



PERGAMON

Neural Networks 11 (1998) 1277–1303

Neural
Networks

1998 Special Issue

Modeling parietal–premotor interactions in primate control of grasping

Andrew H. Fagg*, Michael A. Arbib

USC Brain Project, Computer Science Department, University of Southern California, Los Angeles, CA 90089-2520, USA

Received and accepted 2 April 1998

Abstract

Visual information is processed in the posterior parietal cortex for the hypothesized purpose of extracting a variety of affordances for the generation of motor behavior. The term affordance is used to mean that visual cues are mapped directly to parameters that are relevant for motor interaction. In this paper, we present the FARS model of the cortical involvement in grasping, a model which focuses on the interaction between anterior intra-parietal area (AIP) and premotor area F5. The model represents the role of other intra-parietal areas, working in concert with inferotemporal cortex and F5, to provide AIP with a full range of information from which affordances may be derived. The model also suggests how task information and other constraints may resolve the action opportunities provided by multiple affordances. Our model demonstrates not only that posterior parietal cortex is a network of interacting subsystems, but also that it functions through a pattern of “cooperative computation” with a multiplicity of other brain regions. Finally, through the use of several novel tasks, the model allows us to make specific predictions regarding neural firing patterns at both the single unit and population levels, which aids in our further understanding of information encoding in these brain regions. © 1998 Elsevier Science Ltd. All rights reserved.

Keywords: Neural networks; Grasping; Computational model; Parietal cortex; Premotor cortex; Affordances

1. Affordances and action-oriented perception

Gibson (1955, 1966) observed that the pattern of optic flow, or the movement of features across the retina from moment to moment, contained valuable information that could be used to guide navigation reactively through the environment without prior recognition of objects. Gibson used the term *affordances* to mean parameters for motor interaction that are signaled by sensory cues without invocation of high-level object recognition processes. In our own work, we have stressed the importance of the study of *action-oriented perception* (Arbib, 1972)—emphasizing both the role of affordances and the importance of the organism’s current goals, tasks and internal state in determining what is perceived. We see this approach as an antidote to an overly exclusive focus on object recognition as the goal of human processing. This approach was influenced in no small part by the confluence of ethological, mammalian and human studies in the symposium on “Locating and identifying: two modes of visual processing” (Ingle et al., 1967). The dichotomy of locating and identifying was later linked to primate cortical anatomy in the work of

Ungerleider and Mishkin (1982) who distinguished two systems in the extrastriate visual processing of shape. Both systems have their origins in the primary visual cortex, V1, but one extends ventrally, $V1 \rightarrow V2 \rightarrow V4 \rightarrow$ inferior temporal cortex (IT), and is generally assumed to subserve object recognition. The other extends dorsally from V1 to the posterior parietal cortex (PP) and is characterized as mediating spatial memory. The ventral path (especially the inferotemporal cortex) and the dorsal path (especially the parietal cortex) were thus characterized as the “what” (pattern recognition) and “where” (object location) systems respectively.

Emphasizing that action extends beyond mere localization, Goodale and Milner (1992) studied the ability of a patient (DF), with a ventral lesion, to carry out a variety of object manipulations even though unable to report verbally on, or even pantomime, the object parameters used to guide these actions. On this basis, they suggested that, while the ventral system may still be viewed as serving visual perception, the dorsal system may be more properly called the how system—location (“where”) being only one of many properties needed to determine how to interact with an object. Conversely, Castiello and Jeannerod (1991) studied impairment of grasping in a patient (AT) with a lesion of the visual pathway that left the PP, IT, and the

* Corresponding author. Present address: Computer Science Department, University of Massachusetts, Amherst, MA 01003, USA

pathway $V \rightarrow IT$ relatively intact, but grossly impaired the pathway $V \rightarrow PP$. This patient is, in some sense, the “opposite” of DF—she can use her hand to pantomime the size of a cylinder, but cannot preshape appropriately when asked to grasp it. Instead of an adaptive preshape, she will open her hand to its fullest extent, and only begin to close her hand when the cylinder hits the web between index finger and thumb. Surprisingly, however, when the stimulus used for the grasp was not a cylinder (for which the semantics contains no information about expected size), but rather a familiar object—such as a reel of thread, or a lipstick—for which the usual size is part of the subject’s knowledge, AT showed a relatively adaptive preshape. One thus infers the existence of an $IT \rightarrow PP$ pathway capable of providing a representation of the approximate (“default”) size of a known object to the ventral system. In summary, while the IT of monkey and human provides at best crude information about how to grasp semantically relevant objects, this representation can be refined by the PP which can extract from the visual input a representation of the object’s precise affordances for grasping and manipulation. Note that, unlike Gibson, we imagine several intervening levels of processing between the retina and the extraction of affordances.

In this paper, we examine a variety of neural regions in the monkey cortex that demonstrate an involvement in the vision-for-grasping process, as well as regions in the premotor cortex that participate in the production of grasping movements. We then develop a computational model of how a number of these regions interact with each other in order to (1) extract relevant visual information from a presented object, (2) select an appropriate grasp for the given object and situation, and (3) execute the grasping

movement. Particular attention is given to the neural-level encoding of object and grasp information in the anterior intra-parietal area (AIP) and area F5 (of the inferior premotor cortex). The model hypothesizes that the AIP is responsible for extracting affordance information from the visual stream, while F5 selects one of the corresponding grasps and then manages its execution. The model successfully accounts for the task-dependent, phasic firing patterns observed in these areas. The affordance interpretation of AIP’s role leads to very specific predictions as to how the “how” information is encoded in this region. We explore a number of predictions in the context of novel grasping tasks.

2. Representation of objects and grasps in monkey cortex

Fig. 1 shows a number of areas of interest in the cortex of the macaque monkey. The inferior parietal lobule (IPL) of the parietal cortex in a monkey receives visual inputs from occipito-temporal areas as well as from the visual field periphery of V3 and V2 (Anderson et al., 1991; Baizer et al., 1991). IPL is functionally subdivided into areas 7a, 7b, the second somatosensory area (SII), and several sub-areas buried in the intra-parietal sulcus, including the lateral intra-parietal area (LIP), the ventral intra-parietal (VIP) area, and the AIP area. These areas have specific sensory-motor functions, including those for saccadic eye movements (LIP), ocular fixation (7a), reaching (mostly 7b) and grasping (AIP). A similar modular organization is seen in the motor sector (agranular cortex) of the frontal lobe, which is related to body movements (except for the supplementary eye fields). In particular, F4 and F5 of the inferior premotor cortex are involved in the control of proximal and distal




Fig. 1. Cortical regions in the macaque (figure adapted from Jeannerod et al., 1995). The lateral intra-parietal sulcus (LIP) is involved in oculomotor activity. The anterior intra-parietal sulcus (AIP) is involved in grasping of objects. It receives input from other areas of the posterior parietal cortex, which provide information about the parameters of the object of interest, including location, orientation, shape, and size. In addition, this region has very significant recurrent cortico-cortical projections with area F5 of the inferior premotor cortex.

movements respectively. F5 is the area of agranular frontal cortex involved in grasping in the monkey (Gentilucci et al., 1988) and forms the rostral part of inferior area 6. Its main anatomical connections include AIP and the hand field of the precentral motor area (Muakkassa and Strick, 1979; Matelli et al., 1986).

The following subsections provide a brief review of the neurophysiological database that grounds the FARS model. After presenting a protocol for studying grasp-related neural activity, we show that there are cells in both the AIP and F5 whose firing codes for grasp properties rather than object properties. What the data do not tell us is how these cells are interconnected with each other and a broader network of cortical and subcortical regions to yield a wide variety of grasping behaviors. It is the task of the rest of the paper to offer neurally grounded hypotheses which allow us to form the causally complete FARS model, present a selection of simulations obtained with the model, and then offer several predictions for new experiments designed both to test the model and to extend our insights into parietal–premotor interactions in primate control of grasping.

2.1. A protocol for studying grasp-related neural activity

An important experimental protocol has been developed by Hideo Sakata and coworkers (Taira et al., 1990) to study object- and grasp-related properties of neurons (Fig. 2): a ready signal is given by the turning on of light L1. The monkey responds by placing his hand upon the touch pad that is located directly in front of him. Contact with the touch pad is indicated by the KEY trace. When L1 changes color (the trigger stimulus), the monkey begins the execution of the reaching/preshape movement toward the object that is located in front of him. A second light (L2) turns on as the monkey pulls away from the touch pad. The monkey grasps the object and manipulates it by pulling or pushing in the appropriate direction (SW is also referred to as the object phase). This position is held until L2 changes color




Fig. 2. The Sakata protocol, as described in Taira et al., (1990). Key phase: monkey's hand is in contact with the touch pad. Object phase (SW): monkey grasps the switch. See text for details.

(the secondary trigger stimulus), at which time the monkey releases the object and moves his hand away. If the task is performed properly, the monkey is then rewarded with a squirt of juice.

2.2. Object and grasp coding in AIP

A variety of objects were presented to monkeys trained to perform the Sakata protocol. For each object, four different conditions were used: movement in the light (to manipulate the object), movement in the dark (when the monkey has already seen the object and manipulated it in the light), fixation of the object in the light (with no movement), and fixation in the dark. One set of objects included buttons that had to be pushed, knobs to be pulled, and a joystick that had to be grasped and pushed in a particular direction (Taira et al., 1990). A second set of objects studied includes plates, cones, cubes, cylinders, spheres, and rings of various sizes (Murata et al., 1993).

In monkeys trained to grasp objects requiring different types of grip, about half the AIP neurons related to hand movements fired almost exclusively during one type of grip, with precision grip being the most represented grip type (Taira et al., 1990; Sakata and Kusunoki, 1992; Murata et al., 1993). Many cells in AIP also demonstrate object-specific activity. However, it is important to note that some cells actually respond to several different objects, with varied activity levels. Most often, cells are actually selective to varying degrees for several objects (that often share some common characteristics). Some cells demonstrate specificity toward the size of the object to be grasped, while showing a certain degree of independence from the type of object; other cells demonstrate independence from the size of the object. The fact that there are cells that are both size-selective and size-independent indicates that within a population of cells that code for a particular object, a sub-population of these cells are responsible for capturing size of the object. Finally, a small number of cells shows modulation based upon the object's position and/or orientation in space.

About 40% of the AIP neurons discharge equally well if the appropriate grasping movement is made either in the light with the monkey looking at the object, or in the dark. These cells are referred to as "motor dominant neurons" (Taira et al., 1990). The remaining neurons discharge more strongly ("visual/motor neurons") or exclusively ("visual dominant neurons") in the light. Half of the visual dominant neurons, and a part of the visual/motor neurons become active when the animal fixates the object in the absence of any movement. For these last neurons, the visually effective object and the type of grip coded (assessed in the dark) coincide. They appear, therefore, to match the visual representation of the objects with the way in which the objects are grasped.

Most neurons in AIP also demonstrate phasic activity related to the motor behavior. The identifiable phases

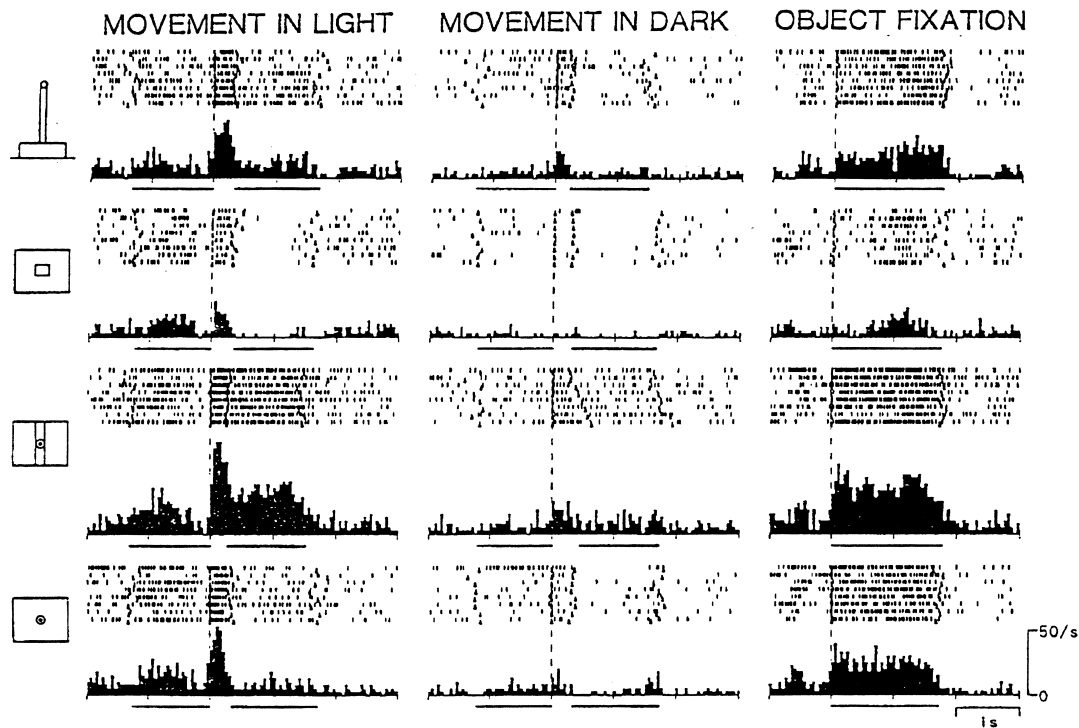


Fig. 3. Visual AIP cell that is most active during fixation or manipulation of a pull-knob inset into a groove (row 3), while somewhat active for conditions involving a joystick (row 1) and a non-recessed pull-knob (row 4). The *key* and *object* phases of movement are indicated for the grasping conditions by the horizontal bars located below the histograms. Figure used by permission of H. Sakata.

include: set (key phase), preshape, enclose, hold (object phase), and ungrasp. Cells participate in varying degrees during different phases of the movement, but are usually most highly active during the preshape and enclose phases of movement. Very importantly, once a cell becomes active, it typically remains active until the object is released.

The AIP cell depicted in Fig. 3 illustrates a number of these properties. All traces are recorded from a single cell, under one of 12 different conditions: four different objects (joystick, push button, pull-knob in groove, and pull-knob), and three different tasks (grasping in the light, grasping in the dark, and fixation of the object in the light). For the grasping tasks, the traces are aligned on the initiation of movement; in the fixation cases, the traces are aligned on the time that the object became visible. The key and object phases of movement are indicated for the grasping conditions by the horizontal bars located below the histograms. The cell depicted in the figure responds heavily to a pull-knob that is recessed into a vertically cut slit, and partially responds to a joystick and a non-recessed pull-knob. Cell activation occurs during both the visual fixation and movement in the light conditions, but responds at an insignificant level during movement in the dark. For this reason, this cell is classified as visual. Note that this cell's activation level is also modulated by the phase of movement. It is most active during preshape and grasp (between the key and object phases), and shows a moderate level of activity prior to movement initiation (the key phase), and in some conditions, following the establishment of the grasp. The drop

of activity at the end of the key phase is thought to relate to the monkey's direction of attention away from the object.

