5.2, and 6.1) adopt the very simple *leaky integrator* model of the neuron, in which the internal state of the neuron is described by a single variable, the *membrane potential* $m(t)$ at the spike initiation zone. The time evolution of $m(t)$ is given by:

$$\tau \frac{dm(t)}{dt} = -m(t) + \sum_i w_i X_i(t) + h \quad (1)$$

with resting level h , time constant τ , $X_i(t)$ the firing rate at the i th input, and w_i the corresponding synaptic weight. We usually define the firing rate as a continuously varying measure of the cell’s activity. The *firing rate* is approximated by a sigmoid function of the membrane potential, $M(t) = \sigma(m(t))$, except in some modules of the imitation model (Section 5.2).

Another form of simplification is that, whereas brain imaging implicates a large number of brain regions in even apparently simple tasks, most specific models of biological neural networks usually include models of only a small selection of the brain regions for which neurophysiological data have been gathered for the task under consideration. In view of this, Section 6.1 quickly reviews several studies from the Arbib laboratory to indicate the challenges in integrating models of different portions of the nervous system to yield an integrated model that provides insight into all the brain regions implicated by human brain imaging in a given family of tasks.

1.4. Schema theory

A “connectionist” model employing an *artificial* neural network has no greater relevance to neurobiology than any other functional model. Indeed, in some cases such models may be positively misleading if they lead to spurious identifications of elements of, say, a backpropagation network with specific neurons of the visual system. It is thus important to have a language in which to express functional decompositions of the implementation of some behavior in such a way that neuralization of the components is open to analysis, rather than being built in. Arbib (1981) showed how to analyze perceptual structures and distributed motor control in terms of interacting schemas (functional units); schema theory has been extended to cognitive

systems as well (see, e.g. Arbib, Conklin, & Hill, 1987). Two important tenets of schema theory are:

1. Schemas may be defined at different granularities. For example, the schemas defining some overall task like visual control of hand movements are at a far higher level than the motor schemas for specific goal-related actions such as grasping, tearing, or holding. However, the high-level schemas contribute to each of the motor schemas for specific actions. It is thus dangerous to say that specific neurons “encode” an action when in fact it may take correlated activity in neurons in a variety of brain regions to “commit” the brain to the specific action. (We will need to recall this caveat when we discuss our preliminary attempts to model the mirror system in Section 6.1.)
2. Reinforcing this point, a schema defined functionally will in general require the interaction (“cooperative computation”) of multiple brain regions for its implementation; conversely, a given brain region will contribute to multiple schemas. Thus although we will speculate below on how various schemas may be assigned to specific brain regions in the macaque (and thus, presumably, to a homologous region in the human), we expect that further research will lead to more subtle insights into the anatomical distribution of each schema.

1.5. Boxology, schemas, and synthetic PET

The danger of models based on human brain imaging, then, is that they tend to focus on a few “boxes” based on brain regions associated with “blobs” of exceptionally high blood flow, rather than analyzing the cooperative computation of multiple brain regions. For analysis directly at the level of such data — augmenting the data of brain imaging with a variety of neurological insights — a schema-based model may be most appropriate, offering in the end a causally complete model of how some task may be implemented, achieving success to the extent to which activity of certain of the high-level schemas in the overall coordinated control program (Arbib, 1981) can be identified with activity observed through brain imaging. However, to relate brain imaging data to “the activities of the underpinning neural networks”, we offer a method — Synthetic PET imaging (Arbib, Biscoff, Fagg, & Grafton, 1995) and its generalizations — which uses neural models based on, e.g. primate, neurophysiology to predict and analyze results of brain imaging for a variety of human behaviors. This approach is very different from the approach of covariance structural modeling (CSEM) developed by many authors from McIntosh and Gonzalez-Lima (1994) to Taylor, Krause, Shah, Horwitz, and Mueller-Gaertner (2000). CSEM makes no use of animal data to constrain our insights into the human brain. Instead, it uses randomness and linearity assumptions to infer task-dependent patterns of

connectivity from the covariances observed in the brain imaging data. However, such “effective connectivity” is a task-dependent “projection” of actual synaptic connections. It is thus a task for future research to understand how to relax the CSEM assumptions and synthesize the “true connectivity” of brain regions from the effective connectivities revealed in a wide range of tasks.

1.6. The shape of the paper

Section 2 introduces the Synthetic PET method and discusses its extension to fMRI (functional magnetic resonance imaging). Section 3 then exemplifies Synthetic PET methodology by reviewing our earlier work on a detailed biological neural network model for visually directed grasping (the FARS model; Fagg & Arbib, 1998) and reports on the comparison of Synthetic PET predictions with our PET studies of human reaching and grasping. As a background for our second case study, Section 4 introduces data on the mirror system in monkey (neural circuitry which allows the brain to recognize actions as well as execute them), a system which we believe holds the key to the human capacity for both imitation (discussed at length in Section 5) and language (discussed briefly in Section 6). With this background, Section 5 summarizes our recent fMRI study of a simple form of imitation and relates it to a loosely biological model (i.e. made of modules in which each is a neural network identified with a brain region, but in which each module is implemented as a connectionist, rather than neurophysiologically constrained, neural network). Finally, Section 6 looks at the many facets of global modeling, with a sampler of detailed modeling of the kind needed to ground future global models (modeling the mirror system, the basal ganglia and sequential behavior, and the cerebellum and the coordination of reach and grasp), a brief look at our views on the relevance of mirror neurons and imitation to language, and a concluding perspective on homologies, brain imaging and neural networks.

2. Synthetic brain imaging defined

In order to provide a causal account of brain function constrained by data from both primate neurophysiology and human brain imaging, we need a method of “Synthetic Brain imaging”. We first recall the definition (Arbib et al., 1995) of *Synthetic PET imaging* — using neural models based on primate neurophysiology to predict and analyze results from PET (Positron Emission Tomography) brain imaging taken during a variety of human behaviors. The key hypothesis is that PET is correlated with regional cerebral blood flow (rCBF) and that this in turn correlates with the integrated synaptic activity in a region, and thus reflects in part neural activity in regions *afferent* to the region studied, rather than intrinsic neural activity of the region alone. However, the method can potentially accommodate other hypotheses on single cell correlates of imaged activity, and can thus be

applied to other imaging techniques, and so we also discuss the extension of the methodology to Synthetic fMRI. These two approaches will be exemplified in Section 3 [Grasp Control and (Synthetic) PET], and Section 5 [Imitation and (Synthetic) fMRI]. For other approaches to relating neural modeling and functional brain imaging, and cognition, see Horwitz, Tagamets, and McIntosh (1999), Tagamets and Horwitz (1998) and Taylor et al. (2000).

2.1. Defining synthetic PET

Mapping the activity simulated in neural network models of interacting brain regions — based on, say, single-cell recordings in behaving monkeys — into predictions of activity values to be recorded from corresponding regions of the human brain by imaging techniques such as PET involves two primary problems: localization, and modeling activation.

(i) *Localization*: Each array in the neural network model represents a neural population in a region identified anatomically and physiologically in the monkey brain. A Synthetic PET comparison requires explicit hypotheses stating that each such region A is homologous to a region h(A) in the human brain such that — within the tasks under consideration — A and h(A) perform their tasks in the same way. In some cases, such homologies are well defined. In other cases, the existence or identity of such a homology is an open question. Thus, the comparison of a Synthetic PET study with the results of a human brain scan study will, inter alia, be a test of the hypothesis “h(A) in human is homologous to A in (a given species of) monkey”, and comparison of synthetic and human studies may suggest a new homology to be tested in further studies.

(ii) *Modeling activation*: PET typically measures regional cerebral blood flow (rCBF). Arbib et al. (1995) hypothesize that the counts acquired in PET scans are correlated with local synaptic activity in a particular region (Brownell, Budinger, Lauterbur, & McGeeer, 1982; Fox & Raichle, 1985) and call this measure the “raw PET activity”. However, PET studies typically do not work directly with these values, but with the comparative values of this activity in a given region for two different tasks or behaviors. We thus define our computational method in two stages:

(a) Compute $rPET_A$, the *simulated value of raw PET activity*, for each region A of our network while it is used to simulate the monkey’s neural activity in some given task:

$$rPET_A = \int_{t_0}^{t_1} \sum_B w_{B \rightarrow A}(t) dt \quad (2)$$

where A is the region of interest, the sum is over all regions B that project to A, $w_{B \rightarrow A}(t)$ is the synaptic activity (firing rate \times |synaptic strength|) summed over all the synapses from region B to region A at time t , and the time interval from t_0 to t_1 corresponds to the duration of the scan.

(b) Compare the activities computed for two different tasks. The comparative activity $PET_A(1/2)$ for task 1 over

task 2 for each region A is then given by:

$$PET_A(1/2) = \frac{rPET_A(1) - rPET_A(2)}{rPET_A(2)} \quad (3)$$

where $rPET_A(i)$ is the value of $rPET_A$ in condition i , to compare the change in PET_A from task 2 to task 1. In the present study we use a different measure, defining the change in *relative synaptic activity* for region A from task 1 to task 2 with $\max(rPET_A(1), rPET_A(2))$ replacing $rPET_A(2)$ in the denominator of Eq. (3), which gives a more robust measure of relative activity. In either case, the result is a Synthetic PET comparison, which presents our prediction of human brain activity, as based on neural network modeling constrained by monkey neurophysiology and known functional neuroanatomy.

Note that we are comparing *synaptic* activity for each region A, not the neural activity of A (though, as we shall stress at the end of Section 2.2, it is an open research issue to determine how often the difference between these measures is significant). Indeed, Taylor et al. (2000) state that although few experiments have been performed to test the suggestion (Arbib et al., 1995; Jueptner & Weiller, 1995; Tagamets & Horwitz, 1998) that inhibition be treated in the sum with the same (positive) weight as excitation, “the few good ones have supported this view (see Horwitz & Sporns, 1994; Jueptner & Weiller, 1995, for reviews).” As a computational plus (going beyond the imaging technology), we may also collect the contributions of the excitatory and inhibitory synapses separately, based on evaluating the integral in (1) over one set of synapses or the other. Using Synthetic PET, we can break apart different factors that contribute to the measure of synaptic activity so that they can be studied independently. This can allow a much more informed view of the actual PET data that are collected, possibly shedding light on apparent contradictions that arise from interpreting rCBF simply as cell activity (an example has been demonstrated in Arbib et al., 1995).

2.2. Extending synthetic PET to synthetic fMRI

For the present paper, we shall assume that Synthetic PET and Synthetic fMRI (functional MRI) are defined by precisely the same formula except for a possible difference in temporal and spatial resolution. In this section, we briefly discuss some necessary steps for future research directed towards more realistic approaches to various modalities of Synthetic Brain Imaging (SBI).

Our current measure of “raw PET activity”, based on a linear function of the total of the absolute value of synaptic activity, already (as we shall demonstrate in Section 3) yields *qualitatively* useful results in evaluating the sign and small versus large magnitude of activities seen in PET comparisons. However, we do not claim that this first approximation yields *quantitatively* accurate predictions. We note, as a target for further research on Synthetic Brain Imaging, the interest of evaluating a variety of more

quantitative fits based on (possibly nonlinear) combinations of cell firing rates, synaptic change, and synaptic activity per se. We also need to add a stochastic analysis to account for the variation in PET activity seen in the same subject on different trials.