2.3. Grasp coding in F5

Rizzolatti et al. (1988) described various classes of F5 neurons, each of which discharge during specific hand movements (including grasping, holding, tearing, and manipulating). The largest class is related to grasping. The temporal relations between neuron discharge and grasping movements vary among neurons. Some of the units fire only as the monkey's hand encloses around the object (i.e. during finger flexion). Other neurons start to fire with finger extension and continue to fire during finger flexion. Finally, others are activated in advance of the movement initiation (these are referred to as "set" cells), and sometimes cease to discharge only when the object is grasped. Of the neurons active during grasping, 85% were found to be selective for one of precision grip, finger prehension, or whole hand prehension, with precision grip the most frequent (Rizzolatti et al., 1988). "Visual" responses are observed in about 20–30% of F5 neurons, with two separate classes being distinguished. One set of neurons responds to the presentation of graspable objects; another class of cells, referred to as "mirror" neurons, responds both when the monkey makes a grasping movement, and when the monkey observes the experimenter or another monkey make a similar grasping movement (di Pellegrino et al., 1992).

Fig. 4 depicts two F5 cells that are active during a precision pinch grasp, using either the contralateral or ipsilateral

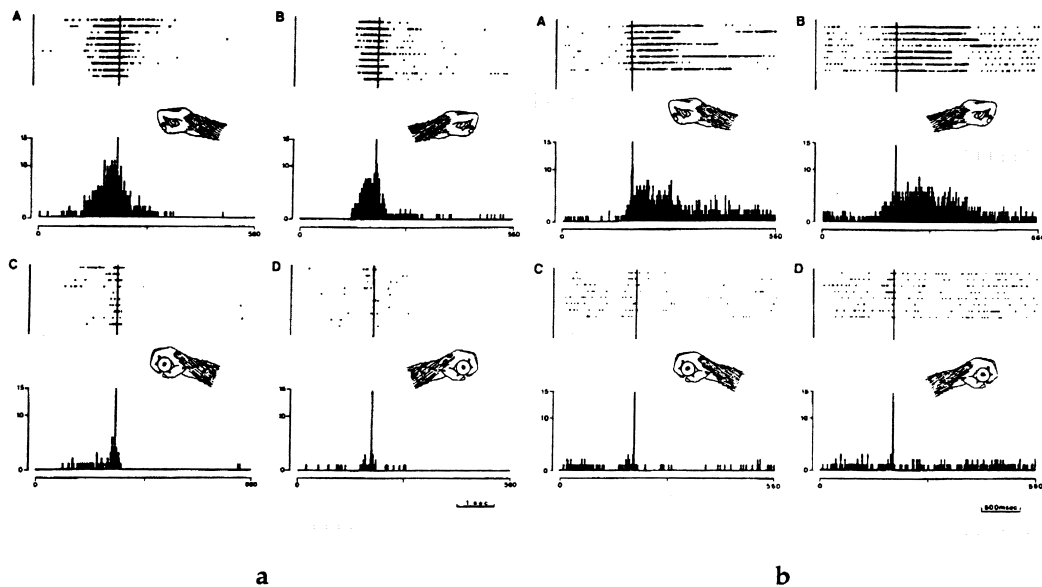


Fig. 4. Two F5 neurons that are active prior to (a) and following (b) object contact. Both cells code for precision grip in both the contralateral (subpanels A) and ipsilateral (B) hands. The neurons do not respond significantly to a palm-opposition-type grasp (C: contralateral; D: ipsilateral). The histograms are centered at the time when the monkey made contact with the object. From figs. 4 and 7 of Rizzolatti et al. (1988).

hand (subpanels A and B respectively). One cell (Fig. 4(a)) begins to fire well in advance of contact with the object, which indicates that it is involved in control of the preshape phase of the grasping movement. The other (Fig. 4(b)) begins to fire shortly before contact is made and continues to fire through the duration of the grasp. When a power grasp is used, the cells are not significantly active (subpanels C and D).

3. The structure of the model

With these data in hand, we now outline an explicit computational model, the FARS (Fagg–Arbib–Rizzolatti–

Sakata) model of the grasping process, based on simplified but biologically plausible neural networks and implemented on the computer to produce detailed simulations (Fagg, 1996). Fig. 5 shows the conceptual core of the model, and Fig. 6 shows the hypothesized pathways between all the regions represented in the complete model. The AIP is hypothesized here as the first stage in the grasp programming process, responsible for integrating object-related information from both the dorsal and ventral systems. We see this process of separating out motor-related features as a way of explaining why cells in the AIP reflect both object- and grasp-related activity patterns.

Fig. 5 illustrates an example of the multiple affordances




Fig. 5. 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).

associated with different ways of grasping an object. In this case, three different features are highlighted. When visual input from the mug is initially presented to AIP via both the dorsal and ventral streams, it computes—according to the FARS model—the set of affordances. The corresponding set of grasps is passed to F5. As a function of additional constraints, F5 selects one of the specified grasps. These constraints include visual information (from the affordances extracted by the AIP), task information, instruction stimuli, and working memories of recently executed grasps. This decision is broadcast back to the AIP, which shunts the other affordances, leaving only the affordance that corresponds to the selected grasp. F5 is then responsible for the high-level execution and subsequent monitoring of the planned preshape and grasp. During the execution of the grasp, the affordance represented by the AIP is continually reinforced by inputs from the active grasp program in F5. These projections allow F5 to update an AIP active memory with information about the grasp that is actually executed, which might differ in some ways from the grasp that was initially specified. Note that we use the term active memory to mean more than the simple holding of information for some delay period so that it can be used later (e.g. Funahashi et al., 1993). Rather, we add that memory may be continuously updated as the result of other computations (in this case, as F5 monitors the progress of the grasp execution). This is similar to the dynamic remapping mechanism we saw in our study of saccades, in which motor afferent signals were used to update a map of targets of potential eye

movements as a saccade was being executed (Dominey and Arbib, 1992).

Based on the spectrum of activity seen in F5 neurons during grasping, we hypothesize that each such neuron codes only a piece of the grasp that is being executed. By piece of a grasp, we mean that the representation of the grasp is distributed across a set of neurons, whose collective responsibility is to produce the appropriate grasp. This distributed coding within F5 is both in time (representing different phases) and in action space (different grasps). With distributed representations, it is important to note that individual neurons often take on multiple roles. We see this in the temporal domain in those neurons that are active during multiple phases of movement. This multiplicity of roles has also been seen in the action domain. Some neurons have been observed in F5 that are active during two distinct motor acts: grasping with the hand and grasping with the mouth. The key point is that despite this lack of discrete coding of actions at the neural level, the system is still able to produce discrete actions as a result of the collective behavior of the neurons. At the same time, because we observe discrete classes of behaviors, this does not necessarily imply that the discreteness should be reflected at the level of the individual neuron.

Fig. 6 gives a high-level overview of the FARS model. In the next few sections, we will discuss the inputs and outputs of the AIP/F5 loop. In addition, we discuss the details of the wiring process for the PIP/IT/AIP/F5 circuit. More details may be found in Fagg (1996).

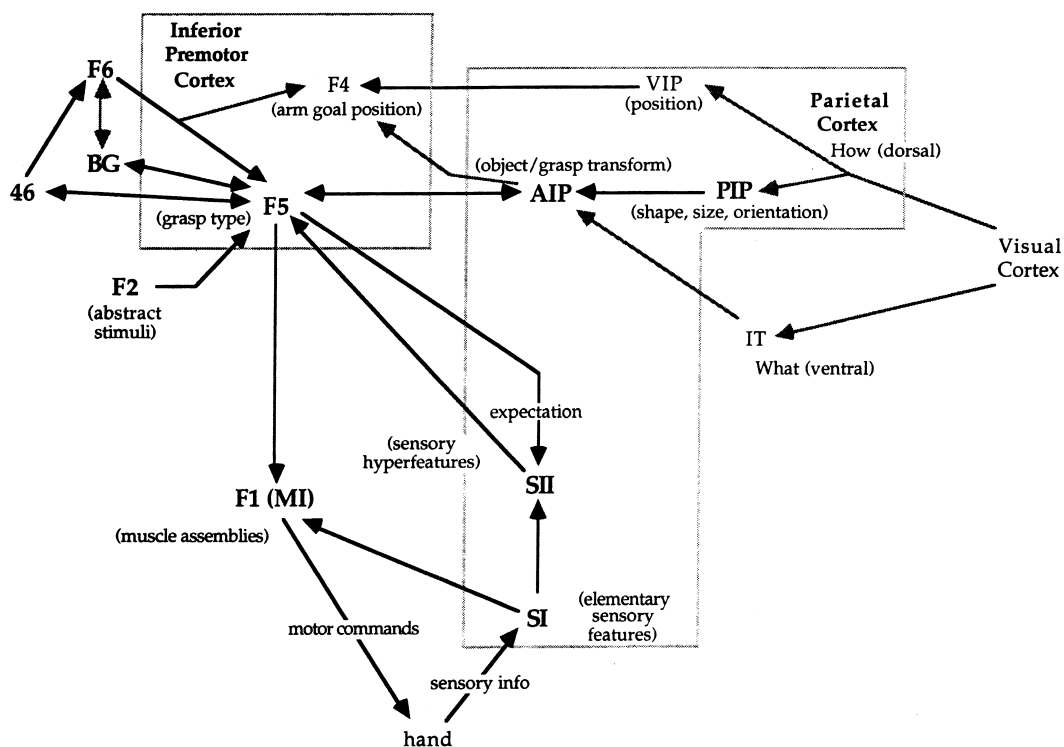


Fig. 6. The complete FARS model. The primary areas are IT (inferotemporal cortex), VIP (ventral intraparietal), PIP (posterior intraparietal), AIP (anterior intraparietal), area 46, BG (basal ganglia), F1 (primary motor cortex), F2 (dorsal premotor cortex), F4 and F5 (of the inferior premotor cortex), and F6 (pre-SMA). Regions and connections in bold are included in the model implementation.

3.1. Modeling the ventral and dorsal visual pathways

The location of target objects is represented in the ventral intraparietal area (VIP) (Colby et al., 1993), using a broadly tuned population code. This suggests that the VIP codes target object position and orientation in space using a population code in a peripersonal coordinate system. This target position is passed to F4, which represents the arm goal position. Within F4, set cells are responsible for setting up the initial reach program, and then movement- and hold-related cells are responsible for executing the movement and maintaining the arm position respectively. However, in the FARS model, we model only grasp programming, and do not model VIP and F4. We have modeled the coordination of reach and grasp elsewhere (Hoff and Arbib, 1993) at the level of interacting schemas; relating such models to neurophysiological data is a goal for current research.

A region neighboring VIP, the posterior intraparietal area (PIP—also known as cIPS) codes object-centered information (Sakata et al., 1997). Here, populations of cells have been shown to fire in response to different shapes that are presented to the monkey. Within a single population of cells, some show specificity towards different parameter ranges (e.g. a cylinder cell might be responsive only to cylinders of width 10–20 mm). More generally, sub-populations of shape cells code the relevant parameters using a population-like coding scheme. On the other hand, individual cells in the IT appear to capture complex combinations of object features (such as shape and color of the object). This, combined with the observation that object representations appear to be rather sparse (Logothetis and Pauls, 1995), implies that relatively different objects are coded by disjoint groups of neurons.

The FARS model does not include the pathways whereby retinal input affects PIP or IT activity, nor does it reproduce the detailed neurophysiological responses seen in these regions. Instead, the model explicitly provides, for each of the visual stimuli used in our simulations, a neural activity code in both the PIP and IT so that these may provide a plausible (though not biologically detailed) encoded visual input to the AIP which is, both literally and metaphorically, where the action starts in our model. In particular, we simplify the IT coding by ignoring the view-direction selective cells observed by Logothetis et al. (1995) and assign a single unit to each object with which the model is “familiar”. The different object classes are represented in the PIP by separate populations of neurons (e.g. cylinder, sphere, block, etc.). Within each population, we use a unit that codes for the general recognition of the shape itself, as well as sub-populations of units to code for the object’s parameters. The parametric information is coded using a population code. Each of the neurons responds maximally to a particular parameter value (referred to as a preferred value). The activity of a parametric neuron is related to the difference between its preferred value and the actual value

being coded. In the case of this model, the activity level is related through a Gaussian function with a specified standard deviation.

Each fundamental object class maintains its own set of parameter sub-populations. As a result, the diameter of a cylinder and the diameter of a sphere are coded using non-overlapping sets of cells. However, this is not to say that the space of all objects can be partitioned into discrete sets. Rather, we imagine these sets as defining a space of objects; objects that mix properties of two classes will have a PIP representation that recruits from both sub-populations.

As an example, a pencil and stick of roughly the same shape and size are coded by the PIP using identical activity patterns. However, the respective codes are orthogonal in the IT. This simplification in the model captures the notion that these objects are conceptually different from one another. In the more general case, we would expect that the IT would represent conceptual classes of objects in an overlapping manner (e.g. the set of all different mugs). On the other hand, a sphere is coded in a manner that is completely orthogonal to the stick and pencil in both PIP and IT.

An affordance is “programmed” into the model by establishing an association from the object’s visual representation in PIP and IT to visual-responsive units in the AIP that are selective to the grasp type and, in some cases, to a specific grasp aperture. In the programming process, the decision as to whether or not to make a connection to an individual AIP unit is determined probabilistically. Thus, only a subset of “appropriate” AIP cells for a particular grasp will respond to an affording visual representation.

Recall from Section 1 the case of the patient AT for whom, no information being available from the dorsal stream (PIP), the projection from the IT can provide, in some cases, the necessary grasp type and parameters. The mapping from object identity in the IT to the AIP is thus provided whenever the nature of the object implies the suitability of a specific set of affordances. This mapping is then somewhat simpler than the mapping from the PIP, since unique objects (and hence object configurations) already give rise to non-overlapping activity patterns in the IT. As a result, the identification of the object can, where the affordances are indeed implicit in the semantics of the object, be mapped directly to both the grasp type and the aperture of grasp. In practice, this projection is important for biasing particular grasp configurations as a function of the specific context presented by the object.

3.2. Interaction of AIP and F5 during execution of the grasp

Fig. 7 demonstrates, at a schematic level, the interaction between populations of cells in AIP and in F5 that is postulated by the FARS model, with particular emphasis placed on the active memory and phasic activation mechanisms. Three AIP units are shown: a visual-dominant cell

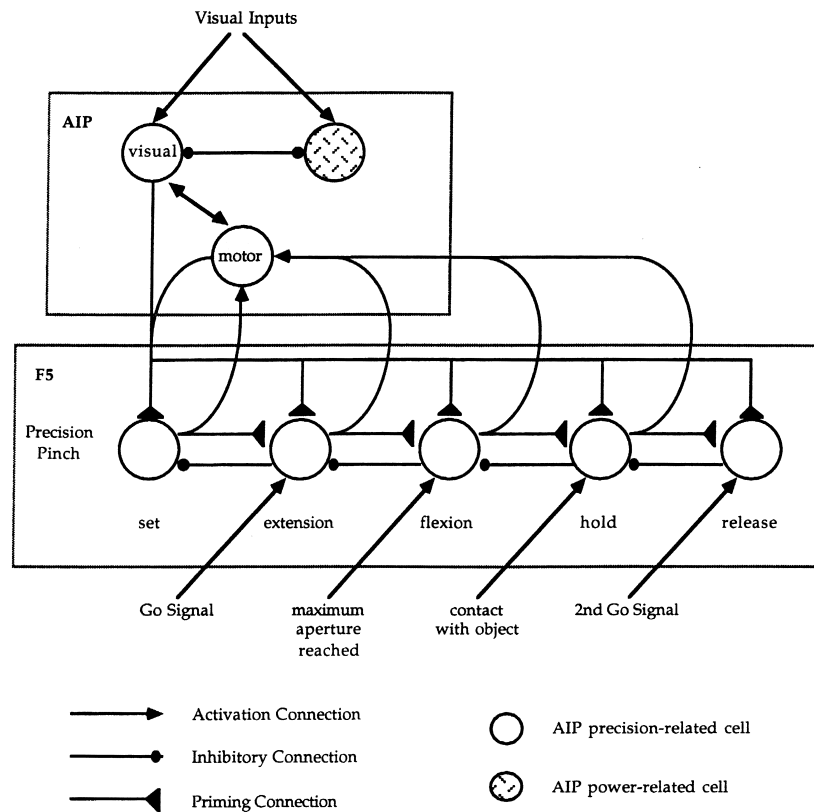


Fig. 7. Interaction between AIP and F5 populations (circles) during execution of the wait–reach/grasp–hold–release program for the Sakata protocol. Activity within F5 cascades from left to right as the program is executed. At each program phase, the state is reported back to the AIP motor-type population. However, the actual circuitry is more complex: the *effective connections* between program states are not coded within F5 but are managed by the combined action of pre-SMA (F6) and the basal ganglia (BG).

which recognizes objects that require a precision pinch, a motor-dominant cell of the same type, and a visual-dominant cell that recognizes objects requiring power grasps. The five F5 units participate in a common grasping program (in this case, for the Sakata protocol using a precision grasp), but each cell fires during a different phase of the program. Each phase of the motor program primes the next, i.e. makes it responsive to other excitatory inputs without necessarily activating the cell. When sufficient priming is present, the actual transition from one phase to the next is triggered either by external events (go signals given by the experimenter), sensory events (contact with the object), or internal events (opening of the hand to its widest extent during the preshape). Activation of F5 cells for one phase of the movement also serves to inhibit the activity of F5 cells encoding the previous phase. Furthermore, as noted in our earlier discussion of coding distributed across time, F5 cells in the model are not constrained to be active during exactly one movement phase, but are allowed to be active over multiple contiguous phases, consistent with the available single-unit data (Rizzolatti et al., 1988).

Although the activity within the “slice” of F5 in Fig. 7 is shown as cascading from left to right as the program is executed, the actual model circuitry is more complex: the effective connections between program phases are not

coded within F5, but are managed by the combined action of pre-supplementary motor area (F6) and the basal ganglia (BG), which represent the high-level wait–reach/grasp–hold–release program. Such indirection is necessary to ensure that a particular action-component can occur in many different motor programs (and not just sequences—see Arbib (1989) for the more general notion of a coordinated control program), rather than always being linked to a unique next state for each “transition trigger”. A discussion of relevant data on pre-SMA is given below.

Within populations of F5 cells that encode a single grasp and phase, cells exchange excitatory connections, thus supporting their mutual coactivation. On the other hand, cells belonging to discordant populations send connections to each another via inhibitory interneurons. These connections ensure that only F5 cells belonging to a single grasp are allowed to achieve a significant level of activation at any one time. When several affordances select multiple grasps in F5, selection of a single grasp is enforced by this competitive mechanism. As we shall see below, this selection process may be biased by other inputs into F5. Also note that, within these constraints, the existence of a connection from one specific cell to another is determined probabilistically. This is intended to model the fact that a single cell is only directly connected to a small subset of cells with which it may be coactive.

The cells within the AIP follow an interconnection scheme similar to that for F5: populations of cells that participate within similar affordances exchange excitatory connections, while inhibitory connections are exchanged between different populations. Excitatory connections are established between a cell in AIP and one in F5 when the encoded affordance matches the grasp. Again, only a subset of candidate connections are actually established.

F5 sends connections to F1 and SII. For a given grasp and movement phase, the active F5 cells recruit those F1 cells that produce the needed distal movements. The projections to SII prime units that detect specific sensory hyperfeatures that will either confirm that the grasp is executing as expected, or that an unexpected event has occurred. In the latter case, F5 will use the return signals to shift execution to a more appropriate program. More details may be found in Section 4.2.

We close this section by noting the highly simplified model of F1 (MI), SI and the hand used in the FARS model. We do not model the arm (see our earlier discussion of F4); the hand is a kinematic model consisting of 15 joints (three for each finger and the thumb). SI contains two classes of units that represent the proprioceptive and kinesthetic state of the hand. The proprioceptive information is represented using a population code; the kinesthetic information is simply binary, capturing the contact/no-contact distinction. Units in F1 drive movements of specific joints as a function of their inputs from SI. F5 specifies a particular, coordinated movement of the hand by selecting a subset of the F1-mediated mappings from the SI state to joint movement. Each SII unit receives from SI both proprioceptive and kinesthetic inputs from all parts of the hand. Because these units have a high threshold, they only become active when the hand is in a particular configuration and is receiving a particular set of tactile inputs.

3.3. Supporting roles played by additional brain regions

We now briefly note the role of other brain regions in the FARS model. We stress again the distinction between empirical data and the model—the former include data on cell firing which suggest possible hypotheses on the functional roles of a given brain region; the model must then make explicit those details of cellular connectivity, etc., that could explain the cellular behavior and mediate the posited functions. Simulations with the model then yield explicit predictions for empirical testing, with the results of new experiments serving to support the model or to suggest ways in which the model can be improved.

3.3.1. SII and expectations

In the grasp versus point comparison in a PET study of humans, Grafton et al. (1996a) find a marked increase of activity in SII, located in the roof of the parietal operculum. The augmented response in this site could not be attributed to a simple sensory effect, such as additional somatosensory

stimulation during grasping, since there was no corresponding increase of rCBF in the primary sensory cortex. Instead, they proposed that the response is related to higher-order tactile feedback or tactile expectation (Arbib et al., 1985) that is used for representing the shape of the target objects and for tactile learning. In detailed anatomical studies in the macaque, Preuss and Goldman-Rakic (1989) found an interconnected network of forelimb and orofacial representations involving the ventral premotor cortex (area 6v), orbitofrontal opercular areas, the opercular portion of area 2, SII, the central insula and area 7b. The functional significance of the strong interconnection between these ventral frontal motor areas and perisylvian somatic areas including SII is not known. Based on the connection with limbic cortex, Preuss and Goldman-Rakic (1989) proposed that the ventral frontal cortex might contain a working memory mechanism of recently encountered objects along with appropriate commands for grasping the objects.

On this basis, SII of the FARS model generates a tactile expectation consonant with the grasp selected by F5 so that any discrepancy between expectation and contact-based feedback can trigger reprogramming of the grasp. This process is implemented in the model through excitatory priming connections sent to SII from F5 cells that are active during the flexion phase of movement and code for a specific aperture. SII units respond to a particular combination of tactile input and hand configuration. When primed, the cells are able to respond quickly to the detection of this combination of inputs. When not primed, although the cells are still able to respond, the response time is much slower. SII units that become active send excitatory signals to F5 units that are involved in the hold phase of the grasp program corresponding to the detected hand configuration. Further details will be discussed in Section 4.2.

3.3.2. SMA (F3 and F6), the BG, and the representation of program sequences in primates

The SMA has been implicated in the planning and execution of complex movements. Human PET studies show a larger degree of activity in SMA during execution of a complex sequence of finger movements than for simple flexion of an individual finger (Roland et al., 1980). A number of studies have also demonstrated an effect in SMA related to mental imagery of complex tasks (Decety et al., 1994; Grafton et al., 1996b).

Based on cytoarchitectonic and microstimulation evidence, Luppino and coworkers argue that there are actually two distinct areas within what has been traditionally called the SMA (Luppino et al., 1990, 1991, 1993). These two regions are referred to as SMA-proper (F3; the caudal region), and pre-SMA (F6). F3 is somatotopically organized and has heavy projections to the limb regions of F1, as well as direct projections to the limb-related portions of the spinal cord. F6 does not project to the spinal cord, and has only moderate projections to areas F3 and F2 (the dorsal

premotor cortex) (Luppino et al., 1990). However, there is a very heavy projection to area F5.

During execution of a pre-learned sequence of saccades, Petit et al. (1996) observed elevated levels of rCBF in both the pre-SMA and the striatum. In a finger movement sequencing task, Boecker et al. (1998) observed increases in rCBF in both pre-SMA and the anterior globus pallidus that correlated with the increasing complexity of the movement sequence. A correlation between sequence complexity and rCBF changes in SMA-proper was not observed.

Recording studies by Rizzolatti have shown F6 neurons that become active when an object that the monkey is about to grasp moves into his peripersonal space (Rizzolatti, 1987). The interpretation of this neural response is that this class of neuron is responsible for generating an internal go signal when it is appropriate for the monkey to begin a reaching movement. In addition, Tanji and Shima (1994) have examined cellular activity in both F3 and F6 during execution of a sequence of hand movements. Not only did they observe cells that responded differentially to the individual movements made during the sequence, but the cells were also sensitive to the order in which the movements were executed. Similar cell behavior was observed in the caudate nucleus by Kermadi et al. (1993) in a task involving production of sequences of either oculomotor or reaching movements.

Based upon an analysis of striatal cell behavior during execution of a button-pressing task, Kimura (1993) concluded that this area was involved in the selection of an appropriate motor program as a function of the behavioral context. Furthermore, in their study of sequence generation in Huntington's disease patients, Georgiou et al. (1995) concluded that the BG are involved in the process of switching from one movement segment to another based on internal cues, but are not responsible for the execution details within each segment.

On the basis of these studies, the FARS model postulates that area F6 (pre-SMA) is responsible for (1) making task-specific decisions about the appropriate motor actions to execute, and (2) representing the high-level sequence of actions to perform (e.g. the sequence of actions required to perform the Sakata grasping task). The BG are thought to provide the machinery for selecting amongst multiple motor programs, and for managing the transition from one step in a motor sequence to the next. In the model, F6 first prepares the ventral premotor regions for execution of the coming reach and grasp by priming both F4 and F5. It then detects the go signal given by the experimenter and initiates execution of the program. Initiation of the release phase of movement is also handled in this way. We do not explicitly model the involvement of F3.

3.3.3. Dorsal premotor cortex (F2)

In a task in which a monkey must respond to the display of a pattern with a particular movement of a joystick, some neurons in F2 respond to the sensory-specific qualities of the

input. However, many units respond in a way that is more related to the motor set that must be prepared in response to the stimulus. F2 is thus thought to be responsible for the association of arbitrary stimuli (e.g. an instruction stimulus) with the preparation of motor programs (Evarts and Wise, 1984). When a muscimol lesion is induced in this region, the monkey loses the ability to make the association correctly, although he is still able to make normal movements of the joystick (Kurata and Hoffman, 1994). In humans, the dorsal premotor cortex has been implicated in conditional grasp selection, in which the color of an LED indicates which of two grasps must be executed (Grafton et al., 1998).

In the FARS model, F2 sends projections to area F5, biasing the selection of the grasp for tasks whose choice is conditional upon presentation of a specific instruction stimulus (IS). This biasing is implemented through low-magnitude excitatory connections from F2 cells encoding the IS to F5 set cells encoding the appropriate grasp. When multiple grasps are afforded by a single object, this biasing is sufficient to ensure that the IS-specified grasp will win the competition.

3.3.4. Area 46

Area 46 has been implicated as a working memory in tasks requiring information to be held during a delay period (Quintana and Fuster, 1993). This memory can participate in the learning of tasks involving complex sequences of movements (Dominey, 1995). Anatomically, this region projects to F6, and also exchanges connections with area F5 (Luppino et al., 1990). In humans, Decety et al. (1994) have recently discovered that area 46 is involved when a subject is asked to imagine herself grasping an object. In addition to this region, they observe activity in area 44 (a possible F5 homologue), as well as at a site along the intra-parietal sulcus (Grafton et al., 1996b).

In the FARS model, area 46, working in conjunction with F6, is responsible for supplying any task-dependent biases for the selection of the grasp in F5. This selection can be based upon the task requirements (such as what is going to be done after the grasp), or based upon a working memory of a recently executed grasp. The biasing can be done at the level of the class of grasp (e.g. power versus precision), or can also include the parameters of the grasp (e.g. width of the aperture). Both options are implemented using low-magnitude excitatory connections into F5.

3.4. Wiring the FARS model

The high-level behavior of individual cells in the FARS model is specified a priori in the form of behavioral descriptors. The connection from one cell to another is automatically determined based upon the compatibility of the descriptors of the cells. The resulting network is intended to give rise to the "general" cell behaviors that have been observed in monkeys (especially for the AIP and F5).

Further variety in behavior is achieved by randomly pruning this network, which allows us to address issues of distributed representation in neural populations. In this section, we give a brief overview of the behavioral descriptors, the number of modeled cells in each region, and the matching process used to determine cell compatibility.

The behavioral descriptors for F5 units are grasp type (precision pinch or lateral pinch), aperture (NONE, 10, 15, ... 45 mm), and phase of activation (Set, Extension, Flexion, Hold, and Release). The activation phase for F5 units can involve multiple contiguous phases. The AIP cells are described as follows: grasp type, aperture, visual/motor orientation (ranges continuously between 0 and 1), and phase of initial activation (Set, Extension, Flexion). Once active, the AIP cells that are involved in the current grasp will continue to be active for the duration of the reach/grasp/hold process.

The number of modeled cells in each brain region is shown in Table 1; the breakdown of the AIP and F5 cells as a function of behavioral descriptor is shown in Table 2. Owing to computational efficiency limitations, the different experiments that are performed with the model rely on different network configurations. The number of cells in AIP and F5 (grasp type, etc.) change from one experiment to another; however, the ratios of basic cell types (e.g. F5 distribution across the different phases) remain the same.