It also worth noting that future developments in SBI will probably need to be based on neural network models extended to include the “*vampire model*” of the neuron, i.e. looking at interactions between neurons, glia, and blood vessels to model more explicitly the way in which neurons “suck blood”, embedding the study of their information processing in a richer model of the metabolic underpinnings of these computations. For example, when the activity of a brain region increases, the, e.g. PET signal for that region increases not only in intensity but also in volume. However, this need not imply that more neurons are being recruited, only that increased blood flow to the same set of neurons means that an increased rCBF signal will be detected in nearby blood vessels connecting to and from the blood vessels of the stated region. For example, if the probability of a square millimeter of tissue sharing blood supply with another square millimeter drops off linearly with the distance between them, then the diameter of a circle of above threshold rCBF will increase linearly as the rCBF within a target region increases. Of course, more detailed studies of activation would have to be based on sophisticated 3D maps of blood supply to complement the maps of patterns of neural innervations used in the Synthetic PET method. Unfortunately, these maps are not yet available. Moreover, it must be acknowledged that the relation between neuronal firing, synaptic activity, and metabolic activity is still not entirely understood and that different models have been proposed, based on both theoretical considerations and empirical findings (Raichle, 1998). As far as synthetic PET and synthetic fMRI are concerned, we argue below that we can apply similar formulas to both approaches on the basis of the following considerations:

The Blood Oxygenation Level-Dependent (BOLD) fMRI signal is currently interpreted as mainly driven by a regional change in the ratio between deoxyhemoglobin and oxyhemoglobin, due to the increased influx of oxygenated blood in a given brain region with increased metabolic demand. The blood influx is estimated to be larger than is required by the increased metabolic demand. Thus, a relative decrease in deoxyhemoglobin (i.e. after oxygen has been removed to serve metabolism), compared to oxyhemoglobin, would result in signal increase due to the different magnetic properties of the two molecules. A change in the ratio between deoxyhemoglobin and oxyhemoglobin, however, could, in principle, result from increased blood volume, determining an increase in the water fraction around deoxyhemoglobin. Thus, a change in blood volume without a change in blood flow may determine the same signal increase determined by an increase in blood flow. With regard to our paper, and to the plausibility of adopting the synthetic PET approach to

model BOLD fMRI signal, a fair analysis of a wealth of empirical evidence (the detailed review of which is beyond the scope of the present article) suggests the following conclusions:

1. The BOLD signal is probably composed of two main factors, blood flow changes and blood volume changes (plus some other minor factors probably contributing very little to the signal, thus negligible here) (Hess, Stiller, Kaulisch, & Scheich, 2000).
2. In spite of this, the BOLD signal co-localizes fairly well with the PET signal and with TMS mapping (i.e. using transcranial magnetic stimulation) of cortical representation of functions (Paus, 1999; Ro, Cheifet, Ingle, Shoup, & Rafal, 1999), and, at least in the monkey, with changes in the local field potential (N. Logothetis, personal communication) that are considered the expression of neural activity at population level. Also, quantitative estimates of single-unit activity in macaques and BOLD fMRI in humans in response to motion coherence in the MT/V5 complex suggest that BOLD fMRI is proportional to firing rate (Rees, Friston, & Koch, 2000).
3. Taken together, these data suggest that blood volume effects do not dramatically alter the relation between the PET signal and BOLD signal on one side, and BOLD signal and regional neural activity on the other.
4. Moreover, quantitative differences between integrated synaptic activity and neuronal firing are presumably small and negligible. This may be due to the following considerations:
 - (4a) most of the regional metabolic demand due to neuronal activity is eminently synaptic (Yarowsky, Kadokaro, & Sokoloff, 1983);
 - (4b) the relationship between firing rate and metabolic synaptic activity is approximately linear (Kadokaro et al., 1987);
 - (4c) the overwhelming majority of synapses are local (Douglas, Koch, Mahowald, Martin, & Suarez, 1995);
 - (4d) a high level of cortical inhibition, resulting in low firing rate and high metabolic activity, is unlikely to be long lasting (van Vreeswijk & Sompolinsky, 1998).

However, possible discrepancies between measures based on synaptic activity and neuronal firing — and measures which also explore the metabolic costs of synaptic change — will be a topic for future consideration. In the present paper, we will first report on a Synthetic PET study of grasping behavior in which we use the synaptic activity measure, whereas in the Synthetic fMRI experiments for imitation described below we calculated the integral over time of the neural activity and not the synaptic one, since in the DRAMA network used in that model the synaptic weight have a different meaning than in the classical Perceptron model and would not relate directly to the neural activity.




Fig. 1. According to the FARS model, AIP uses visual input to extract affordances, which highlight the features of the object that are relevant to grasping it. F5 then applies various constraints to select a grasp for execution and to inform AIP of the status of its execution, thus updating AIP's active memory. The areas shown are AIP (anterior intraparietal cortex), area F5 (of the ventral premotor cortex), and regions providing supporting input to F5, namely F6 (pre-SMA), area 46 (dorsolateral prefrontal cortex), and F2 (dorsal premotor cortex).

3. Grasp control and (synthetic) PET

The first Synthetic PET study (Arbib et al., 1995) focused on Synthetic PET predictions based on a model of Dominey and Arbib (1992); of mechanisms in the monkey brain for control of saccadic eye movements. Here we summarize a follow-up study (Arbib, Fagg, & Grafton, 2000), based on the FARS (Fagg–Arbib–Rizzolatti–Sakata) model of the grasping process of parietal–premotor interactions in the visual control of monkey grasping movements. We first summarize key properties of the FARS model, then show how we used it to generate Synthetic PET predictions which were then compared to a new PET study of human grasping (Grafton, Fagg, & Arbib, 1998).

3.1. The FARS model

The cells of area F5 of the macaque inferior premotor cortex are often selective for the type of grasp made by the monkey (Rizzolatti et al., 1988). Grasps observed during these experiments include precision pinches (using the tips of the index finger and thumb), lateral pinches (thumb against the side of the index finger), and power grasps (four fingers opposing the palm). In addition, the firing of these cells typically correlated with a particular phase of the ongoing movement.

F5 exchanges cortico-cortical connections with the anterior intra-parietal area of parietal cortex (AIP), whose cells demonstrate a variety of both visual- and grasp-related responses (Taira, Mine, Georgopoulos, Murata, & Sakata, 1990).

The FARS model of the grasping process (Fagg & Arbib, 1998) focuses on the roles of several areas, especially the interactions of F5 and AIP, but also their relationship with

pre-SMA (F6), F2 (dorsal premotor cortex), and the basal ganglia (BG).

3.1.1. A brief model overview

The crucial aspects of the model (see Fig. 1) are the following:

1. AIP serves the dual role of first computing a *set of affordances* for the object being attended (i.e. AIP highlights properties of the object relevant for physically interacting with it), and then maintaining an *active memory* of the selected affordance as the corresponding grasp is prepared and executed.
2. F5 integrates a variety of constraints to decide on the single grasp that is to be executed. These constraints include visual information (from the affordances extracted by AIP), task information (from pre-SMA, F6), instruction stimuli (from dorsal premotor cortex, F2), and a working memory of recently executed grasps. When the movement is triggered, F5 is responsible for the high-level execution and subsequent monitoring of the planned preshape and grasp.
3. As indicated in Fig. 2, F6 (pre-SMA) represents the high-level execution of the sequence, while phase transitions dictated by the sequence are managed by the basal ganglia (BG). For a task (“the Sakata protocol”) in which the monkey was presented with an object, then grasped the object in response to a go signal, held the object, and finally released the object after a secondary go signal, the following phases in F5 activity were influenced by BG control mechanisms: preparatory (set), finger extension, finger flexion, holding, and release.
4. Fagg and Arbib (1998) offered both a computational analysis and an analysis of empirical data in support of the hypothesis that not only is F5 responsible for unfolding (with BG supervision) the grasp in time during the execution of the movement, but that F5 also sends recurrent connections back to AIP to update AIP's *active memory* for the grasp that is about to be executed or that is being executed by F5.

3.1.2. Control of conditional behavior

The many details of the FARS model (Fagg & Arbib, 1998) are beyond the scope of the present review. Here we focus on its performance in a *conditional* task in which the grasp is not known prior to a given trial, and is only determined by an arbitrary *instruction stimulus* (IS) made available during the course of the trial (e.g. an LED whose color indicates one of two grasps). The dorsal premotor cortex (F2) is thought to be responsible for the association of arbitrary IS with the preparation of motor programs (Mitz, Godshalk, & Wise, 1991; see Fagg & Arbib, 1992 for a related model). In the FARS model, then, F5 combines the view of the object (with its multiple affordances signaled by

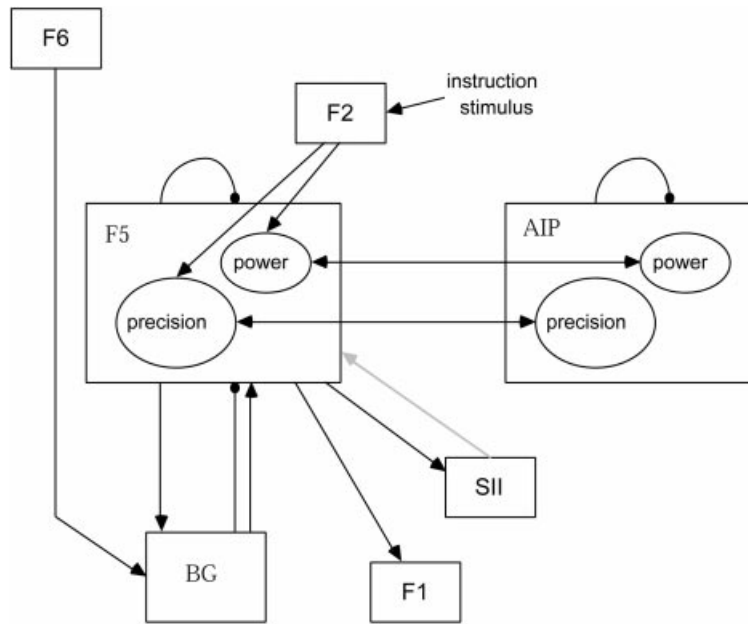


Fig. 2. Another view of FARS model architecture. Arrows indicate excitatory connections between regions; filled circles indicate inhibitory connections. The precision pinch and power grasp pools in F5 and AIP are connected through recurrent excitatory connections. The precision pinch pool contains more neurons than other grasps. The dorsal premotor cortex (F2) biases the selection of grasp to execute as a function of the presented instruction stimulus.

AIP) with the instruction stimulus (F2) to specify which grasp is to be executed.

Fig. 2 presents a schematic view of the model's architecture. Arrows indicate excitatory connections between regions; filled circles indicate inhibitory connections. The dorsal premotor cortex (F2) biases the selection of grasp to execute as a function of the presented instruction stimulus. The precision pinch and power grasp pools in F5 and AIP are connected through recurrent excitatory connections. Since Rizzolatti noted that more F5 neurons are related to the precision pinch than to other grasps, the precision pinch pool in the FARS model contains more neurons than other grasp pools. Arbib et al. (2000) show how this affects the Synthetic PET measure in these and downstream regions, and discuss the implication of these results for reconciling the FARS model with actual PET studies of rCBF in comparing the precision pinch task against the power grasp task. However, in what follows we concentrate on the comparison of a conditional with a non-conditional task.