For example, experiments involving only a distinction between a precision and a lateral pinch (of a particular aperture) only involve cells that could potentially participate in a trial (see Table 2, column A). On the other hand, experiments involving many possible object sizes, and hence many corresponding grasp apertures, involve a different collection of F5 and AIP cells (column B). Based upon the battery of experiments performed with the FARS model, we believe that a model involving a full set of cells would perform similarly.

The behavioral descriptors are used to establish the connections between pairs of cells. We give several specific examples here; a more complete description of the connections between different types of cell may be found in Appendix A. Two F5 cells that match in their behavioral descriptors—grasp type, phase, and aperture—exchange excitatory connections. Furthermore, an F5 cell that encodes no aperture will exchange excitatory connections with cells encoding a specific value of aperture—as long as there is a match in the grasp type and phase. Both classes of connections ensure that cells encoding the motor program and phase are coactive during execution of that program. A similar set of restrictions is used in establishing connections amongst cells in the AIP. Connections from AIP to F5 also require a match in grasp type, phase, and aperture; from F5 to AIP, the AIP cell must also have a motor orientation.

Table 1
Number of modeled cells in each region

Layer	Number of units	Comments
PIP	183	
(Object classes)	3	Types: cylinder, sphere, block
(Parameters per object)	1—sphere	Parametric dimensions. The three parameters for a block are length, width, and height
	2—cylinder	
	3—block	
General shape neurons per class	1	Encodes only recognition of shape, with no size information
Neurons per parameter per class	30	Gaussian population code representing size of object along a single dimension
AIP	110–232	Number depends upon the number of grasps (type and aperture) represented. See Table 2
F5	430–750	Depends on number of grasps (type and aperture size) represented. See Table 2
F1	480	Divided amongst the 15 hand joints
SI	178	
Tactile cells	28	Two force-sensitive cells for each finger pad (three for each finger; two for the thumb). One cell of the pair increases in firing rate with force, the other decreases
(Hand joints)	15	Number of modeled joints in the hand
Proprioceptive neurons per joint	10	Gaussian population code for joint position
SII	6–24	One detector for each grasp type and aperture represented in the model
BG	10	Two units for each motor program phase: Set, Extension, Flexion, Hold, and Release
F2	2	One for each grasp type that can be associated with an arbitrary stimulus
F6	5	
READY	1	Detection of ready signal
GO	1	Detection of go signal—initiates movement
GO 2	1	Detection of 2nd go signal—initiates release of object
Grasp bias	2	Task-related bias of grasp type
Area 46	430–750	Same number as in F5

3.5. Information flow through the FARS model

With this background, we can now chart the flow of information through the FARS model as the Sakata protocol is performed. The following description refers to the cells

pictured in Fig. 7, with Fig. 8 showing the actual time course of several F5 cells seen in computer simulations.

Initially some set of visual parameters that is specific to a particular object is encoded within the PIP and IT as described above; these patterns activate the appropriate

Table 2

Breakdown of modeled cells in F5 and AIP as a function of behavioral description for two different experimental configurations. Column A shows the case where a precision and a lateral pinch are represented (with only one aperture); column B corresponds to the configuration in which the many possible apertures are represented for the precision pinch

Description	A	B	Notes
F5 cells (total)	430	750	The larger number of F5 cells is necessary to capture AIP Cells (total) all five phases, whereas phase is not represented spatially (to a significant degree) in AIP.
	110	232	
F5 cell grasp orientation			Ratio of precision to lateral pinch is roughly based upon the observation by Rizzolatti et al. (1988) that the precision pinch is the most heavily represented grasp in F5. Note that there is no lateral pinch representation in experiment ‘‘B’’
Precision pinch	242	750	
Lateral pinch (side opposition)	188	—	
F5 cell parameter coding			A specific representation of grasp aperture has not been examined experimentally, but the FARS model hypothesizes that such a representation is necessary. A wide range of apertures are represented in experiment ‘‘B’’
General (no specificity)	170	82	
Aperture	260	668	
F5 cell phasic responses			Proportion of cells in each phasic class is roughly based upon Rizzolatti et al. (1988). Note that if a cell is active for more than one phase, it is counted multiple times. The largest number of cells is active during movement (Extension and Flexion phases)
set	56	100	
extension	197	278	
flexion	202	286	
hold	65	139	
release	50	115	
AIP parameter coding			Murata et al. (1993) observed both shape- and size-selective responses in AIP. Size-selective responses appeared to follow a population-encoding scheme. These ratios are extrapolated from F5 observations (Rizzolatti et al., 1988)
General (no specificity)	51	24	
Aperture	59	208	
AIP grasp orientation			
Precision pinch	63	232	
Lateral pinch	47	—	
AIP phasic responses			All three onset behaviors have been observed by Taira et al. (1990); the rough proportions are captured here
set	46	98	
extension (early)	46	98	
flexion (late)	18	36	
AIP visual/motor orientation			Taira et al. (1990) observe primarily visual/motor cells, and only a small number of pure visual or motor cells. Note the difference across experiments in the ratio of visual dominant versus motor dominant cells. The model is robust to these differences
Pure visual	11	19	
Visual dominant ($V > M$)	49	135	
Motor dominant ($V < M$)	30	49	
Pure motor	20	29	




Fig. 8. F5 activity during execution of a precision grasp. (a) The top two traces show the joint angle of the most proximal joints of the thumb and index finger, respectively (graphed with opposite signs so as to convey a clear view of aperture changes). The next five traces represent the average firing rate of five F5 neurons (set-, extension-, flexion-, hold-, and release-related). The remaining five traces represent the various external (Ready, Go, Go2) and internal (SII) triggering signals. To illustrate the temporally distributed coding of F5 cells, part (b) shows a number of F5 cells whose activity spreads over more than one phase.

visual-dominant cells in the AIP—Fig. 7 shows a typical cell which is precision-grasp related. A competitive process takes place within the AIP, shunting the activity of cells (in this case side-opposition related cells) that correspond to grasps not appropriate for the presented object. The activation of the AIP visual cell primes a set of cells within F5 that correspond to the grasp to be used (in this case the precision grip). When F6 detects the Ready signal (700 ms after the presentation of the object), the F5 set-related cell becomes active, signaling a preparation for execution of the precision grasp. The set-related activity prepares the sub-program for execution by priming the cell representing the extension phase of the motor program.

When the trigger (Go) signal is given by the experimenter (detected by F6), the extension-related cell begins to fire, causing the preshape of the hand to begin (downward motor commands are not shown). At the same time, the activation of this cell forces the set-related cell to turn off (via the backward inhibitory connection), and the flexion-related cell is primed.

At the time that the fingers reach an extension appropriate for the size of the object to be grasped (as detected by an SII unit), the flexion-related unit is activated, marking the beginning of the enclose phase of movement. In response to this activation, the extension-related cell is turned off and the hold-related cell is primed. The next phase transition occurs when a tactile stimulus is detected (by an SII unit); the hold-related neuron activates, shunting the flexion-related cell, and priming the release-related cell. The hold neuron continues to be active until the external secondary go signal is received (“Go 2” of Fig. 8). This signal causes the

activation of the release neuron and the shunting of the hold neuron.

At each phase of the program, the corresponding F5 unit sends excitatory input to the motor related cell in the AIP. The recurrent priming signal from this AIP cell serves as an active memory of the motor program that is currently being executed. This memory ensures that at any point in the execution of the grasping program, the F5 cells that correspond to the precision grasp (in this example) are activated during the next phase of movement, and not cells that participate in other grasps.

3.6. AIP visual/motor distinction

In the Sakata AIP experiments, a distinction is made between cells that are visual-dominant, motor-dominant, and visual/motor, as determined by the degree to which they participate during a pure fixation task (in the light), or a grasping task in the dark (Taira et al., 1990; Murata et al., 1993). Here we present the model with the same tasks; we first detail these two protocols, and how the model is structured such that the tasks may be performed. We then demonstrate a range of AIP cell behaviors exhibited by the model.

3.6.1. The fixation task

For the fixation in the light condition, the system is presented with the object as in the original protocol; however, neither the ready nor the go signals are ever provided. Hence, the entire 8.4 seconds of the task are spent fixating on the object and no movements are generated.

3.6.2. The dark grasping/working memory task

Before performing the dark grasping task, the model first grasps the object in the light. During the performance of the task in the light, area 46 constructs a memory of those F5 cells that participate in the grasp. This memory is implemented in the model using a latching mechanism. Once the initial grasp is complete, the visual inputs are removed (PIP and IT are cleared), and a second trial is initiated with the presentation of the ready and go signals. During this second trial, area 46 provides excitatory inputs to those F5 cells that were active during the first trial. With this added bias, the F5 cells corresponding to the grasp executed on the previous trial achieve a higher level of activity, and are able to shunt F5 units belonging to other grasps. Note that the area 46 working memory provides a static description of the grasp that was recently executed. By static we

mean that the temporal aspects of the grasp are not stored, providing only a memory of those units that were active at some time during the execution.

3.6.3. Simulation results

Figs. 9–11 show three different types of modeled AIP cells during performance of each of the three tasks. During the fixation task, the model's pure motor cells do not respond to the visible object in the absence of a grasping movement (Fig. 9). In addition, movement in either the light or the dark results in identical activity traces for motor cells. For a pure visual-related cell (Fig. 10), even though this unit does not receive direct projections from any F5 units, its activity level is modulated by the phase of movement. The phasic modulation of this cell during performance of the grasp is due to a combination of excitatory and inhibitory

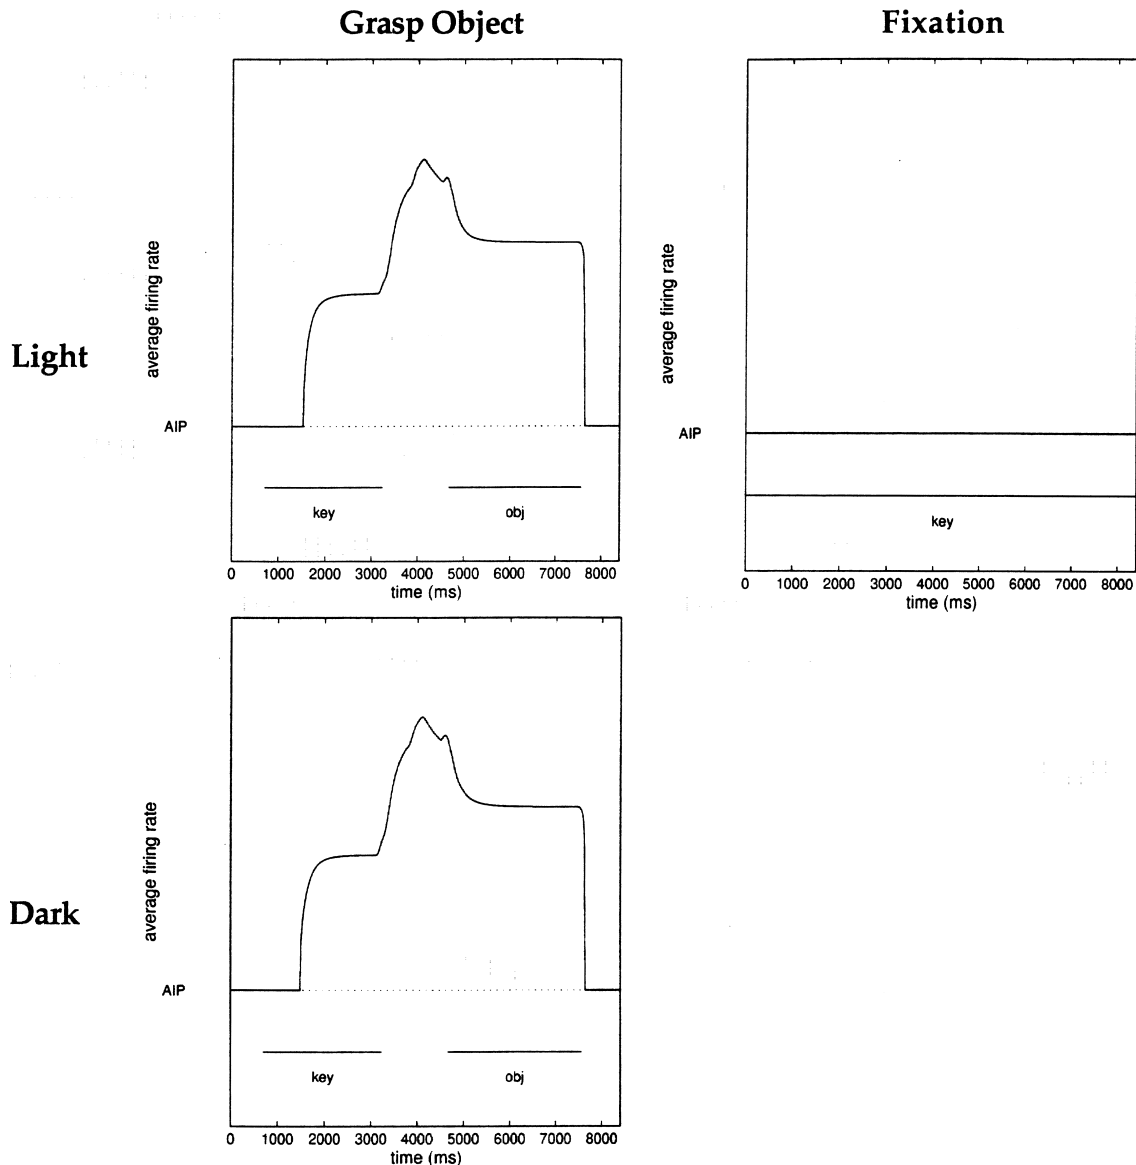


Fig. 9. A pure motor-related AIP cell during performance of three different tasks. The *key* phase corresponds to the presentation of the "ready" signal; the *obj* phase is the time during which the model is grasping the object.

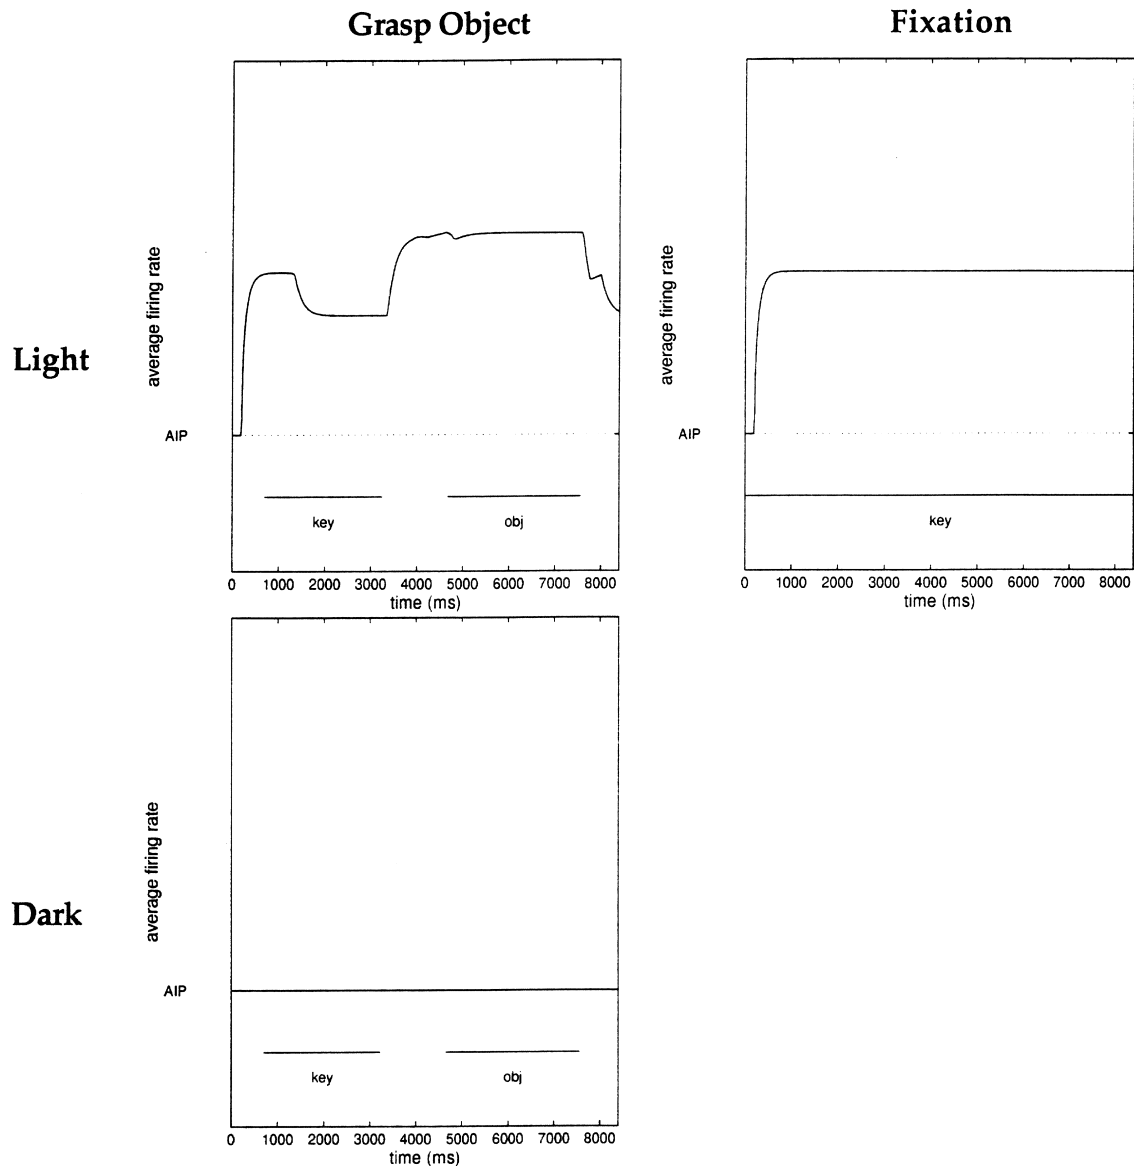


Fig. 10. A pure visual AIP unit. Despite the fact that this unit does not have direct inputs from F5, there is still a modulation of the cell's activation as a function of motor program phase.

signals received from other active AIP cells, some of which are modulated directly by F5 units.

Finally, for an AIP unit whose motor-related responses are modified by visual inputs (hence a visual/motor cell; Fig. 11), during the dark condition, the overall activity of the unit is significantly reduced. As a result, the unit is not able to remain above threshold once the network reaches the hold phase of the motor program.

In summary, the FARS model reproduces at a qualitative level the variety of cell behaviors observed in both the F5 and AIP during several tasks. The visual and motor cell distinction made by Sakata in AIP comes out of the model as a function of an AIP cell's primary inputs. Although most cells receive both visual inputs from PIP and recurrent motor inputs from F5, one class of input sometimes dominates—resulting in the cell being labeled

as either visual or motor. Note that visual cells can show phasic-related responses during the movement-in-the-light condition. We explain this as resulting either from low-magnitude, recurrent F5 inputs, or from inputs from other AIP cells that are being driven phasically. In neither case are these inputs of high enough magnitude to excite the cell without a baseline of activation provided by visual input.

One particular feature of note from the monkey data is the high level of AIP activation that occurs during the preshape and enclose phases of the task. We explain this behavior in the model by first noting Rizzolatti's observation that more F5 grasping cells are active during these two phases of movement than any other phase (Rizzolatti, 1987). This distribution of cell types is reflected in the model. The larger number of active cells during movement implies a higher

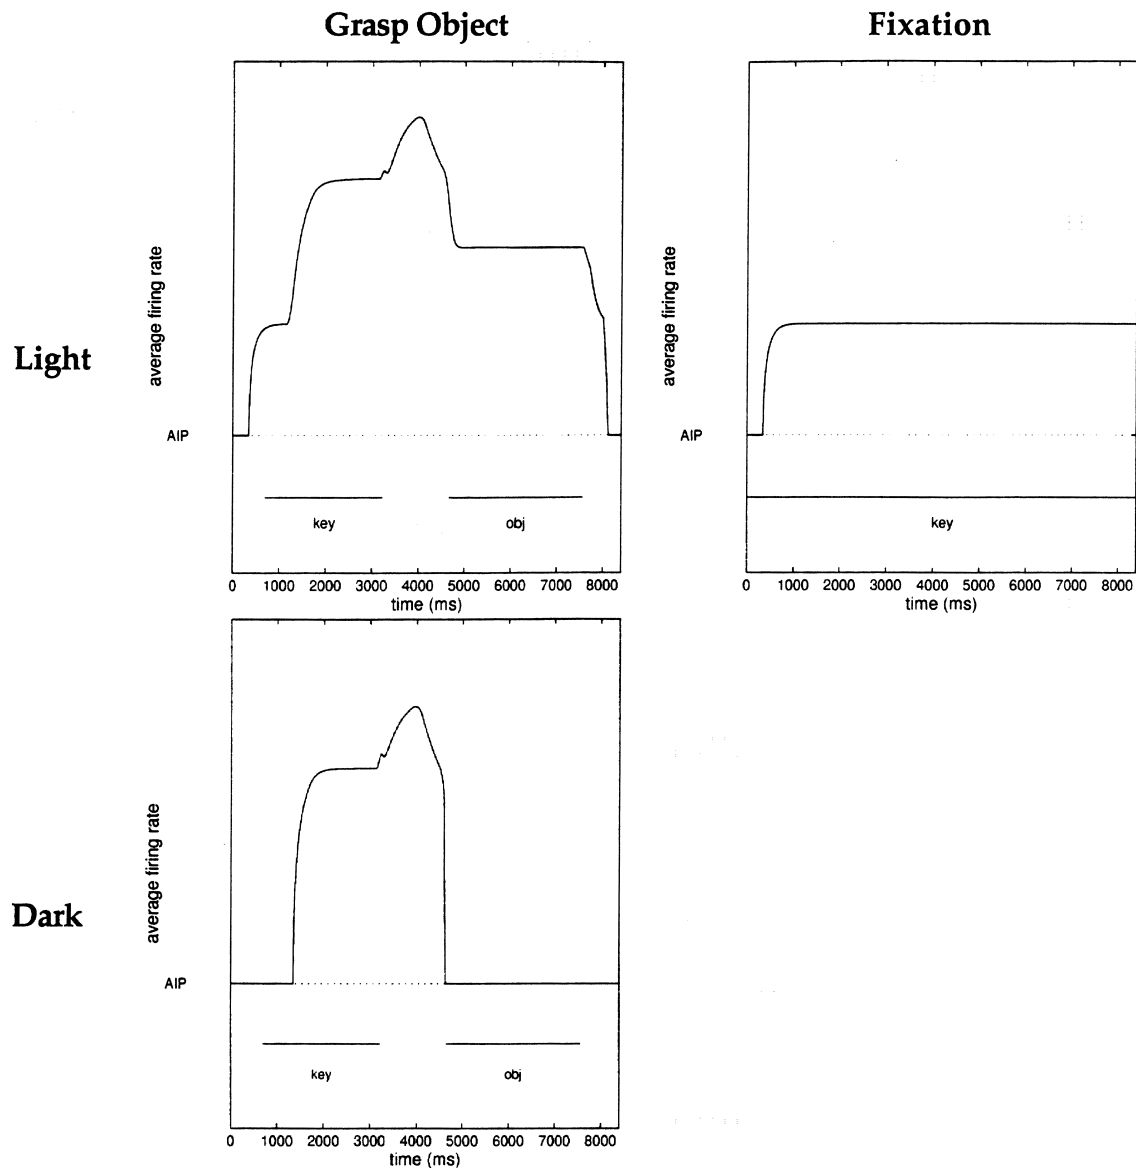


Fig. 11. A visual/motor AIP unit. This unit shows a high level of activation for both motor conditions, and a moderate level of activation for the fixation task.

level of activation being fed back to AIP, which results in the higher activation of individual units in that area.

4. Predictions made by the model

In this section, we present a series of simulation experiments designed to probe further how object and grasp information is utilized and represented in AIP and F5. The behavior of the model in these novel tasks provides a set of specific predictions at both the single-unit and population levels as to what we might expect if the same experiments were performed in monkeys. The first task examines how fine parameters of the grasp (in this case the grasp aperture) might be coded in both F5 and AIP. Then, we look at how these two regions respond when there is a mismatch between the expected object size and the actual object size.

Finally, we examine the issue of grasp versus object coding in F5 and AIP, and ask what aspects of the object description are reflected in these two regions.

4.1. Representation of object and grasp parameters

When humans reach to grasp an object, the maximum aperture achieved by the hand during preshape is related to the anticipated size of the object (Paulignan et al., 1991; MacKenzie and Iberall, 1994). How is this information encoded in the grasping motor programs? Murata et al. (1993) reported that the activation of some AIP cells is modulated by the size of the object that is presented to the monkey. In the model, we extrapolate this observation to F5, assuming that sub-populations of grasp-specific units encode the aperture of grasp, using a population encoding scheme. The AIP cells that are responsive to small objects




Fig. 12. Cylinders of different widths map to a precision grasp of varying aperture size.

exchange excitatory connections with F5 cells that encode small grasps. Within both F5 and AIP, the sub-populations of size specific cells support mutual activation through excitatory connections. Competition between disparate sub-populations is enforced via connections from inhibitory interneurons local to these regions.

When the model is presented with a set of cylinders which range in diameter from 10 to 40 mm (Fig. 12), the maximum aperture achieved during preshape by the index finger and thumb reflects the anticipated diameter of the cylinder (Fig. 13). Within the AIP and F5, some cells behave in a size-specific manner, whereas others are not sensitive to this parameter (but are sensitive to other details, such as the type of grasp). Fig. 14 shows one F5 cell that is active during the set phase of the motor program, but only in cases where a precision pinch of narrow aperture is executed. Note that the cell is responsive to a range of aperture sizes (10–25 mm).

4.2. Object size perturbation

We earlier hypothesized that the AIP maintains an active memory of the grasp that is about to be, or is currently

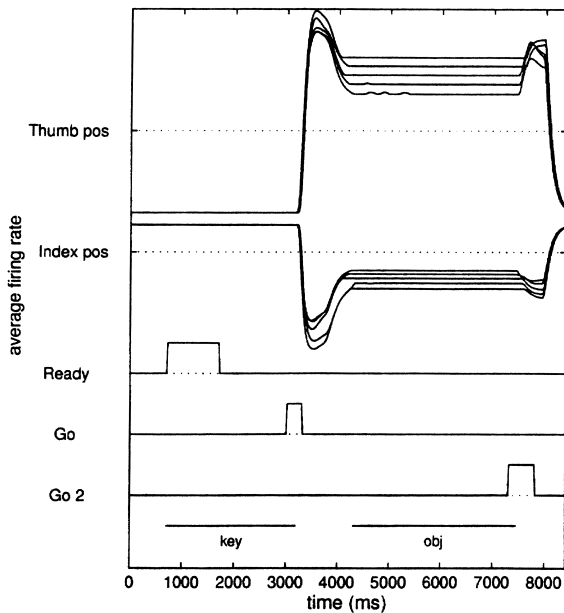


Fig. 13. Thumb and index finger temporal behavior as a function of cylinder size.

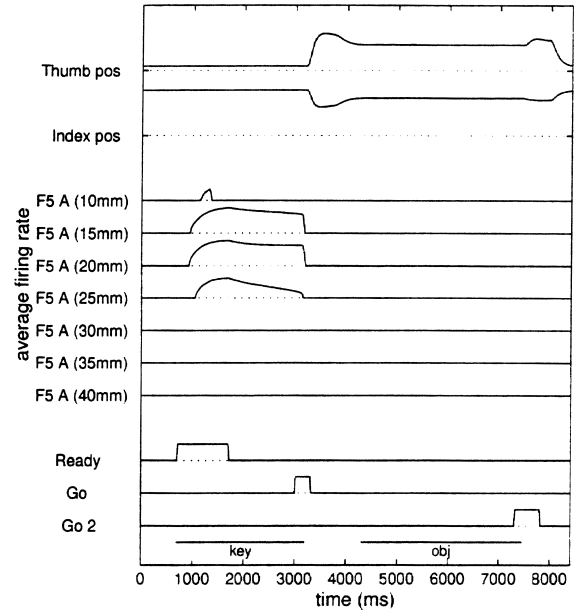


Fig. 14. F5 cell responses during precision grasps of seven different apertures. This particular cell is active only for narrow precision pinches.

being, executed. This requires that any deviations in the executed grasp from that which was planned, must somehow be detected and then reported back to AIP. Take for example the case in which there is a miscalibration between the visually estimated size of the object and the appropriate grasp aperture (as in the size perturbation task of Gentilucci et al. (1995)). In the model, as noted earlier, object size is sensed by SII at the time of contact with the object, so

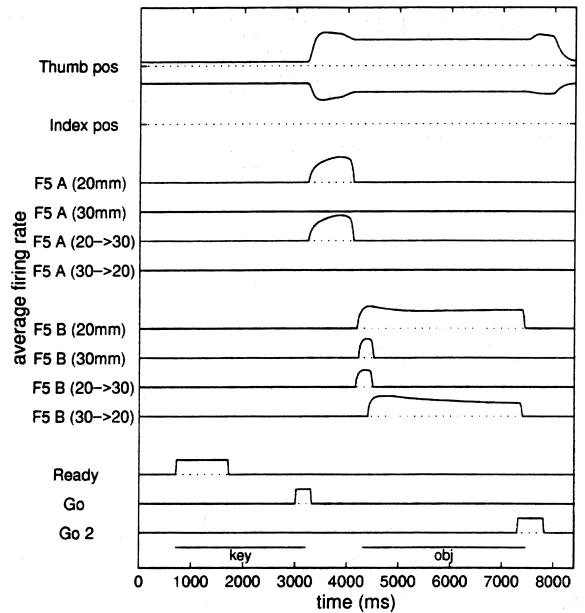


Fig. 15. F5 movement-related cell (A) and a hold-related (B) cell during the perturbation experiment. 20 mm/30 mm traces correspond to presentation and grasping of a 20 mm and a 30 mm cylinder respectively; traces labeled 20 → 30 and 30 → 20 indicate perturbation trials, in which a 20 mm cylinder is switched for a 30 mm cylinder, and a 30 mm cylinder for a 20 mm one respectively.

that a perturbation in object size results in a shift in the SII activity pattern. For small perturbations, the activity patterns for the perturbed and unperturbed cases can be overlapping. In either case, this shift in SII contact triggers input to the F5 layer which causes a different set of F5 hold units to be activated than would normally be the case. This new set of active F5 neurons comprises units that are tuned for holding the object size actually encountered.