3.2. Synthetic PET results

Arbib et al. (2000) conducted two different Synthetic PET experiments, providing predictions for what we expect when the experiments are performed in the human. In both experiments, the modeled subject is asked to grasp a single object using one of two grasps. In the first experiment, we examined (i) the effects of knowing which grasp to use prior to the onset of recording (non-conditional task), and (ii) only being told which grasp to use after a delay period (conditional task). In the latter task, an instruction stimulus in the form of a bi-colored LED informs the subject

which grasp should be used. Here we summarize the comparison of conditional and non-conditional tasks.

The most significant predicted change was the level of activity exhibited by area F2 (dorsal premotor cortex). Its high level of activity in the conditional task is due to the fact that this region is only involved when the model must map an arbitrary stimulus to a motor program. In the non-conditional task, the region does not receive IS inputs, and thus its synaptic activity is dominated by the general background activity in the region.

The additional IS inputs in the conditional task have a second-order effect on the network, yielding small changes in synaptic activity in F5, BG, and AIP (see Arbib et al., 2000 for details). Increased synaptic activity in F5 is due to the additional positive inputs from F2. These inputs also cause an increase in the region's *activity level*, which is passed on through excitatory connections to both AIP and BG (recall Fig. 2).

It is important to recall that synaptic activity does not have the same meaning as *neural activity*. This can be seen by examining the definition of $w_{B \rightarrow A}(t)$ (see Eq. (2)). The absolute value of the synaptic strength contributes positively to this measure — so increases in either positive or negative signals into a region will be reflected as an *increase* in synaptic activity. Neural activity, on the other hand, depends monotonically on the synaptic strength: increases in positive signals result in increases in neural activity, whereas increases in negative signals result in decreases of activation. An important ability of the Synthetic PET technique is that the positive and negative contributions to the Synthetic PET measure can be differentiated in the simulation. This information, combined with knowledge of the




Fig. 3. Left hemisphere localization of task related effects. PET statistical comparisons of the pooled data across subjects (in red, $P < 0.005$) are superimposed on a single subject's MRI scan centered in the same coordinate space. The view is left superior oblique, and denotes differences between conditional grasp selection versus fixed grasping.

gross anatomy (especially the sign of connections between regions), can aid in inferring changes in neural activity across tasks. Although the positive contributions to F5 and AIP essentially dominate the full PET measure, we also saw small increases in the negative inputs into these regions. These inhibitory signals are due to negative inputs from local recurrent connections in the respective areas (in the case of F5, BG also contributes additional negative inputs). This serves as additional evidence that both F5 and AIP experience increases in their overall neural activity.

3.3. Human brain imaging

To provide human PET data relevant to the above Synthetic PET predictions, Grafton et al. (1998) asked subjects to repeatedly perform grasping movements over a 90-s scanning period. The targets for grasping were located at three stations mounted on an experimental apparatus, with each station consisting of both a rectangular block that could be grasped using a power grasp, and a pair of plates (mounted in a groove on the side of the block), which could be grasped using a precision pinch (thumb and index finger). A force sensitive resistive (FSR) material, mounted on the front and back of the block, detected when a solid power grasp had been established. The pinch plates were attached to a pair of mechanical micro-switches which detected when a successful precision pinch had been executed. For each station, the block and plates were mounted such that the subject could grasp either one without requiring a change in wrist orientation. A bi-colored LED at each station was used to instruct the subject as to the next target of movement. A successful grasp of this next target was indicated to the subject by a change in the color of the

LED. The subject then held the grasp position until the next target was given. Targets were presented every 3 ± 0.1 s.

Four different scanning conditions were repeated three times each. In the first, subjects repeatedly performed a power grasp to the indicated block. The target block was identified by the turning on of the associated LED (green in color). When the subject grasped the block, the color of the LED changed from green to red. For the second condition, a precision pinch was used. The target was identified in the same manner as the first condition. In the third grasping condition (conditional task), the initial color of the LED instructed the subject to use either a precision pinch (green) or a power grasp (red). When contact was established, the LED changed to the opposite color. In the fourth (control) condition, the subjects were instructed to simply fixate on the currently lit LED, and not make movements of the arm or hand (prior to the scan, the arm was placed in a relaxed position). The lit LED changed from one position to another at the same rate and variability as in the grasping tasks. Prior to scanning, subjects were allowed to practice the tasks for several minutes. Grafton et al. (1998) provide results for Grasp versus Rest, Precision versus Power Grasp, Conditional versus Non-Conditional Grasp. Here we present only the third case:

Fig. 3 reflects differences of conditional grasp selection (power or precision based on color cues) as compared to an average of the fixed grasping conditions (power and precision tasks): $\text{Cond} - (\text{Power} + \text{Precision})/2$. The upper arrow indicates a large area of significance in the left superior frontal sulcus corresponding to the dorsal premotor cortex. As noted earlier, this region in monkey is thought to be involved in the arbitrary association of stimuli with the preparation of motor programs. The lower arrow indicates increased CBF in the left inferior parietal lobule and intra-parietal sulcus. Because this comparison is counterbalanced for the amount of movement made during execution of the tasks, there is no difference observed in the motor execution areas.

3.4. Comparison of PET and synthetic PET

The model predicts that the conditional task should yield much higher activation in F2 (dorsal premotor cortex), some activation of F5, and a slight activation of AIP. The human experiment confirmed the F2 result, but failed to confirm the predictions for F5. Furthermore, in human we see an activation of the inferior parietal cortex, along the intra-parietal sulcus, which is perhaps an AIP homologue.

Can we make use of the negative F5 result to further refine the model? Consider the functional connectivity of these regions in the model (Fig. 4, left). In the model, the strength of the projection from F2 to F5 is essentially a free parameter. In other words, there is a wide range of values over which the model will correctly perform the conditional and non-conditional tasks. The implication is that, by tuning this parameter, we can control this projection's contribution

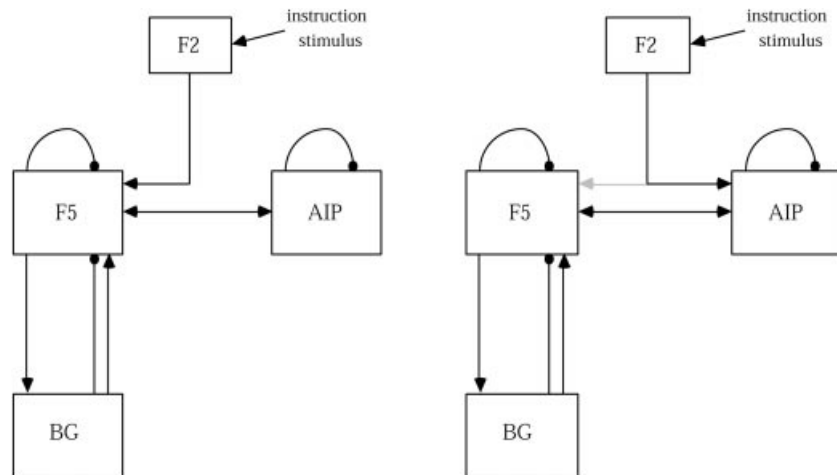


Fig. 4. Previous functional model (left; compare Fig. 2) and updated functional model (right). In the revised model, the information from F2 flows (primarily) into the circuit through a projection into AIP.

to the synaptic activity measure in F5. However, the difference in AIP synaptic activity from the non-conditional to the conditional task will always be less than the difference observed in F5. Why is this the case? By increasing the projection strength from F2 to F5, we observe an increase in both F5 synaptic *and* cell activity. The increase in F5 cell activity, however, is attenuated by local, recurrent inhibitory connections. Thus the excitation that is then passed on to AIP via F5 does not reflect the full magnitude of the signal received from F2.

The conclusion is that, although we can adjust the free parameter to match one or the other observations in the human experiment (of either F5 or AIP changes), the model cannot reflect both at the same time. One possibility for repairing this problem in the model is to reroute the F2 information so that it enters the grasp decision circuitry through AIP (or both AIP and F5), rather than exclusively through F5 (Fig. 4, left). This would yield an increase in activity in AIP due to F2 activation with only an attenuated signal being passed on to F5, resulting in only a small increase in F5 synaptic activity. Note that we do not necessarily assume that there exists a direct cortico-cortical connection from F2 to AIP or F5, but only that there is a functional connection (which potentially involves multiple synapses).

The low-level details of the FARS grasping model (Fagg & Arbib, 1998) were derived primarily from neurophysiological results obtained in monkey. The synthetic PET approach extracts measures of regional synaptic activity as the model performs a variety of tasks. These measures are then compared to rCBF (regional cerebral blood flow) observed during human PET experiments as the subjects perform tasks similar to those simulated in the model. In some cases, the human results provide confirmation of the model behavior. In other cases, where there is a mismatch between model prediction and human results, it is possible

(as we have shown) to use these negative results to further refine and constrain the model and, on this basis, design new experiments for both primate neurophysiology and human brain imaging (see Arbib et al., 2000 for further details).

4. The mirror system and imitation

As a background for our second case study on Synthetic Brain imaging, on Imitation and (Synthetic) fMRI in Section 5, we now review data on the mirror neuron system for grasping in the monkey brain, and then discuss its implications for the study of imitation in humans. We shall briefly discuss the relevance of mirror systems and imitation for the evolution of language in Section 6.3.

4.1. The monkey mirror system for grasping

In Section 3.1, we saw that cells of area F5 of inferior premotor cortex of the macaque are often selective for the type of grasp made by the monkey (Rizzolatti et al., 1988), while the anterior intra-parietal area of parietal cortex (AIP) has cells demonstrating a variety of both visual- and grasp-related responses (Taira et al., 1990). Here we extend the F5 database by noting the discovery by Gallese, Fadiga, Fogassi, and Rizzolatti (1996) and Rizzolatti, Fadiga, Gallese, and Fogassi (1996) of a subset of F5 grasp cells which they called *mirror neurons*. Like other F5 neurons, mirror neurons are active when the monkey performs a particular class of actions. However, in addition the mirror neurons become active when the monkey observes the experimenter or another monkey performing an action. In most of the mirror neurons, there is a clear relation between the coded observed and executed action. The actions studied so far include grasping, manipulating and placing. The congruence between the observed and executed action varies. For some of the mirror neurons, the congruence is

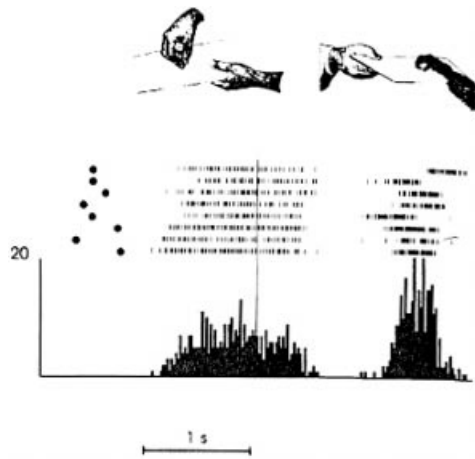


Fig. 5. Activity of a cell during action observation (left) and action execution (right). There is no activity during either initial presentation of the object or as the tray is brought towards the monkey (from Gallese et al., 1996).

quite loose; for others, not only must the general action (e.g. grasping) match but also the way the action is executed (e.g. power grasp) must match as well. To be triggered the mirror neurons require an interaction between the experimenter and the object. The sight of the experimenter or the object alone does not trigger mirror activity.