This shifting of activity from one grasping program to another is demonstrated in Fig. 15. Average firing rate traces for two units (A and B) are shown for four different conditions (grasping of a 20 mm cylinder, a 30 mm cylinder, perturbation of 20 to 30 mm and perturbation from 30 to 20 mm). Unit A is a movement-related cell, and hence turns off as contact with the object is made. Because the perturbation is not detected until time of contact, the cell behaves in the same manner for both trials involving vision of the 20 mm cylinder. Unit B, on the other hand, is a hold-related cell. Although its preferred aperture is 20 mm, it also activates for a short period of time when a 30 mm grasp is established.

When the cell is presented with a perturbation, its activity pattern switches to match the size of object that is actually grasped. Note, however, that the onset of the cell (comparing the 20 mm and 30 → 20 mm conditions) is delayed by about 200 ms. This delay is due to a combination of (1) the priming of the SII unit in anticipation of contacting a 20 mm object (which yields quicker activation of the SII unit in the non-perturbed case), and (2) the time required to “derail” (shunt) the ongoing execution of the 30 mm grasping program in F5. Because F5 projects back to AIP, the modeled AIP cells demonstrate a similar range of behavior during the perturbation: some cells show drastic shifts in activity shortly after the point of contact, whereas others show little or no change in behavior. We hypothesize that this updated representation in the AIP may provide the basis for the observed adaptation of preshape over multiple perturbed grasping trials (Gentilucci et al., 1995).

4.3. Separating object from grasp coding in the AIP and F5

Although a number of experiments have observed condition-dependent activity in the AIP and F5 (Rizzolatti, 1987; Taira et al., 1990), it is rather difficult to determine exactly what is being encoded by these regions. Specifically, we are interested in the question of whether these regions have a tendency to capture grasp- or object-related information. Many of the experiments performed in monkeys have varied both object and grasp together. Here we present a pair of modeled experiments in which one of grasp or object is held constant, while the other is varied. The results of these simulations allow us to better tease out the answer to our question, and stand as specific predictions as to what we would expect to see in monkey.




Fig. 16. Two objects that map to the identical grasp.

4.3.1. Multiple objects that afford the same grasp

In the first experiment, the model is presented with two different objects that are roughly the same width (Fig. 16). The objects are chosen such that their representations in PIP are non-overlapping (in this case we are using a small cylinder and a narrow plate), but both objects are graspable using the identical precision pinch. We are interested in looking at how AIP encodes the affordances for these two objects, and to what degree the object-specific properties are preserved in this region.

Fig. 17 shows the population response for all AIP cells under two conditions: grasping the plate or the cylinder in the light. Each point corresponds to a single cell, the coordinates of which are determined by the normalized cell activity during performance of the two conditions. (For each unit, the firing rate is integrated over the course of the task, and divided by the maximum of this measure over all cells and both conditions.) Points that fall along the diagonal correspond to cells that do not vary significantly in their activity between the two conditions. In contrast, points that deviate from the diagonal correspond to cells that do vary from one condition to the other. In Fig. 17(A), the AIP population demonstrates a range of object responses. Although a few cells demonstrate large differences in response to the two objects, many are equally active for both. As a comparison, almost every F5 cell (Fig. 17(B)) is equally responsive to both objects.

How does the variety of AIP cell responses relate to the visual/motor cell distinction? To study this question, we compare the population responses for the plate and cylinder conditions under two tasks: fixation in the light (Fig. 17(C)), and grasping in the dark (17(D)). The fixation responses are a direct result of connections from visual areas into AIP, and show significant differences in behavior between the two objects. Cells that fall along the diagonal are those that receive inputs from both cylinder and plate PIP/IT representations; units that are plotted near the horizontal or vertical axes are activated by only the plate or cylinder respectively. On the other hand, under the dark movement task, the active AIP cells (which are necessarily motor or visual/motor cells) do not show object specificity. This is because AIP activation in this case depends on F5 input which was itself determined by the working memory for grasp encoded in area 46.

In general, units that receive a significant amount of visual input (and thus are classified as visual-related cells)

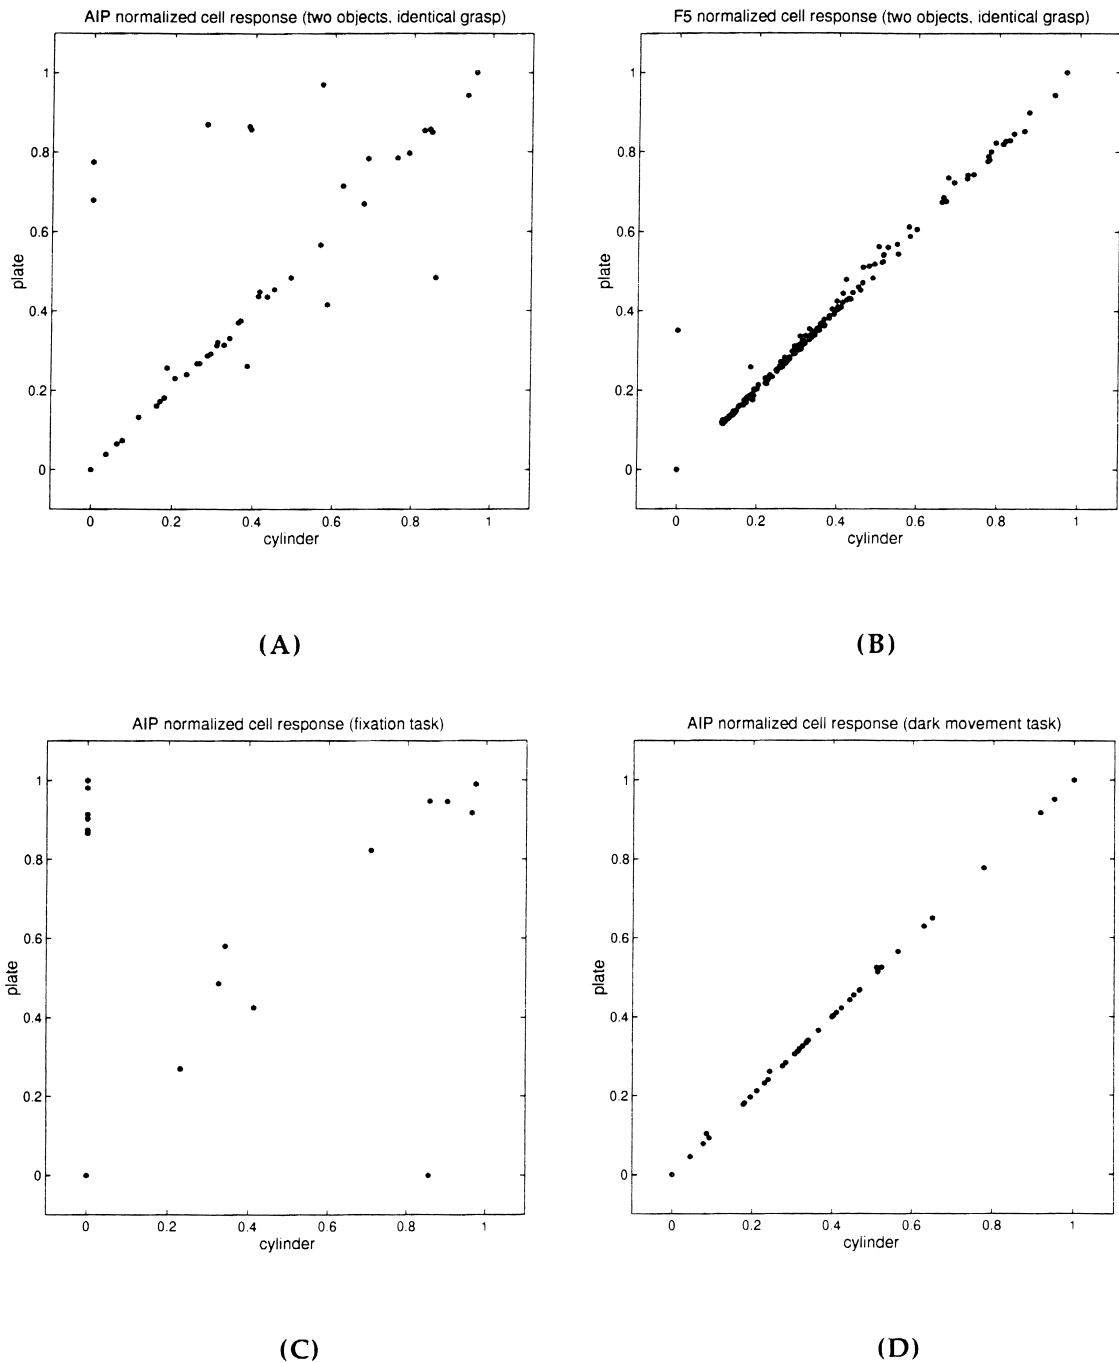


Fig. 17. Comparison of population responses towards two different objects (but identical grasps). Lighted movement task, AIP (A) and F5 (B) cells; and AIP populations during fixation (C) and dark movement (D) tasks.

are more likely than motor-related units to display differences in how they represent the two different objects. As discussed earlier, when a mapping is established from an object representation in the PIP (or IT) to a set of appropriate AIP units, the decision as to whether or not to make a connection from the visual representation to an individual AIP unit is determined probabilistically. In this experiment, because the two objects are represented in the PIP using non-overlapping populations of neurons, only a subset (about half) of the possible AIP units will respond to both

objects (Fig. 18). This aspect of the wiring process (random connections from the PIP to the AIP) accounts for the primary clusters of cells observed in Fig. 17(A,C) (i.e. diagonal, horizontal, and vertical). Within these clusters, deviations from one cell to another, especially in magnitude of response (as indicated by distance from the origin), are due to the randomization in connections within and between the AIP and F5 areas.

For the most part, however, the object distinction is lost at the level of F5, which is the source of the motor-related




Fig. 18. Visual-related AIP cells receive object-specific inputs. Motor-related cells receive recurrent inputs from F5, which do not demonstrate object-specific activity.

activity in AIP. This happens despite the fact that only a subset of possible projections between AIP and F5 are established (also determined probabilistically). The reason is that the motor-related activity flows through the loop formed by the AIP and F5. As the activity makes successive cycles through the loop, more and more cells (from both AIP and F5) are recruited, until most of the cells that belong to a single grasp program become active. As a result, the model predicts very few differences in object coding in the AIP cells that are classified as motor-related.

4.3.2. Multiple grasps afforded by the same object

In general, a single object affords many possible grasps, one of which is selected at the time of execution as a function of the current context (which may include task requirements, position of the object in space, and obstacles). How is this one-to-many mapping captured in the F5/AIP circuitry? In the following simulation experiment, the model is presented with a single object (a small cylinder), and asked to perform one of three tasks (Fig. 19). For each task, the model grasps the cylinder. The three different tasks are:

1. grasp the cylinder using a precision pinch (“non-conditional/precision”);
2. grasp the cylinder using a power grasp (“non-conditional/power”);
3. as a function of an instruction stimulus (the color of a light), grasp the cylinder using either a precision pinch (“conditional/precision”) or a power grasp (“conditional/power”).

When the grasp is known ahead of time (tasks 1 and 2), it is assumed that some higher level planning region




Fig. 19. A single object mapping to two possible grasps. Before execution, one grasp must be selected based upon the current context (e.g. based upon an IS).

predisposes the selection of the correct grasp. In the model, it is area F6 that performs this function. However, when the correct grasp is not known a priori (task 3), some other region must be invoked that is capable of associating an arbitrary IS with the appropriate grasp. As discussed earlier, we implicate the dorsal premotor cortex (F2) in this association process (Mitz et al., 1991; Fagg and Arbib, 1992). The protocol used in this experiment is similar to the one used in the previous set of experiments. The only difference is that for task 3 an instruction stimulus is presented 800 ms after the ready signal.

Fig. 20 shows the behavior of two set-related cells in F5. For both the precision and power conditional cases, the cell activity is identical until the IS is presented. This is expected for all cells, since all system inputs are the same, other than for the IS. After the IS, we see a divergence in the activity pattern. Because these cells are both involved in encoding the precision grasp, they continue to increase in their activity level when the precision condition is specified, while their activity is shunted for the power grasp case. Note that cell A is shunted almost immediately, while cell B remains active until the go stimulus is received. Examining the non-conditional traces, we see that cell A is only active for the precision grasp case. During the power grasp case, this cell is significantly inhibited by other F5 units and thus never becomes active. Cell B, on the other hand, is active in both cases, even though the correct grasp is known before the trial begins (at $t = 0$).

In monkey, we predict that there will be similar behavior at the single-unit level: in the conditional task, prior to the IS, we should see partial set activity corresponding to the possible grasps for this task; following the IS, this activation pattern should give way to strong activation of those cells that will participate in the selected grasp. In AIP, we predict that motor-related cells will show a similar behavior to their F5 counterparts. Visual-related cells, however, will demonstrate less sensitivity to the selected grasp. A more detailed account of this experiment, along with a comparison study of PET activity in human grasping may be found in Arbib et al. (1998).

4.4. Implications of the affordance interpretation

Let us now examine more closely the meaning of object coding in the AIP. One hypothesis put forward in this paper—the idea that AIP extracts affordances from the visual input—has very specific implications as to what aspects of the object are actually captured in the AIP representation. When objects are grasped in particular ways, certain dimensions of the objects are not as relevant as others to selecting the specifics of the grasp (including the grasp type and aperture). In this simulation experiment, we present the model with a series of four different box-shaped objects, each of which must be grasped along the indicated horizontal axis (in Fig. 21). The two boxes in the left column are to be grasped using a precision pinch of aperture 10 mm,

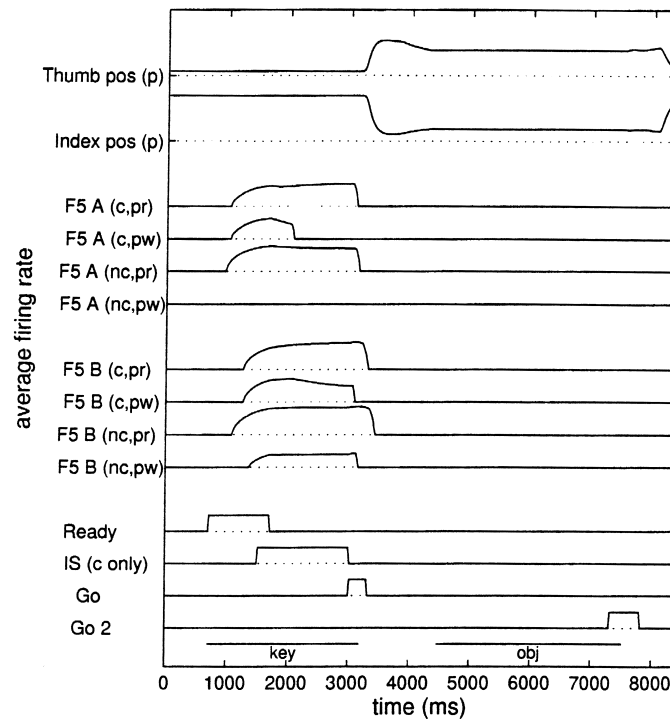


Fig. 20. Two F5 units (A, B) in response to the four conditions: (c,pr), (c,pw), (nc,pr), and (nc,pw). c = conditional; nc = non-conditional; pr = precision pinch; pw = power grasp.

although they differ significantly in size in the other two dimensions. Likewise, the boxes in the right column must be grasped using a precision pinch of aperture 20 mm. The key question, then, is how these four objects are coded by the AIP cell population.

Fig. 22(A) compares the visual-related population responses for the two objects in the leftmost column of Fig. 21 (same aperture). All active AIP units demonstrate a significant level of activity for both objects, with little variation




Fig. 21. Four boxes of different dimensions. Grasping is performed along the horizontal axis. The two blocks in the left column are grasped using a precision pinch with 10 mm aperture; the blocks on the right require a 20 mm aperture.

between the two. The small amount of variation is due to the fact that a few cells are sensitive to object dimensions other than the width along the grasp axis. However, comparing the responses towards the two objects in the bottom row of Fig. 21, we see a much wider distribution of responses, with some cells responding exclusively to one object or the other (Fig. 22(B)). This results from the need to capture the differences in object widths along the axis of the grasp, so that the appropriate grasp aperture may be provided to F5.

Fig. 22(C,D) demonstrates the motor-related AIP population responses for the same pair of objects. The motor-related responses tend to shed any object-specific properties not related to the grasp. For objects requiring identical grasps, the motor-related responses are identical (Fig. 22(C)); objects mapping to different grasps demonstrate grasp-specific modulation (Fig. 22(D)). In this latter case, despite the need to represent different grasps, many AIP motor-related cells are shared between the two objects (which are plotted along the diagonal). These shared units are of two types: (1) general precision grasp units (with no aperture specificity), and (2) aperture-specific cells that respond to both conditions. The latter case is due to the population code that is used to capture the aperture parameter. Thus, the cells that fall along the diagonal capture the commonalities between the two situations, whereas the remaining cells capture the necessary differences.

5. Discussion

The planning and execution of grasps requires the

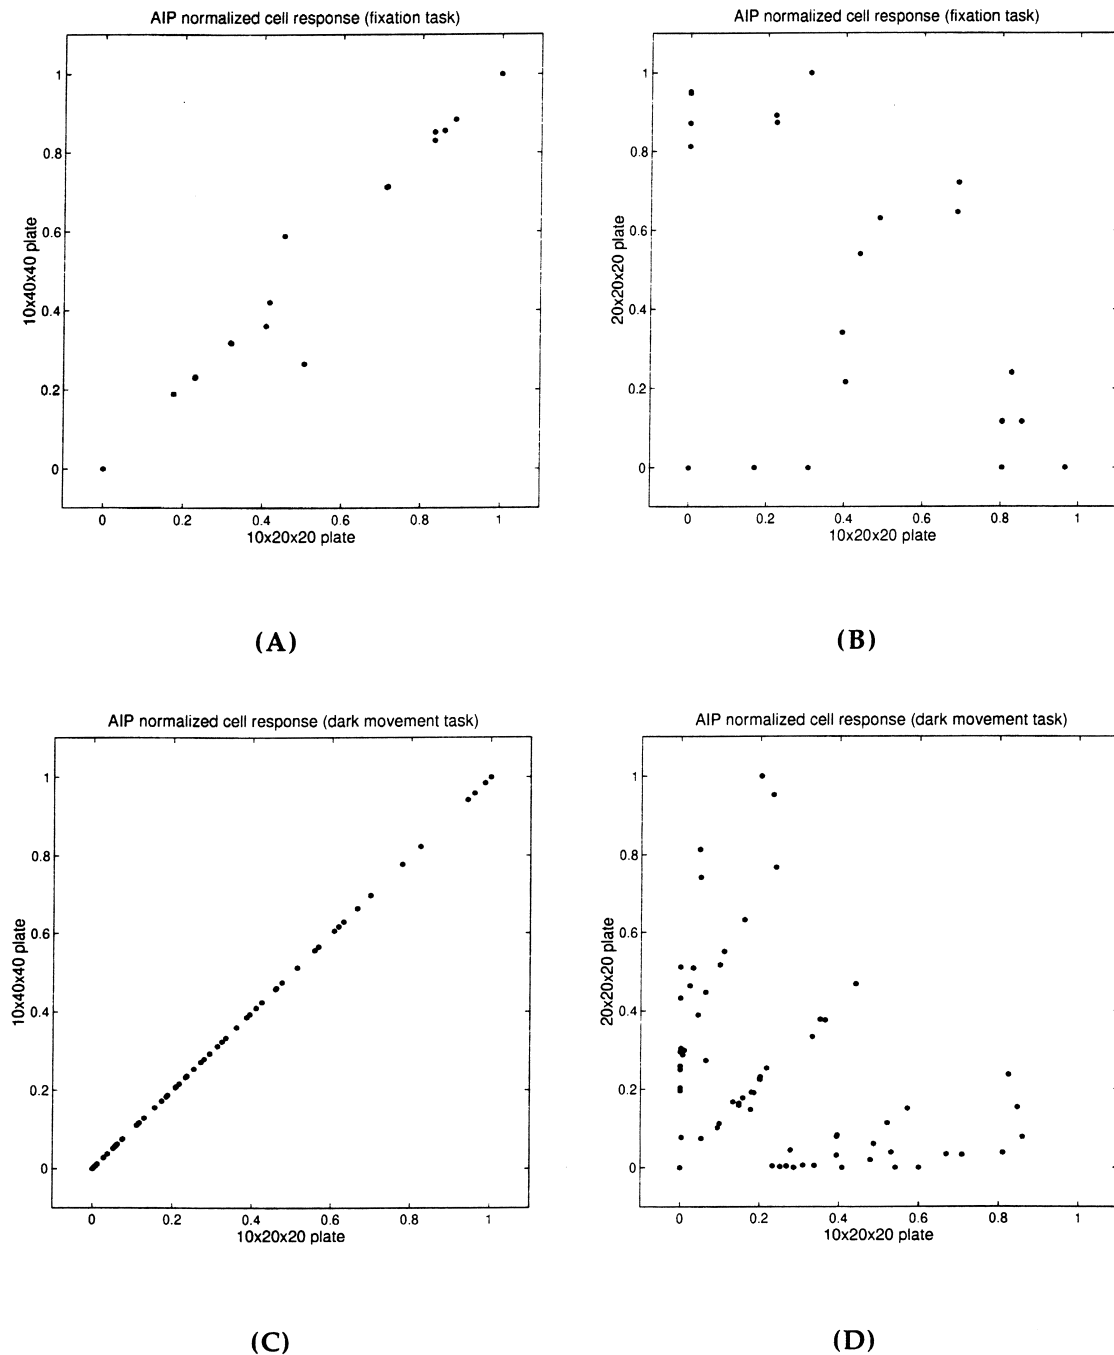


Fig. 22. Comparison of AIP visual responses for objects of the same (A) and different (B) widths; and AIP motor-related responses (dark movement condition) for objects of same (C) and different (D) widths.

integration of multiple sources of information. The visual system provides the object's identity, as well as an estimate of its shape, size, and localization in space. Although this information can be used to specify the fine details of a grasp plan, a single object actually affords many different grasps. The FARS model is a neural-level model of the cortical processes that are involved in the generation and execution of grasp plans. The model makes use of "cooperative computation" amongst several brain regions, including AIP, F5, and a variety of supporting areas. AIP's computa-

tion of visual affordances is complemented by corollary discharge from F5, which in the model is hypothesized to resolve multiple potential grasps by information on task constraints (from F6), working memory (from area 46) and instruction stimuli (from F2). As a result, AIP can function as an active memory in which a set of affordances, some of which are initially activated by visual input, are updated as the plan of action unfolds. We also suggest that IT can give a broadly tuned code for affordances when object knowledge is appropriate, and that when visual

input to AIP is available AIP can refine or replace IT input with a more finely tuned code for affordances. We argue, too, that action generally involves plans that take advantage of, but are not wholly driven by, actions currently afforded by visual stimuli, noting the role of pre-SMA and BG in sequence management.

The model addresses the issues of neural-level encoding of object, grasp, and phasic information, and shows how populations of cells can capture the critical information that is required to perform the necessary computations.

5.1. Hypotheses

In addition to reproducing single-cell behavior that has been observed experimentally in F5 and AIP in monkey, the model makes several more general hypotheses about the computations that are being performed by these neural regions.

1. The neural activity patterns observed in AIP reflect the extraction of affordances (action-relevant features) from the visual representation of the attended object.
2. During execution of the grasp, AIP maintains an active memory of the affordance that corresponds to the grasp that is being executed. This active memory is maintained in part by recurrent projections from F5.
3. During grasp plan formation, F5 selects from the possible grasps, as represented in working memory or in AIP (corresponding to the affordances). This selection process takes into account task-related constraints (received from areas 46, F2, or F6).
4. F5 is responsible for the high-level execution and monitoring of the grasp, based in part on haptic object sensing by SII.
5. Parameters of an object or grasp are reflected in subpopulations of F5 and AIP cells.

5.2. Predictions

These hypotheses lead to a number of predictions about the single cell- and population-level behaviors under several novel tasks.

1. In the case where two distinct objects are presented and subsequently grasped in an identical manner, hypotheses (1) and (2) imply that motor-related cells in AIP will tend not to reflect differences in the two situations. However, visual-related cells (which derive many of their inputs from PIP and IT) can reflect these differences, and hence appear to demonstrate an object-specific code.
2. Hypothesis (1) implies that different parameters of an object will be inequitably reflected in AIP. Specifically, we expect to see primarily those parameters that are relevant for programming of the grasp. Object parameters not used in the grasp programming process will not significantly affect AIP cell activity.

3. In the case where an object is to be grasped in one of two ways, but information about which of the two grasps that must be selected (e.g. via an instruction stimulus) is withheld for some delay period, hypothesis (3) implies that we would expect both possible grasps to be prepared initially, and hence we predict activity patterns in F5 that correspond to both possibilities. When the IS is given, the population of cells not corresponding to the selected grasp is inhibited. Hypothesis (2) implies that AIP motor-related cells will behave in a similar manner.
4. Because F5 monitors the ongoing grasp (hypothesis (4)), when there is a miscalibration between the visually estimated and actual size of an object, we predict a shift in the executing program at the time that the more precise tactile information becomes available. Hypothesis (2) implies that this shift in program will also be reflected in the motor-related cells of AIP.

Acknowledgements

Funding for this work has been provided in part by a grant from the Human Brain Project (with funding from NIMH, NASA, and NIDA) under a P20 Program Project Grant HBP: 5-P20-52194 for work on “Neural Plasticity: Data and Computational Structures” (M. A. Arbib, Director), by the National Science Foundation under Grant No. IRI-9221582, and by a fellowship to A. H. Fagg from the University of Southern California, School of Engineering, and the Computer Science Department.

Appendix A

In the implementation of the model, computing units are constructed from collections of simulated neurons, which are referred to as p-units, short from primable units. All neural regions within the model are constructed from a set of these p-units. A single p-unit is made up of a set of four neurons (Fig. 23). The primary function of the p-unit is performed by the signal unit. This unit fires in response to the presentation of a specific set of trigger inputs. Trigger inputs can be, for example, go signals (as is the case for F5) or some combination of object features (as in features represented in AIP). Activation of the signal unit can cause the output unit to fire, which results in the generation of an output from the p-unit. When this happens, we refer to the p-unit as being active. The p-unit output, depending upon the brain region, is connected either to other p-units (in the same or other brain regions) or to actuators.

The simple feature detection that is performed by the signal unit is subject to two levels of constraint, without which the p-unit cannot become active. The first constraint is that the p-unit must be primed. Priming signals (detected by the priming unit) are typically received from high-level

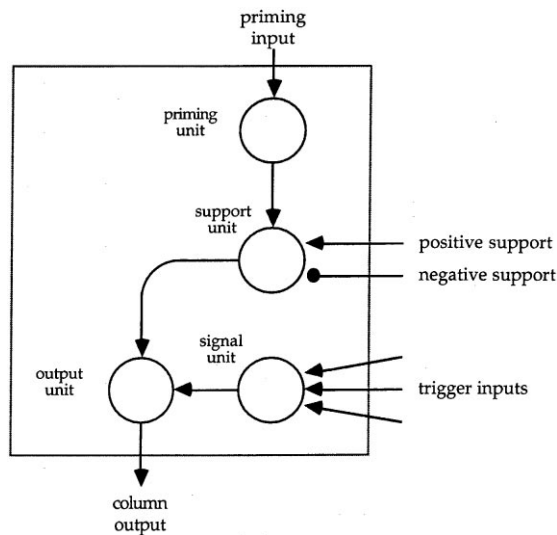


Fig. 23. The p-unit architecture.

regions. For example, when F6 begins to prepare a grasp program, it will prime all the p-units within F5. In this architecture, priming signals are considered to be non-specific; the priming of F5 by F6 prepares the system to perform a grasp, but does not necessarily specify the details of the grasp.

The second level of constraint, received through the positive/negative support inputs, is used in two different ways. First, in order to enforce the coactivation of p-units that participate in the same motor program (e.g. a power grasp), positive connections are established from each p-unit output to the support input of the other p-units. Likewise, when two p-units are involved in different

motor programs (e.g. a power and a precision grasp), mutually inhibitory connections are established. Second, when a higher-level region (e.g. area 46) selects a very specific motor program within another region (a grasp in F5 as a function of memory), it sends positive signals to the support input of the p-units participating in the motor program. The constraints are considered to be satisfied only when the p-unit is both primed and supported. When this occurs, the p-unit is prepared to become active when the correct trigger input is received.