Fig. 5 demonstrates the dual response property of mirror neurons. The recorded neuron in the figure is silent during the presentation of the object, but starts firing when the experimenter picks up the object. The neuron interestingly does not fire during the time the tray is moved towards the monkey (the monkey knows that he is going to pick up the object) and finally it starts firing again when the monkey, himself, picks up the object. A range of congruence (the degree of strictness of the match between the action observed and the executed action) is found in mirror neurons. The most congruent responses require the observed action and the executed action to use the same grasp type.

Fogassi, Gallese, Fadiga, and Rizzolatti (1998) find that F5 is not the only area that has mirror neurons. The rostral part of the inferior parietal lobule of the macaque monkey (area 7b or PF) also has neurons with similar mirror properties.

Recent human PET studies indicate that a similar mirror neuron system exists in the human brain. Grafton, Arbib, Fadiga, and Rizzolatti (1996a) scanned subjects under three conditions: object viewing (the control condition), observing a precision grasping of common objects, and imagining themselves doing the same grasp. The imagined minus control and observation minus control results were compared. Differences were seen in both lateral activations and medial/dorsal activations. In the observation minus control condition, activity locations were left STS (the cortex of superior temporal sulcus), left inferior frontal area 45, and the left parietal area 40. Also there was some activation found in the rostral part of the left intraparietal sulcus. Grafton et al. pointed out that the location of the

activated sites during grasping observation in humans corresponds rather well to that of monkey cortical areas containing neurons that selectively discharge during hand action observation. Finally they conclude that in monkeys and humans a similar cortical circuit is involved in representing observed grasping. Unfortunately we cannot be sure that the activities are due to mirror neuron system since there is no convincing study showing that the same regions get activated *selectively* for self executed grasps (Grafton, Fagg, Woods, & Arbib, 1996).

4.2. Imitation in humans

Imitation is a complex behavior requiring the integration of a multitude of sensory stimuli and of a variety of movements. This integration process is affected by attentional, emotional, social, and environmental factors. All these factors modulate learning by imitation. Still, some fundamental principles of learning by imitation can be extrapolated by the investigation of its neural mechanisms. These principles can be used to define the mental representations underlying learning by imitation in neural, computational, psychological, and social terms.

In spite of the complexity of imitative behavior, there are data (still somewhat controversial) supporting the claim that even soon after birth human infants can imitate facial and some manual gestures (Meltzoff & Moore, 1977). However, it is important to stress that there may be different mirror mechanisms for different acts, and that imitation for grasping requires months to develop. Smiling is probably an innate releasing mechanism for smiling (whether the infant smiles at the parent or the parent smiles at the infant) depending, like laughing and weeping, on “built-in” circuitry. We distinguish this from building a novel repertoire and then building further to assemblages of variations.

Our hypothesis is that an appropriate mirror system (observation/execution matching system) orchestrates the various components involved in the sensorimotor transformations required by imitation of behaviors more complex than smiling, laughing and weeping. More complex behaviors are not simply learned and replicated as unified non-decomposed patterns. In terms of neural control, the brain generates internal models (neural mechanisms that can mimic the input–output characteristics of the learning system: Jordan & Rumelhart, 1992; Wolpert, Ghahramani, & Jordan, 1995). To learn, a system must be able to generalize beyond the set of sensory stimuli and movements that has been experienced in the past. This is possible only through the internal representation of sensory consequences of motor commands (the forward model) or, inversely, through the extrapolation of motor commands from the observation of motor behavior (the inverse model).

The representation of action supported by mirror neurons seems to be more a representation of action goals than of motor outputs (Arbib & Rizzolatti, 1997). We must thus seek to establish how these goals may be combined to form

a repertoire of purposeful and adaptive behavior. Moreover, action goals are hierarchically organized in human imitative behavior. Studies on imitation in children show that preschoolers represent the most salient goal of the action to be imitated, but often ignore lower goals (Bekkering, Wohlschlaeger, & Gattis, 2000) (incidentally, this is the same behavior of patients with aphasia, Head, 1920; supporting the possible links between imitation and language to be taken up briefly in Section 6.3). Goals are easily embodied by objects (Mary grasps a cup). Thus, the role of objects in action observation and learning by imitation seems to be a crucial one and needs to be clarified.

In Section 5.1, we present the results of an fMRI study of an imitation task that involves a subject observing a very simple finger movement and responding with a movement which in its entirety is already in the subject's own repertoire. Skill acquisition more generally (consider, for example, learning a new dance), however, rests on a far more complex form of imitation that involves "parsing" a complex movement into more or less familiar pieces, and then performing the corresponding composite of (variations on) familiar actions. Note the insistence on "more or less familiar pieces" and "variations". Elsewhere, we (Arbib, 1981) have introduced the notion of a coordinated control program, to show how a new behavior could be composed from an available repertoire of perceptual and motor schemas (the execution of a successful action will in general require perceptual constraints on the relevant movements). However, skill acquisition not only involves the formation of new schemas as composites of old ones, it also involves the tuning of these schemas to match a new set of conditions, to the point that the unity of the new schema may over-ride the original identity of the components. For example, if one is acquiring a tennis stroke and a badminton stroke through imitation, the initial coordinated control program may be identical, yet in the end the very different dynamics of the tennis ball and shuttlecock lead to divergent schemas. Conversely, a skill may require attention to details not handled by the constituent schemas of the preliminary coordinated control program. Fractionation may be required, as when the infant progresses from "swiping grasps" at objects to the differentiation of separate schemas for the control of arm and hand movements. Later, the hand movement repertoire becomes expanded as one acquires such novel skills as typing or piano playing, with this extension matched by increased subtlety of eye–arm–hand coordination. Thus we have three mechanisms (at least) to learn completely new actions: forming new constructs (coordinated control programs) based on familiar actions; tuning of these constructs to yield new encapsulated actions, and fractionation of existing actions to yield more adaptive actions as tuned, coordinated control programs of novel schemas.

However, the study of imitation mechanisms based on the building up and/or the fractionation of coordinated control programs is a task for the future. Instead, Section 5.1 will

present fMRI data on human imitation of simple finger movements, while Section 5.2 will offer a preliminary global model of imitation based on the joint-by-joint reproduction of an observed movement which Section 5.3 will build upon to provide a Synthetic fMRI study of grasping and finger tapping.

5. Imitation and (synthetic) fMRI

5.1. An fMRI study of "simple" imitation

In a recent study of "simple" imitation in humans (Iacoboni et al., 1999), we compared fMRI activity during imitation with fMRI activity during two control tasks in which subjects were required to perform the same finger movements as in the imitation task, but in response to symbolic or spatial cues. We reasoned that if some human brain areas had mirror properties during imitation, then these areas should become active during action execution regardless of how the action was elicited, and should also become more active during imitation. This is because the action to be performed was also observed and so a brain area with mirror properties would be simultaneously activated by the motor command to be executed and by the visual input representing the observed action. Moreover, these areas should become active during action observation compared to rest.

Fig. 6 shows a comprehensive view of activated areas during imitation versus rest in this experiment. Among these activated areas, we found two areas with mirror properties, i.e. they were simultaneously activated by the motor command to be executed and by the visual input representing the observed action. One was located in the inferior frontal cortex of the left hemisphere and corresponded to Brodmann area 44 (Broca's area). The other was located in the rostral part of the posterior parietal cortex in the right hemisphere (area PE/PC). We proposed that the left inferior frontal cortex (area 44) describes the observed action in terms of its motor goal, whereas the right posterior parietal area (PE/PC) codes the precise kinesthetic aspect of the movement (Iacoboni et al., 1999). Note, however, that the lateralization of the activations is rather relative. In fact, when statistical thresholds are lowered, bilateral activations are observed in both inferior frontal and posterior parietal cortices.

The demonstration of mirror properties in Broca's area during imitation is in support of the hypothesis that language evolved from a basic mechanism not originally related to communication: the mirror system with its ability to generate and recognize actions (Rizzolatti & Arbib, 1998; see also Section 6.3).

We also observed mirror properties in the activity of an area located in STS. The activity in this region did not reach full significance after correction for multiple comparisons and was not reported in the original paper. However, there




Fig. 6. Cortical and subcortical activations during imitation versus rest. Within these areas, some (inferior frontal cortex, anterior part of the posterior parietal cortex, and STS) had activity that is characteristic of areas with mirror properties.

are two reasons to believe that this area is likely to belong to the human mirror system. First of all, we have subsequently observed similar mirror properties in a follow-up experiment in the same region. Fig. 7 shows the activity of this region in both experiments, for identical tasks. Second, in two PET experiments (Grafton, Arbib et al., 1996;

Rizzolatti et al., 1996) subjects observed the experimenter grasping a three-dimensional object. Grasp observation significantly activated the STS. Neurons that become selectively active in the STS during the sight of moving hands and faces but not of moving objects were also described in single unit studies.

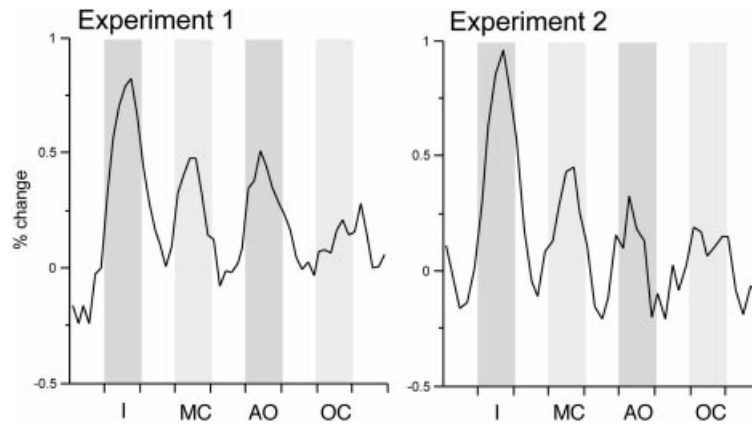


Fig. 7. Activity in STS showing mirror properties. This activity was observed in two separate experiments comprising 12 subjects in the first experiment and 10 different subjects in the second experiment. I: imitation; MC: motor control; AO: action observation; OC: observation control.

The mirror area in STS is more posterior than the STS region found active in previous studies of action observation in humans and in macaques. (See Section 6.1 for further discussion of data on STS.) It is possible that the STS region, a region critical for the understanding of social signals, is constituted of a rostral sector with sensory properties only and a caudal sector with mirror properties. It has been proposed that the anterior STS neurons that are activated by the sight of moving hands and faces analyze the actions of others. This anterior ‘sensory-only’ STS region may feed into the more posterior ‘mirror’ STS region during imitation. The anterior STS region would provide a visual analysis of the actions of others, whereas the more posterior STS regions would provide a first level of matching between observation and execution. The other two mirror areas described in our paper would describe the observed action in terms of its motor goal (Brodmann Area 44), and would code the precise kinesthetic aspect of the movement (PE/PC). It is possible that each of these other mirror areas has a dedicated cortical area for preprocessing of the action to be imitated during action observation. In fact, in our fMRI experiment on imitation we did observe greater activation for action observation than for observation of the symbolic or spatial cues in two areas, one in the dorsolateral prefrontal cortex (that would feed into Brodmann Area 44) and one in the caudalmost part of the intra-parietal sulcus (that would feed into PE/PC).

We saw (Section 4.1) that mirror neurons are found in the macaque brain for object-oriented grasping actions and show a narrow tuning with respect to the type of object and the way it is treated but quite a broad tuning for the trajectory of the movement or the motor effector involved (Gallese et al., 1996). In our fMRI study we used an intransitive action, and no objects were involved. The presence of real objects, however, alters the motor component of the task and makes it more difficult to interpret imaging data. However, dots or marks on a table can be used as end-points of imitated actions. These dots or marks are known to alter the imitative behavior of preschoolers and adults. In a recent

study, the experimenter initiated hand movements to one of two dots on a table and asked children to imitate his actions. In another condition the same movements were performed and were imitated, except that this time the movements were directed to the same places on the table but no dots were visible. In the latter condition children imitated well the movements of the experimenter, including their laterality, whereas when dots were visible children would often reach for the correct dot with the incorrect hand (Bekkering et al., 2000). This is because when observing object-oriented actions, the object activates a ‘direct’ motor program; i.e. a motor program that leads most directly to the effect the action has on the object, like grasping, pointing, covering, etc. Without objects, the relative positions of the fingers become the goal of the action. These goals also call the ‘direct’ motor program, but now the ‘direct’ motor program is matching the observed non-object oriented action. This predicts that goal-oriented areas, such as Broca’s area, should be *more active during the imitation of object-oriented actions* than during the imitation of non-object-oriented movements. The reasoning is that if Broca’s area is the homologue of the monkeys mirror neuron system (tuned for object-oriented actions), then Broca’s area should also be better tuned for object-oriented actions. For the parietal area PE/PC we might expect that the visible presence of goals would not modulate its activity and thus goal-oriented actions and intransitive actions to be imitated will activate this region similarly. Preliminary data from ongoing fMRI experiments performed in the Iacoboni laboratory seem consistent with these predictions.

Taken together, these data suggest a first draft of an information processing model of human imitation according to which the early ‘description’ of the action is performed in STS. From this region, inputs would be transferred to the posterior parietal cortex for precise coding of the kinematics of the action. This region would in turn send information to Broca’s area for defining the goal of the action. Information would be recursively shuffled among these areas for a tight coupling of these three ‘levels’ of




Fig. 8. The global imitation model consists of three parts for visual recognition, motor control and learning and is composed of biologically inspired modules, namely the superior temporal sulcus (STS), the spinal cord, the primary motor cortex (M1), the dorsal premotor area (PMd) and the cerebellum.

information processing, i.e. motion recognition, kinesthetics, and goal-orientation. In Section 6.1 we outline the structure of a model which develops some of these ideas on the basis of neurophysiological data. However, the present article emphasizes a “global” model (Section 5.2) based on a high-level view of a few large brain regions, each modeled connectionistically without attempting to address details of neurophysiology. We then offer (Section 5.3) a Synthetic fMRI analysis of the model. The integration of these two complementary models, and the development therefrom for a model of complex imitation based on coordinated control programs, is a goal for future research.

5.2. Imitation: a global model

Our Global Imitation Model (Fig. 8) is a preliminary model of primate imitative ability (Billard, 2000). Our aim was to build a comprehensive, but simplified model of the visuo-motor pathway behind learning by imitation, from processing real video data to directing a complete dynamic simulation of a humanoid.

This model is biologically inspired in its function, as its composite modules have functionalities similar to that of specific brain regions, but the modules are composed of artificial neural networks whose neurons are little constrained by biological data. It is loosely based on neurological findings in primates and incorporates abstract models of some brain areas involved in visuo-motor control, namely STS, the spinal cord, the primary motor cortex (M1), the dorsal premotor area (PMd) and the cerebellum. The model is implemented in a dynamic simulation of a 65 degrees of freedom avatar which allows us to approximate the motions of a humanoid.

The STS module takes as input the Cartesian coordinates of each joint of the demonstrator’s limbs in an exocentric frame of reference. It then transforms these coordinates to a new set of coordinates relative to an egocentric frame of reference. Our assumption of the existence of orientation-sensitive cells in an egocentric frame of reference in STS is sustained by neurological evidence in monkeys (Perrett, Harries, Mistlin, & Chitty, 1989; Perrett, Harries, Bevan et al., 1989) and humans (Andersen, Lawrence, Bradley, & Xing, 1997; Kertzman, Schwarz, Zeffiro, & Hallett, 1997; Vallar et al., 1999). The vision system also incorporates a simplified attentional mechanism which triggers whenever a significant change of position (relative to the position at the previous time step) in one of the limbs is observed. Note that, at this stage of the modeling and given the simplicity of this module, the attentional module does not relate to any specific brain area. The attentional mechanism creates an inhibition, preventing information flow from M1 to PMd and further to the cerebellum, thereby allowing learning of new movements only when a change in the limb position is observed. In previous experiments, we used video recording of human motion to drive the imitation experiments, see Billard (2000).

Motor control in our model is hierarchical with, at the lowest level, the spinal cord module, composed of primary neural circuits (central pattern generators (CPGs); Stein, Grillner, Selverston, & Stuart, 1997), made of *motor neurons* and *interneurons*. The structure of the CPGs is strongly inspired by Ijspeert’s model of vertebrate spinal circuits (Ijspeert, Hallam, & Willshaw, 1999), and encode simple rhythmic movements of arm and legs for open-loop walking, as well as primary motor behaviors, such as retracting and stretch arm movements of the humanoid avatar. The motor neurons in our simulation activate the avatar’s muscles. We model two muscles per degree of freedom per joint. Each muscle is represented as a spring and a damper (Lacquaniti & Soechting, 1986).

The M1 module monitors the activation of the spinal networks. Nodes in M1 are distributed following a topographic map of the body. Learning of movements is done in the PMd and cerebellum modules. These modules are implemented using the Dynamical Recurrent Associative Memory Architecture (DRAMA) (Billard & Hayes, 1999) which allows learning of times series and of spatio-temporal invariants in multi-modal inputs. Finally, the decision module controls the passage between observing and reproducing the motor sequences, i.e. it inhibits PMd neural activity due to STS (visual) input to flow downwards to M1 (for motor activation). It is implemented as a set of if-then rules and has no direct biological inspiration.

Neurons in the PMd module respond to both visual information (from STS) and to corresponding motor commands produced by the cerebellum. The STS-PMd-M1 interconnection is a simplified version of the mirror-neuron model described in Section 6.1.

Each submodule of the model is implemented at a connectionist level. In the spinal cord module, we use Eq. (1) to model each neuron unit as a leaky integrator, whereas in the M1, PMd and cerebellum modules, we follow the DRAMA neural architecture (see Billard & Hayes, 1999 for a complete description). This differs from the Hopfield network in several ways, especially in its different topology, training and retrieval rules. Each connection of the network is associated not only with a weight (as usual) but also time parameter. Both are updated during learning. The time parameters represent the decay of neural activity along the synapses, while the weights keep a memory of the frequency of neural activation (weights are increased following classical Hebbian rules, i.e. whenever two neurons are simultaneously co-activated). The DRAMA network provides a general ability for learning complex time series by recording explicitly the temporal and spatial aspect of the patterns in its two different sets of connection parameters.

5.3. Synthetic fMRI for grasping and finger tapping

We present here two synthetic fMRI experiments on imitation of grasping and finger tapping movements. The

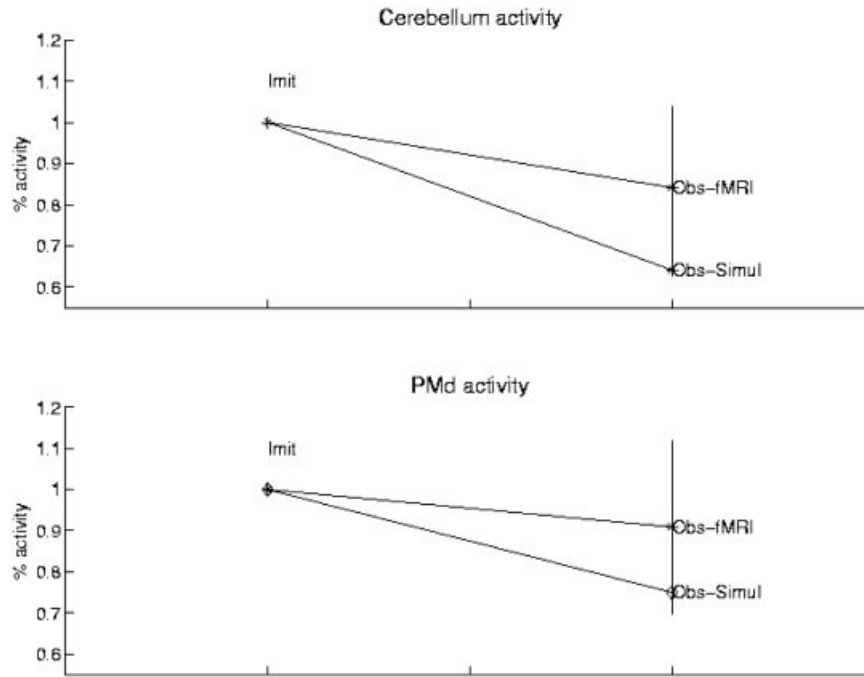


Fig. 9. Neural activity in cerebellum and PMd modules during observation and imitation.

second experiment was meant to reproduce the fMRI experiment described in Section 5.1 on human imitation of an index finger tapping task. We use a pair of demonstrator and imitator avatars. In the grasping experiment, the demonstrator avatar moves its right arm from rest position (alongside its body) to reaching a horizontal and frontal position with all fingers closed in a strong grasp. The finger tapping experiments consists of having the demonstrator lifting upwards four times its index finger. This experiments starts with the avatar’s right arm lifted in the horizontal, frontal position, all fingers relaxed.

During each demonstration, the imitator avatar processes, in the STS module, the simulated Cartesian coordinates of the demonstrator’s limbs and learns the demonstrated motion by processing the information upwards to the PMd and cerebellum modules. Once the demonstration is finished, the imitator avatar is allowed to reproduce the sequence of movements by reactivating the node in the cerebellum which had first been activated during the demonstration. This automatically leads to sequential reactivation of the other relevant nodes in the Cerebellum and PMd modules. The PMd activity projects downwards to M1 and spinal cord to activate the imitator’s muscles. Fig. 9 shows the activity in cerebellum and PMd modules during observation and imitation.

Table 1
Raw data of synthetic fMRI for imitation of grasping and finger tapping

	Grasp	Finger tapping
Cerebellum	0.75	0.64
PMd	0.77	0.75

Computing synthetic fMRI as with our equation for Synthetic PET (Eqs. (2) and (3)) we calculated the raw fMRI activity ($fMRI_A$) of the cerebellum, PMd, M1 and STS modules during observation and imitation tasks. Table 1 shows the raw values for cerebellum and PMd for the two experiments. Data on synthetic grasping are predictive, and could be compared to future fMRI experiments. In the following, we compare synthetic data on finger tapping to the fMRI experiment described in the previous section.

Fig. 10 shows the ratio of activity of the four regions for real and synthetic fMRI data during the finger tapping task. The error bars on the real fMRI data are the standard deviation calculated across the 12 subjects and the four trials (per subject). In both our model and the real data, M1 has the biggest activity. The ratio of activity across the four regions in our model is comprised within the error margin (or close to it, as it is the case for STS). This qualitative similarity of the model to the data is obtained by adjusting the number of connections across modules and the encoding (the number of nodes active) in each module.

Fig. 10 shows the comparative fMRI activity of the cerebellum and PMd modules (those are the only modules activated during both tasks) for imitation against observation tasks in the finger tapping experiment. This is compared to the real fMRI activity of the same regions as measured by Iacoboni et al. (1999). The reason why the ratio in simulation is not equal to one, even though the same nodes in cerebellum and PMd are activated in both tasks, is due to the fact that each node stays at the maximum of activity slightly longer in the retrieving/imitation task than during the observation. This is an artifact due to the neural activation function in the DRAMA architecture which allows for




Fig. 10. Ratio of synthetic and real fMRI activity in cerebellum, PMd, and STS against that of M1.

an imprecision in the timing of the retrieval of the neural activation (Billard & Hayes, 1999). This effect can be observed in Fig. 11, where one sees that the neural activity in cerebellum and PMd during imitation follows a small plateau at the maximum of activity before decaying, which it does not do during observation. This artifact is meant to compensate for random noise in the visual input

during learning. In short, our model assumes an equal (plus a random error) activity between observation and imitation in cerebellum and PMd modules.

The synthetic activity returned consistently lower ratios than the fMRI data. This may be due to two factors. First, as we noted in Section 2.2, BOLD fMRI is probably affected by both blood flow and blood volume changes (Hess et al.,

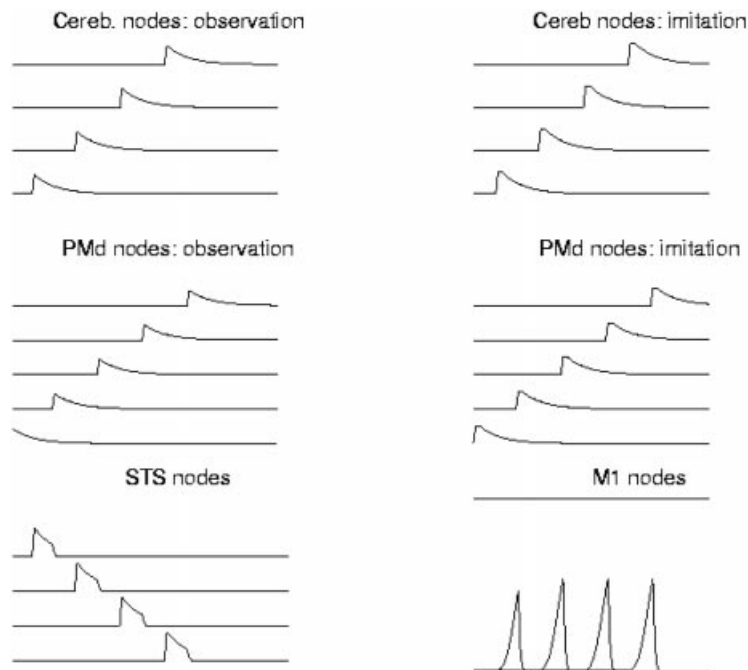


Fig. 11. Ratio of synthetic and real fMRI activity in cerebellum and PMd, observation against imitation task.

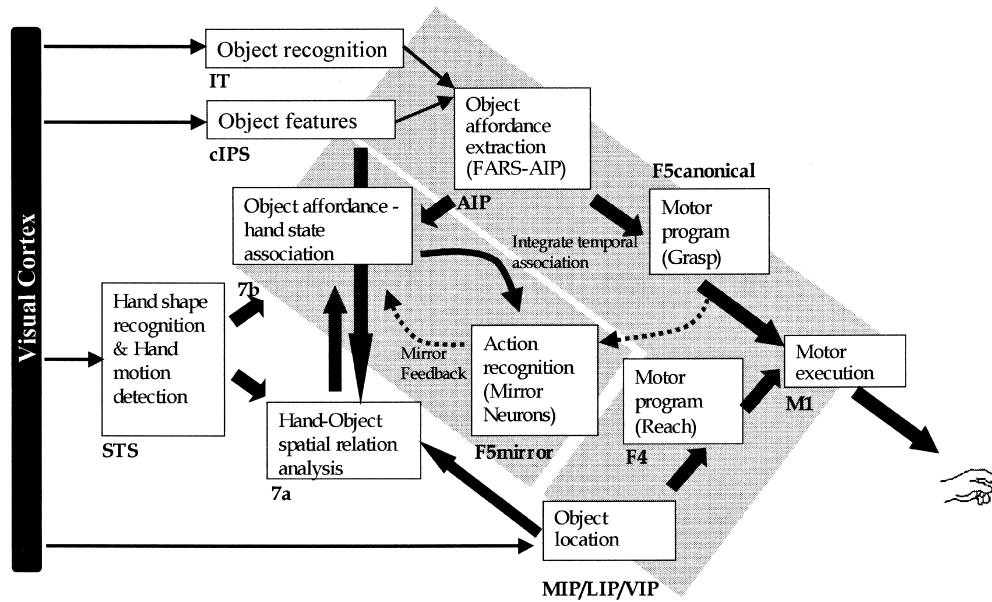


Fig. 12. Schemas of the mirror neuron model with brain regions assigned. The three regions with shared background at right are incorporated from the FARS model (Oztop et al., 2000).

2000). The model does not explicitly account for blood volume effects, given that these effects are not closely linked to neural activity. Second, empirical data on MT neurons have suggested that the neuronal population activity may be mainly driven by low firing rate neurons the activity of which is unaffected by the task or the stimulus (Scannell & Young, 1999). Thus, task-related changes in neural activity may result in small changes in the overall neuronal population activity, determining smaller changes in BOLD signal, compared to expected values based on theoretical considerations or single-unit data. Both factors may conjointly contribute to higher ratios for fMRI data than for synthetic data. The point is not to claim that we have achieved the final model of imitation, but rather to highlight the importance of being able to use multiple sources of data in validating a model, and the utility of Synthetic fMRI in bringing human brain imaging to bear in developing a model of “the supporting neural networks”. Section 6 discusses further issues in developing the research program exemplified in the present article.

6. The many facets of global modeling

6.1. Modeling the mirror system

As is clear from Fig. 6, brain imaging implicates a large number of brain regions in even apparently simple tasks, and we have discussed only a few of the regions in Fig. 6 explicitly in this article. In view of this, we quickly review several studies from the Arbiblaboratory (of course the work of many other laboratories could have been marshaled to

make the same point) to indicate the immense challenges in integrating models of different portions of the nervous system which address detailed neurophysiological data from studies of primates and other species. Such studies are necessary to complement the high-level views of human regional brain activity provided by PET and fMRI studies. We focus on modeling of the Mirror System, then note (*very* briefly) the challenges of modeling other systems relevant to brain activity during imitation, such as the role of basal ganglia in sequential behavior, and the role of cerebellum in the coordination of reach and grasp.

To complement the model of Section 5.2, we offer a high-level view of a model (Oztop, Rizzolatti, & Arbib, 2000) of the function and development of the mirror system in terms of interacting schemas which should be equally applicable to studies of human and monkey. The schemas to the left of Fig. 12 implement the visual system of the model. The *Hand shape recognition* schema recognizes the hand shape (e.g. precision grasp) whereas the *Hand motion detection* schema locates the hand moving in space and sends related information to the *Hand–Object spatial relation analysis* schema which also receives object-related signals from the *Object features* schema. The *Hand–Object spatial relation analysis* schema is needed because, in order to elicit mirror neuron activity in the monkey, seeing the agent and object is necessary but not sufficient — a hand mimicking a matching grasp would fail to elicit mirror neuron activity unless the hand’s trajectory were taking it toward an object with a grasp that matches one of the affordances of the object. The output of this visual analysis is relayed to the *Object affordance–hand state association* schema. This, together with the *Action recognition (mirror neurons)* schema forms the

core mirror (learning) circuit, marked by the gray slanted rectangle in Fig. 12, which mediates the development of mirror neurons. Note that the *association* schema is also the target of the *Hand shape recognition* schema as well as the *object affordance* signals supplied by the AIP module of the FARS model. In this way, the association schema has access to complete (visually observed) *hand state* information.

Details of the implementation of these schemas are beyond the scope of this article. Instead, we offer some preliminary suggestions (also shown in Fig. 12) on the neural localization of the various schemas.

Object Recognition schema: We follow the usual assumption that this is localized in inferotemporal cortex (IT).

Object features schema: Sakata, Taira, Kusunoki, Murata, and Tanaka (1997) and Sakata, Taira, Murata et al. (1997) discovered binocular visual neurons in the lateral bank of the caudal intraparietal sulcus (cIPS), and the neighboring area V3a. These neurons respond preferentially to a bar, plate or solid object in a particular orientation in space. They also found neurons selective for a particular axis of the objects (AOS neurons) and for particular orientation of the surfaces (SOS, surface orientation selective neurons). Such 3D features of the objects are certainly relevant to a grasp affordance (e.g. axis of object determines the wrist angle of the hand). We thus localize the *Object features* schema in cIPS.

Object affordance extraction schema: Following the FARS assumption, we localize this in AIP.

Motor program and Action recognition schemas: The F5 neurons of the FARS *Motor program (Grasp)* schema are the “canonical” (non-mirror) neurons; while the new *Action recognition* schema includes the F5 mirror neurons. F4 is implicated as coding the target position in the FARS model and we relate this to the *Motor program (Reach)* schema, though we do not yet include this schema in our simulations. For completeness reasons we included not only area F4, but cIPS and IT as well (cIPS replaces PIP in the FARS model, where it was hypothesized that PIP codes object features upon which AIP can extract affordances, whereas IT codes the identity of the object being recognized).

Object Location schema: We distribute this between areas MIP, LIP and VIP. These areas (see Colby & Goldberg, 1999 for a review) triple encode the space around the animal. Although it is hard to assign a single type of encoding to each of these regions a rough separation is possible. VIP is more concentrated on encoding the ultra-near space (less than 5 cm from the face); neurons in MIP are specialized for responding to stimuli within reaching distance; and LIP represents the further space that we explore best with our eyes rather than our hands.

The *Hand shape recognition* and *Hand motion detection* schemas: We propose that STS implements both schemas. Perrett, Harries, Benson, Chitty, and Mistlin (1990) and Perrett, Mistlin, Harries, and Chitty (1990) report neurons in STS responsive to goal directed hand motion. Some of the

STS area cells show responses to goal directed hand motion in a translation/scale/rotation invariant way. Andersen, Asanuma, Essick, and Siegel (1990) show that motion processing input originating from middle temporal area (MT) is channeled via medial superior temporal area (MST) or LIP. MT projects to MST and to other areas in the parietal cortex concerned with visuospatial function. Visual input from V1 is further elaborated in MT, where the firing pattern of neurons reflects the speed and direction of motion of visual targets. Barnes and Pandya (1992) report that area 7a (PG-Opt) is reciprocally connected to STS and suggests that the visuospatial analysis that is associated with posterior intraparietal lobule (which includes 7a) could be amplified in the multimodal regions of STS (TPO and Pga). Therefore, the neurons of multimodal areas of the STS could be envisioned as being concerned with analyzing the position of the organism in relation to the environment (Barnes & Pandya, 1992).

Object affordance–hand state association schema: We localize this in area 7b because recent findings on the existence of mirror and mirror-like neurons in area 7b and connection of this area with the mirror neuron region of area F5 (Fogassi et al., 1998) indicates an intimate relation between 7b and F5 mirror neurons. Anecdotal notes (Fogassi, personal communication, 1999) indicate that there exist 7b mirror-like neurons, which fire for simple arm/hand actions. We propose that PF mirror-like neurons form the building blocks of F5 mirror neurons. The similarity of STS and 7b responses to active hands, combined with the connectivity pattern of STS and area 7 makes the STS–7b circuit suggested in Fig. 12 a plausible approximation to the primate hand shape-motion recognition circuit.

Hand–Object spatial relation analysis schema: The kind of computation we are pursuing may be implemented by multiple regions but, as a speculation, one can propose 7a as a possible area involved in this computation. Area 7a is certainly involved in some sort of spatial coding. Stein (1991) suggests that 7a represent extra-personal space. Siegel and Read (1997) show 7a neuron involvement in the analysis of motion evoked during locomotion or by the *manipulation of objects by the hands*. Andersen, Shenoy, Snyder, Bradley, and Crowell (1999) suggests that 7a represent targets in a *world-centered* coordinate frame. Finally Maunsell (1995) suggests an object location coding in 7a, which is capable of coding the location of objects even without visual stimuli ever falling into the receptive fields of some 7a neurons.

For details the core circuit and details of the learning mechanisms, and especially on the role of the hand state in linking observation of own-hand to other’s hand, see Oztop et al. (2000). Our point here is to emphasize how many brain regions must be considered, and to sample the wealth of data which must be taken into account to ground the neural networks which ground a model of the global brain function involved in mirror activity and imitation, complementing the high-level model of Section 5.2.

However, even a model including all the brain regions as shown in Fig. 12 will be inadequate for at least two reasons:

1. As noted in Section 1.4, high-level schemas may contribute to the motor schemas for many specific actions. It is thus dangerous to say that specific neurons “encode” an action when in fact it may take correlated activity in neurons in a variety of brain regions to “commit” the brain to the specific action. Similarly, Fig. 12 is too rigid in assigning schemas to specific regions. However, modeling must start somewhere, and Fig. 12 provides a good starting point for the development of more distributed models.
2. Brain imaging shows activation of areas not included in Fig. 12, such as the basal ganglia (which were included in the FARS model) and cerebellum (which was not). We thus close this section by briefly citing our modeling efforts on basal ganglia and cerebellum.

We have argued that the basal ganglia may be involved in assisting cortical planning centers in some fashion as well as providing sequencing information, and have thus developed (Bischoff-Grethe, Crowley, & Arbib, 2000) a model which includes the control of arm movements as well as saccades, with a recent focus on control of sequences of skeletomotor actions (Bischoff, 1998; Bischoff-Grethe & Arbib, 2000). We postulate that the basal ganglia’s direct and indirect pathways perform two different roles: the indirect pathway inhibits upcoming motor commands from being performed while the current movement is in progress, while the direct pathway projects the next sensory state back to cortex. This informs supplementary motor area region SMA-proper and motor cortex of the expected next state and allows these regions to switch to the next movement of the sequence.

Our models of cerebellar involvement in motor control hypothesize that where basal ganglia is involved in the phasing in and out of various actions, cerebellum is concerned with the tuning and coordination of movements. In our view, the cerebellum does not act directly on the muscles, but rather acts through Motor Pattern Generators (MPGs) — circuits which combine, e.g. trajectory or rhythmic control with local feedback circuitry. We view the cerebellum as divided into *microcomplexes*. Each microcomplex is a general computational module which combines a patch of cerebellar cortex (Purkinje cells and Golgi cell) with the underlying set of cells in the cerebellar nucleus to which the Purkinje cells project, and the inferior olive cells whose climbing fibers provide error signals for motor learning. The “contextual input” is provided by the parallel fibers, the granule cell axons which provide a nonlinear combination of mossy fiber inputs. The job of the Purkinje cells is to learn to pair the parallel fiber output with a pattern of inhibition of the nuclear cells so as to ensure that these cells better tune the MPGs. The parallel fibers are long enough that — through their shared contextual input — Purkinje cells may learn not only to tune individual MPGs to changing

circumstances, but also to coordinate multiple MPGs (such as those for reaching and grasping) so that a complex movement may be achieved in a smooth and integrated way. In short, we do *not* hypothesize that the cerebellar microcomplex learns to *replace* the MPG. Rather, we view the cerebellum as applying its compensations by modulating MPGs, whether cortical or subcortical, and this compensation occurs on multiple time scales. Further, the compensation patterns can be stored and recalled based on higher level task information.

A cerebellar model must be tested and developed in conjunction not only with an MPG model but also with a “plant model” which contains sufficient complexity to “challenge” the cerebellum model. Schweighofer, Spoelstra, Arbib, and Kawato (1998) and Spoelstra, Arbib, and Schweighofer (2000) showed how the cerebellum may compensate for Coriolis forces and other joint interactions in allowing coordinated control of multiple joints in reaching. In the context of the models of reaching and grasping, and of mirror neurons and imitation described above, the problem is to understand how to extend cerebellar modeling for reaching — and we have shown none of the details here — to encompass not only the coordination of reach and grasp in well-learned movements but also to understand how the cerebellum serves to ensure the graceful coordination of arm and hand and body “on the fly” when we carry out a novel behavior, whether in imitation of an observed behavior or in reacting appropriately to novel circumstances.

This quick tour indicates the immense challenges we will have to face in integrating diverse models of different portions of the nervous system if Synthetic Brain Imaging is to reach its full potential in linking models based on detailed neurophysiology to the high-level views of human regional brain activity provided by PET and fMRI studies.

6.2. Homologies, brain imaging and neural networks

The fundamental benefit of Synthetic Brain Imaging is that it allows for specific predictions of PET or fMRI activity in human brain imaging studies, based on neural network models of behavior constrained by animal neurophysiology and arguments from homology. Since the models themselves are a product of functional anatomy, measured single-unit recordings, and behavioral measurements, Synthetic Brain Imaging provides a powerful bridge between all of these approaches. An additional strength of the Synthetic PET implementation is that the contribution of excitatory and inhibitory influences can be teased apart. Because *synaptic activity* is not the same as *neural activity*, being able to distinguish excitatory from inhibitory influences can be an aid to inferring neural activity from the rCBF measure, possibly clarifying apparent contradictions in rCBF data.

In Section 2.2 we suggested that quantitative differences between integrated synaptic activity and neuronal firing are

often negligible. However, there are as yet no models which adequately link neural activity to the BOLD signal, and so the differences between fMRI and PET methods, which we have minimized in the present study, deserve much future study. In particular, we need to develop the “vampire model” of the neuron to relate synaptic and neuronal activity and the metabolic costs of synaptic change to rCBF.

Synthetic Brain Imaging is sufficiently flexible that it will be possible to have network implementations spanning multiple species. Homologies and differences between species (cf. Bota & Arbib, 2000) can then be tested more rigorously using predictions generated by the Synthetic Brain Imaging, while the human data provide another form of validation of neural network models derived from monkey data. It is certainly satisfying when the human results provide confirmation of the behavior of a model of diverse interacting biological neural networks, but far more is learned when there is a mismatch between model prediction and human results. The challenge is then to use these negative results to further refine and constrain the model and, on this basis, design new experiments for both primate neurophysiology and human brain imaging. The resulting data can then catalyze the development of new global models which synthesize our evolving understanding of the human brain in relation to the details we glean from the anatomy and neurophysiology of the brains of other species.

6.3. Language and “higher cognitive functions”

In “Language Within Our Grasp”, Rizzolatti and Arbib (1998) showed that the mirror system in monkey is the homologue of Broca’s area, a crucial speech area in humans, and argued that this observation provides a neurobiological “missing link” for the long-argued hypothesis that primitive forms of communication based on manual gesture preceded speech in the evolution of language. Their “Mirror System Hypothesis” states that the matching of neural code for execution and observation of hand movements in the monkey is present in the common ancestor of monkey and human, and is the precursor of the crucial language property of *parity*, namely that an utterance usually carries similar meaning for speaker and hearer (or signer and perceiver, etc.). Arbib (2000) refines this hypothesis by suggesting that imitation plays a crucial role in human language acquisition and performance, and that brain mechanisms supporting imitation were crucial to the emergence of *Homo sapiens*.

As already argued, imitation in general requires the ability to break down a complex performance into a coordinated control program of schemas which approximate the pieces of the performance to be imitated. This then provides the framework in which attention can be shifted to specific components which can then be tuned and/or fractionated appropriately, or better coordinated with other components of the skill. This process is recursive, yielding both the mastery of ever finer details, and the increasing grace and accuracy of the overall performance. Arbib (2000) argues

that what marks humans as distinct from their common ancestors with chimpanzees is that whereas the chimpanzee can imitate short novel sequences through repeated exposure, humans can acquire (longer) novel sequences in a single trial if the sequences are not too long and the components are relatively familiar. The very structure of these sequences can serve as the basis for immediate imitation or for the immediate construction of an appropriate response, as well as contributing to the longer-term enrichment of behavior. Of course, as sequences get longer, or the components become less familiar, more and more practice is required to fully comprehend or imitate the behavior. Arbib (2000) then analyzes seven hypothesized stages of evolution: grasping; a mirror system for grasping; a simple imitation system for grasping; a complex imitation system for grasping; a manual-based communication system; proto-speech, characterized as being the open-ended production and perception of sequences of vocal gestures, without implying that these sequences constitute a language; and, finally, with cultural evolution predominating over biological evolution, speech and languages as we currently know them. At each stage, the earlier capabilities are preserved. Moreover, the addition of a new stage may involve enhancement of the repertoire for the primordial behaviors on which it is based.

Development of the full argument would take us beyond the scope of this paper. The point to be emphasized here is that our future work on applying Synthetic Brain Imaging to global brain function can even address areas of cognitive science for which no animal data are available if we employ comparative neuroscience and evolutionary arguments to relate brain regions for “higher cognitive functions” to homologous brain regions in other species.

Acknowledgements

This work was supported in part by a Program Project (P20) grant from the Human Brain Project (P01MH52194), and in part by a Human Frontier Science Program grant to Iacoboni and Arbib. We thank Giacomo Rizzolatti for his contribution to the research program described here. Writing of the paper was also supported in part by an ARC-IREX grant for the study of “Hand Movement Recognition and Language: Assistive Technologies and Neural Mechanisms” while Arbib was a Visiting Professor in the Computer Science Department at the University of Western Australia, and by a personal fellowship from the Swiss National Science Foundation to Aude Billard.

References

- Andersen, R. A., Asanuma, C., Essick, G., & Siegel, R. M. (1990). Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *Journal of Comparative Neurology*, 296, 65–113.

- Andersen, R. A., Lawrence, H. S., Bradley, D. C., & Xing, J. (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience*, 20, 303–330.
- Andersen, R. A., Shenoy, K. V., Snyder, L. H., Bradley, D. C., & Crowell, J. A. (1999). The contributions of vestibular signals to the representations of space in the posterior parietal cortex. *Annals of the New York Academy of Sciences*, 871, 282–292.
- Arbib, M. A. (1981). Perceptual structures and distributed motor control. In V. B. Brooks, *Handbook of physiology, Section 2: the nervous system, vol. II, Motor control, Part 1* (pp. 1449–1480). American Physiological Society.
- Arbib, M. A. (2000). The mirror system, imitation, and the evolution of language. In C. Nehaniv & K. Dautenhahn, *Imitation in animals and artifacts*. Cambridge, MA: MIT Press (in press).
- Arbib, M. A., Bischoff, A., Fagg, A., & Grafton, S. (1995). Synthetic PET: analyzing large-scale properties of neural networks. *Human Brain Mapping*, 2, 225–233.
- Arbib, M. A., Conklin, E. J., & Hill, J. C. (1987). *From schema theory to language*. Oxford: Oxford University Press.
- Arbib, M.A., Fagg, A.H., & Grafton, S.T. (2000). Synthetic PET imaging for grasping: from primate neurophysiology to human behavior. Submitted for publication.
- Arbib, M., & Rizzolatti, G. (1997). Neural expectations: a possible evolutionary path from manual skills to language. *Communication and Cognition*, 29, 393–424.
- Barnes, C. L., & Pandya, D. N. (1992). Efferent cortical connections of multimodal cortex of the superior temporal sulcus in the Rhesus monkey. *Journal of Comparative Neurology*, 318 (2), 222–244.
- Bekkering, H., Wohlschlagel, A., & Gattis, M. (2000). Imitation of gestures in children is goal-directed. *Quarterly Journal of Experimental Psychology, Section A: Human Experimental Psychology*, 53, 153–164.
- Billard, A. (2000). Learning motor skills by imitation: a biologically inspired robotic model. *Cybernetics and Systems Journal* [Special Issue on Imitation in Animals and Artifacts, in press].
- Billard, A., & Hayes, G. (1999). Drama, a connectionist architecture for control and learning in autonomous robots. *Adaptive Behavior*, 7, 35–64.
- Bischoff, A. (1998). *Modeling the basal ganglia in the control of arm movements*. PhD thesis, Department of Computer Science, University of Southern California.
- Bischoff-Grethe, A., & Arbib, M.A. (2000). Sequential movements: a computational model of the roles of the basal ganglia and the supplementary motor area. In preparation.
- Bischoff-Grethe, A., Crowley, M.G., & Arbib, M.A. (2000). Movement inhibition and next sensory state prediction in the basal ganglia. *IBAGS Proceedings*, Chapter 23.
- Bota, M., & Arbib, M. A. (2000). The neurohomology database. In M. A. Arbib & J. Grethe, *Computing the brain: a guide to neuroinformatics*. San Diego: Academic Press (in press).
- Brownell, G. L., Budinger, T. F., Lauterbur, P. C., & McGeer, P. L. (1982). Positron tomography and nuclear magnetic resonance imaging. *Science*, 215, 619–626.
- Colby, C. L., & Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annual Review of Neuroscience*, 22, 319–349.
- Dominey, P. F., & Arbib, M. A. (1992). A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex*, 2, 153–175.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. C., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269, 981–985.
- Fagg, A., & Arbib, M. A. (1992). A model of primate visual-motor conditional learning. *Journal of Adaptive Behavior*, 1 (1), 3–37.
- Fagg, A. H., & Arbib, M. A. (1998). Modeling parietal–premotor interactions in primate control of grasping. *Neural Networks*, 11 (7–8), 1277–1303.
- Fogassi, L., Gallese, V., Fadiga, L., & Rizzolatti, G. (1998). Neurons responding to the sight of goal-directed hand/arm actions in the parietal area PF (7b) of the macaque monkey. *28th Annual Meeting of Society for Neuroscience*.
- Fox, P. T., & Raichle, M. E. (1985). Stimulus rate determines regional brain flow in striate cortex. *Annals of Neurology*, 17, 303–305.
- Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119, 592–609.
- 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., Fagg, A. H., Woods, R. P., & Arbib, M. A. (1996). Functional anatomy of pointing and grasping in humans. *Cerebral Cortex*, 6 (2), 226–237.
- Grafton, S. T., Fagg, A. H., & Arbib, M. A. (1998). Dorsal premotor cortex and conditional movement selection: a PET functional mapping study. *Journal of Neurophysiology*, 79, 1092–1097.
- Head, H. (1920). Aphasia and kindred disorders of speech. *Brain*, 43, 87–165.
- Hess, A., Stiller, D., Kaulisch, T., Heil, P., & Scheich, H. (2000). New insights into the hemodynamic blood oxygenation level-dependent response through combination of functional magnetic resonance imaging and optical recording in gerbil barrel cortex. *Journal of Neuroscience*, 20, 3328–3338.
- Horwitz, B., & Sporns, O. (1994). Neural modeling and functional neuroimaging. *Human Brain Mapping*, 1, 269–283.
- Horwitz, B., Tagamets, M. -A., & McIntosh, A. R. (1999). Neural modeling, functional brain imaging, and cognition. *Trends in Cognition Science*, 3, 91–98.
- Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical mechanisms of human imitation. *Science*, 286, 2526–2528.
- Ijspeert, A. J., Hallam, J., & Willshaw, D. (1999). Evolving swimming controllers for a simulated lamprey with inspiration from neurobiology. *Adaptive Behavior*, 7, 151–172.
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: supervised learning with a distal teacher. *Cognitive Science*, 16, 307–354.
- Jueptner, M., & Weiller, C. (1995). Review: does measurement of regional cerebral blood flow reflect synaptic activity? — Implications for PET and fMRI. *NeuroImage*, 2, 148–156.
- Kadekaro, M., Vance, W. H., Terrell, M. L., Gary, H., Eisenberg, H. M., & Sokoloff, L. (1987). Effects of antidromic stimulation of the ventral root on glucose utilization in the ventral horn of the spinal cord in the rat. *Proceedings of the National Academy of Sciences of the United States of America*, 84, 5492–5495.
- Kertzman, C., Schwarz, U., Zeffiro, T. A., & Hallett, M. (1997). The role of posterior parietal cortex in visually guided reaching movements in humans. *Experimental Brain Research*, 114, 170–183.
- Lacquaniti, F., & Soechting, J. F. (1986). Simulation studies on the control of posture and movement in a multi-jointed limb. *Biological Cybernetics*, 54, 367–378.
- Maunsell, J. H. R. (1995). The brain's visual world — representation of visual targets in cerebral-cortex. *Science*, 270 (5237), 764–769.
- McIntosh, A. R., & Gonzalez-Lima, F. (1994). Structural equation modeling and its application in functional brain imaging. *Human Brain Mapping*, 2, 2–22.
- Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, 198, 74–78.
- Mitz, A. R., Godshalk, M., & Wise, S. P. (1991). Learning-dependent neuronal activity in the premotor cortex. *Journal of Neuroscience*, 11 (6), 1855–1872.
- Oztop, E., Rizzolatti, G., & Arbib, M.A. (2000). Mirror neuron system for grasping: a computational model. Submitted for publication.
- Paus, T. (1999). Imaging the brain before, during, and after transcranial magnetic stimulation. *Neuropsychologia*, 37 (2), 219–224.
- Perrett, D. I., Harries, M. H., Benson, P. J., Chitty, A. J., & Mistlin, A. J. (1990). Retrieval of structure from rigid and biological motion: an analysis of the

- visual responses of neurons in the macaque temporal cortex. In A. Blake & T. Troscianko, *AI and the eye* (pp. 181–199). New York: Wiley.
- Perrett, D. I., Mistlin, A. J., Harries, M. H., & Chitty, A. J. (1990). Understanding the visual appearance and consequence of hand actions. In M. A. Goodale, *Vision and action: the control of grasping* (pp. 163–180). Norwood, NJ: Ablex.
- Perrett, D. I., Harries, M., Mistlin, A. J., & Chitty, A. J. (1989). Three stages in the classification of body movements by visual neurons. In H. B. Barlow, et al., *Images and understanding* (pp. 94–107). Cambridge: Cambridge University Press.
- Perrett, D. I., Harries, M., Bevan, R., Thomas, S., Benson, P. J., Mistlin, A. J., Chitty, A. J., Hietanene, J. K., & Ortega, J. E. (1989). Frameworks of analysis for the neural representation of animate objects and actions. *Journal of Experimental Biology*, *146*, 87–113.
- Raichle, M. E. (1998). Behind the scenes of functional brain imaging: a historical and physiological perspective. *Proceedings of the National Academy of Sciences of the United States of America*, *95* (3), 765–772.
- Rees, G., Friston, K., & Koch, K. (2000). A direct quantitative relationship between the functional properties of human and macaque V5. *Nature Neuroscience*, *3*, 716–723.
- Rizzolatti, G., & Arbib, M. (1998). Language within our grasp. *Trends in Neuroscience*, *21*, 188–194.
- Rizzolatti, G., Camarda, R., Fogassi, L., 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., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, *3* (2), 131–141.
- Ro, T., Cheifet, S., Ingle, H., Shoup, R., & Rafal, R. (1999). Localization of the human frontal eye fields and motor hand area with transcranial magnetic stimulation and magnetic resonance imaging. *Neuropsychologia*, *37* (2), 225–231.
- Sakata, H., Taira, M., Kusunoki, M., Murata, A., & Tanaka, Y. (1997). The parietal association cortex in depth perception and visual control of action. *Trends in Neuroscience*, *20*, 350–357.
- Sakata, H., Taira, M., Murata, A., Gallese, V., Tanaka, Y., Shikata, E., & Kusunoki, M. (1997). In P. Their & H. O. Karnath, *Parietal visual neurons coding three-dimensional characteristics of objects and their relation to hand action, parietal lobe contributions to orientation in 3D space*. Heidelberg: Springer.
- Scannell, J. W., & Young, M. P. (1999). Neuronal population activity and functional imaging. *Proceedings of the Royal Society of London B: Biological Sciences*, *266*, 875–881.
- Schweighofer, N., Spaelstra, J., Arbib, M. A., & Kawato, M. (1998). Role of the cerebellum in reaching quickly and accurately: II. A detailed model of the intermediate cerebellum. *European Journal of Neuroscience*, *10*, 95–105.
- Siegel, R. M., & Read, H. L. (1997). Analysis of optic flow in the monkey parietal area 7a. *Cerebral Cortex*, *7* (4), 327–346.
- Spaelstra, J., Arbib, M. A., & Schweighofer, N. (2000). Cerebellar control of a simulated biomimetic manipulator for fast movements. *Biological Cybernetics* (in press).
- Stein, J. F. (1991). Space and the parietal association areas. In J. Paillard, *Brain and space*. Oxford: Oxford University Press (Chapter 11).
- Stein, P. S. G., Grillner, S., Selverston, A. I., & Stuart, D. G. (1997). *Neurons, networks and motor behavior, A Bradford Book*. Cambridge, MA: MIT Press.
- Tagamets, M. A., & Horwitz, B. (1998). Integrating electrophysiological and anatomical data to create a large-scale model that simulates a delayed match-to-sample human brain imaging study. *Cerebral Cortex*, *8*, 310–320.
- Taira, M., Mine, S., Georgopoulos, A. P., Murata, A., & Sakata, H. (1990). Parietal cortex neurons of the monkey related to the visual guidance of hand movement. *Experimental Brain Research*, *83*, 29–36.
- Taylor, J. G., Krause, B., Shah, N. J., Horwitz, B., & Mueller-Gaertner, H. - W. (2000). On the relation between brain images and brain neural networks. *Human Brain Mapping*, *9*, 165–182.
- Vallar, G., Lobel, E., Galati, G., Berthoz, A., Pizzamiglio, L., & Le Bihan, D. (1999). A fronto-parietal system for computing the egocentric spatial frame of references in humans. *Experimental Brain Research*, *124*, 281–286.
- van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Computations*, *10*, 1321–1371.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, *269*, 1880–1882.
- Yarowsky, P., Kadekaro, M., & Sokoloff, L. (1983). Frequency-dependent activation of glucose utilization in the superior cervical ganglion by electrical stimulation of cervical sympathetic trunk. *Proceedings of the National Academy of Sciences of the United States of America*, *80*, 4179–4183.