The output unit is implemented as a leaky-integrator neuron (Arbib, 1989). Depending upon the brain region, the time constant of response ranges between 150 and 250 ms. The remaining units are simply implemented as linear-threshold neurons, with no temporal dynamics. Figures in this paper that show “cell” activity, plot the firing rate of the output unit. Further information regarding the implementation of the p-units (but there referred to, somewhat misleadingly, as “columns”) may be found in Fagg (1996).

Appendix A.1 Details of the Wiring Process

In this section, we describe in detail the wiring process that takes place between the subpopulations of cells in each region. Table 3 lists the set of rules that are used to decide when it is appropriate for one F5 cell to send a projection to another cell, the sign of that projection, and the probability of establishing the connection if it is appropriate. The rules reference the behavioral descriptors described earlier for F5 and AIP. The behavioral descriptors for the SII consist of grasp type, and phase (Extension, Flexion, and Release). The descriptor for the BG is the movement phase, and the

Table 3

Connections from F5 cells to other populations. Specifies the types of F5 cell sending connections (column 1), the destination of the connections (column 2), and the probability of connection for a pair of matching cells (column 4). The abbreviations are as follows: p = phase (one of Set, Extension, Flexion, Hold, and Release), g = grasp (precision or lateral pinch), a = aperture (either NONE, or a specific aperture, X), * = match anything, and unit # U indicates a specific unit in a population

Region and cell type	Projection to region and cell type	Notes	Prob
F5			
$g = G$	- F5: $g\pi G$	Inhibit other grasp cells	1
$p, g, a = \text{NONE}$	+ F5: $p, g, a = \text{NONE}$	Support cells of same population (Phase, Grasp, and no Aperture)	0.5
$p, g, a = X$	+ F5: $p, g, a = X \pm \Delta$	Support cells of same population (Phase, Grasp, and similar Aperture)	0.5
	- F5: $p, g, a \neq X \pm \Delta$	Inhibit cells of same different aperture populations	1
	+ F5: $p, g, a = \text{NONE}$	Aperture F5 cells support non-aperture cells	0.5
$p = (F,H,R), g, a$	+ AIP: $p = *, g, a, o = M$	F5 cells that are active during the Flexion, Hold, and Release phases, connect to motor-oriented AIP cells that match in grasp and aperture classification	0.2
$p = S, g, a$	+ AIP: $p = S, g, a, o = M$	F5 Set cells connect to motor-oriented, Set AIP cells that match	0.2
$p = E, g, a$	+ AIP: $p = (S,E), g, a, o = M$	F5 Extension cells connect to motor-oriented, Set and Extension AIP cells that match	0.2
$p, g, a = \text{NONE}$	+ F5: $p, g, a = X$	Non-aperture F5 cells support activation of aperture F5 cells	0.5
$p, g, a = X$	+ F1	Projects to those F1 cells that implement the necessary movement for this grasp and phase	0.5
$p = (E,F,R), g, a = X$	+ SII: $p, g, a = X$	Aperture F5 cells that are active for the Extension, Flexion, and Release phases project to specific SII cells	0.5
p	+ BG: p	Inform BG as to the current movement phase	1
unit # U	+ A46: unit # U	Each F5 cell projects to its twin in A46	1

Table 4

Connections from one class of cells to other populations. Specifies the types of cell sending connections (column 1), the destination of the connections (column 2), and the probability of connection for a pair of matching cells (column 4)

Region and cell type	Projection to region and cell type	Notes	Prob
AIP			
$g, a = \text{NONE}$	+ AIP: $g, a = \text{NONE}$	Non-aperture AIP cells of the same class exchange positive connections	0.1
$g, a = X$	+ AIP: $g, a = X$	AIP non-aperture cells project to aperture cells	0.1
	+ AIP: $g, a = \text{NONE}$	Aperture cells project to non-aperture cells	0.1
	+ AIP: $g, a = X \pm \Delta$	AIP cells of the same class (grasp type, and similar aperture) exchange positive connections	0.1
$g = G$	– AIP: $g, a \neq X \pm \Delta$	Non-matching aperture cells inhibit one-another	1
	– AIP: $g\pi G$	Cells that do not represent the same grasp inhibit one-another	1
g, a	+ F5: g, a	AIP cells project to F5 cells that match in grasp and aperture	0.25
PIP/IT			
	+ AIP: $g, a = (X \pm \Delta, \text{NONE}), o = V$	PIP cells project to visual-oriented AIP cells for which the grasp and aperture are appropriate for the object	0.5
SII			
p, g, a	+ F5: $p + 1, g, a$	SII cells detect the end of some movement phases. These cells trigger the activation of F5 cells for the next phase of movement	0.7
$p = P, g = G, a = A$	– SII: $p\pi P$, OR $g\pi G$, OR $a\pi A$	Inhibit other SII cells that encode other hand configurations	1
F6			
$p = S$	+ F5: $p = S$	F6 triggers initiation of the Set phase of movement (detection of READY signal)	1
$p = F$	+ F5: $p = F$	F6 triggers initiation of the Flexion phase of movement (detection of GO signal)	1
$p = R$	+ F5: $p = R$	F6 triggers initiation of the Release phase of movement (detection of 2nd GO signal)	1
g	+ F5: $p = S, g$	F6 biases the selection of the grasp by exciting specific grasp Set cells	1

descriptors used for F6 and F2 are the grasp. F6 also contains cells that represent the external triggering signals given by the experimenter.

The first column of Tables 3–5 describe the subpopulations from which a set of connections may originate; the second column describes the appropriate destinations. In the first row of Table 3, the subpopulations are described as a function of grasp only. A cell that represents a particular grasp sends a negative connection to other F5 cells that do

not encode that particular grasp. The second row states that the subpopulations are defined as a function of phase and grasp, and must consist of cells that do not encode aperture. These cells project positively to others of the same type with a probability of 0.5. The third row states that cells encoding a particular aperture, phase, and grasp project positive connections to cells of the same grasp and phase, as long as the apertures are similar.

Tables 4 and 5 describe the connection rules for the

Table 5

Connections from cells in area 46, F2, BG, and SI to other populations. Specifies the types of cell sending connections (column 1), the destination of the connections (column 2), and the probability of connection for a pair of matching cells (column 4)

A46			
unit #U	+ F5: unit #U	Each A46 cell projects to its twin in F5	1
F2			
g	+ F5: $p = S, g$	F2 biases the selection of the grasp by exciting specific grasp Set cells	1
BG			
p	+ F5: $p + 1$	Prepare F5 cells for the next phase	1
	– F5: $p - 1$	Inhibit last F5 phase	1
SI			
g configuration, $p = (E, F, R)$	+ SII: g, p	The representation in SI of the state of the hand at the end of the Extension, Flexion, and Release phases of movement is associated with the corresponding SII unit	1
—	+ F1	F1 cells receive either position- or force-related inputs from SI	1

remaining regions represented in the FARS model. Note that area 46 and F5 exchange excitatory connections between pairs of cells. It is through these connections that a memory is first latched from F5 into area 46 during execution of the grasp in the light, and then restored to area F5 during grasp execution in the dark.

For the connections from PIP to AIP, the mapping from object to grasp (including aperture) is specified a priori. For each object \rightarrow grasp mapping, the PIP representation of the object is associated with a 50% probability with visual-oriented AIP cells that match the grasp and aperture. Association means that cells that would normally be active in PIP to represent the object provide excitatory connections to AIP cells. Note that the AIP cells that either encode no aperture or the aperture appropriate for the object are eligible to receive PIP inputs.

References

- Anderson, J. A., Rossen, M. L., Viscuso, S. R. & Sereno, M. E. (1991). Experiments with representation in neural networks: object motion, speech, and arithmetic. In H. Haken & M. Stadler (Eds.), *Synergetics of Cognition* (pp. 54–69). Berlin: Springer-Verlag.
- Arbib, M. A. (1972). *The Metaphorical Brain: An Introduction to Cybernetics as Artificial Intelligence and Brain Theory*. Wiley Interscience.
- Arbib, M. A. (1989). *The Metaphorical Brain 2: Neural Networks and Beyond*. New York: Wiley-Interscience.
- Arbib, M. A., Iberall, A. R. & Lyons, D. (1985). Coordinated control Programs for Control of the Hands. In A. W. Goodwin & I. Darian-Smith (Eds.), *Hand Function and the Neocortex. Experimental Brain Research Supplemental*, 10, 111–129.
- Arbib, M. A., Fagg, A. H. & Grafton, S. T. (1998). Synthetic PET imaging for grasping: from primate neurophysiology to human behavior (in preparation).
- Baizer J. S., Ungerleider L. G., & Desimone R. (1991). Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *Journal of Neuroscience*, 11 (1), 168–190.
- Boecker H., Dagher A., Ceballos-Baumann A. O., Passingham R. R., Samuel M., Friston K. J., Poline J. B., Dettmers C., & Conrad B. (1998). Role of the human rostral supplementary motor area and the basal ganglia in motor sequence control—investigations with H-2 O-15 PET. *Journal of Neurophysiology*, 79 (2), 1070–1080.
- Castiello U., & Jeannerod M. (1991). Measuring time to awareness. *Neuroreport*, 2 (12), 787–800.
- Colby C. L., Duhamel J. R., & Goldberg M. E. (1993). Ventral intraparietal area of the macaque—anatomic location and visual response properties. *Journal of Neurophysiology*, 69 (2), 902–914.
- Decety J., Perani D., Jeannerod M., Bettinardi V., Tadary B., Woods R., Mazziotta J. C., & Fazio F. (1994). Mapping motor representations with positron emission tomography. *Nature*, 371 (6498), 600–602.
- di Pellegrino G., Fadiga L., Fogassi L., Gallese V., & Rizzolatti G. (1992). Understanding motor events: a neurophysiological study. *Experimental Brain Research*, 91 (1), 176–180.
- Dominey P. F. (1995). Complex sensory-motor sequence learning based on recurrent state representation and reinforcement learning. *Biological Cybernetics*, 73 (3), 265–274.
- Dominey P. F., & Arbib M. A. (1992). A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex*, 2, 153–175.
- Evarts, E. V. & Wise, S. P. (1984). Basal Ganglia outputs and motor control. In *CIBA Foundation Symposium*, 107 (pp. 83–102).
- Fagg, A. H. (1996). A computational model of the cortical mechanisms involved in primate grasping. Ph.D. Dissertation, University of Southern California, Computer Science Department. Available at: http://www-hbp.usc.edu/HBP/models/conformant/FARS_1.2/Thesis/fagg_thesis_TOC.html
- Fagg A. H., & Arbib M. A. (1992). A model of primate visual-motor conditional learning. *Journal of Adaptive Behavior*, 1 (1), 3–37.
- Funahashi S., Bruce C. J., & Goldman-Rakic P. S. (1993). Saccade-related activity in the primate prefrontal cortex and its role in working memory processes. *Biomedical Research*, 14 (suppl. 1), 85–88.
- Gentilucci M., Fogassi L., Luppino G., Matelli M., Camarda R., & Rizzolatti G. (1988). Functional organization of inferior area 6 in the macaque monkey I. Somatotopy and the control of proximal movements. *Experimental Brain Research*, 71, 475–490.
- Gentilucci M., Daprati E., Toni I., Chieffi S., & Saetti M. C. (1995). Unconscious updating of grasp motor program. *Experimental Brain Research*, 105 (2), 291–303.
- Georgiou N., Bradshaw J. L., Phillips J. G., Chiu E., & Bradshaw J. A. (1995). Reliance on advanced information and movement sequencing in Huntington's-disease. *Movement Disorders*, 10 (4), 472–481.
- Gibson J. J. (1955). The optical expansion-pattern in aerial location. *American Journal of Psychology*, 68, 480–484.
- Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Allen and Unwin.
- Goodale M. A., & Milner A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience*, 15 (1), 20–25.
- Grafton, S. T., Fagg, A. H., Woods, R. P. & Arbib, M. A. (1996a). Functional anatomy of pointing and grasping in humans. *Cerebral Cortex*, 6(2), 226–237.
- Grafton, S. T., Arbib, M. A., Fadiga, L. & Rizzolatti, G. (1996b). 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., & Arbib M. A. (1998). Dorsal premotor cortex and conditional movement selection: a PET functional mapping study. *Journal of Neurophysiology*, 79, 1092–1097.
- Hoff B., & Arbib M. A. (1993). Simulation of interaction of hand transport and preshape during visually guided reaching to perturbed targets. *Journal of Motor Behavior*, 25, 175–192.
- Ingle, D. J., Schneider, G. E., Trevarthen, C. B. & Held, R. (1967). Locating and identifying: two modes of visual processing (a symposium). *Psychologische Forschung*, 31(1 and 4).
- Jeannerod M., Arbib M. A., Rizzolatti G., & Sakata H. (1995). Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends in Neuroscience*, 18 (7), 314–320.
- Kermadi, I., Jurquet, Y., Arzi, M., & Joseph, J. P. (1993). Neural Activity in the Caudate-Nucleus of Monkeys During Spatial Sequencing. *Experimental Brain Research*, 94 (2), 352–356.
- Kimura M. (1993). The role of putamen and caudate-nucleus systems in the voluntary movement. *Biomedical Research*, 14 (1), 37–41.
- Kurata K., & Hoffman D. S. (1994). Differential effects of muscimol microinjection into dorsal and ventral aspects of the premotor cortex of monkeys. *Journal of Neurophysiology*, 71 (3), 1151–1164.
- Logothetis N. K., & Pauls J. (1995). Psychophysical and physiological evidence for viewer-centered object representations in the primate. *Cerebral Cortex*, 5 (3), 270–288.
- Logothetis N. K., Pauls J., & Poggio T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, 5 (5), 552–563.
- Luppino G., Matelli M., & Rizzolatti G. (1990). Cortico-cortical connections of two electrophysiologically identified arm representations in the mesial agranular frontal cortex. *Experimental Brain Research*, 82 (1), 214–218.
- Luppino G., Matelli M., Camarda R. M., Gallese V., & Rizzolatti G. (1991). Multiple representations of body movements in mesial area 6 and the adjacent cingulate cortex: an intracortical microstimulation study in the macaque monkey. *Journal of Comparative Neurology*, 311, 463–482.

- Luppino G., Matelli M., Camarda R., & Rizzolatti G. (1993). Corticocortical connections of area F3 (SMA-proper) and area F6 (pre-SMA) in the macaque monkey. *Journal of Comparative Neurology*, 338, 114–140.
- MacKenzie, C. L. & Iberall, T. (1994). *The Grasping Hand. Advances in Psychology*, 104. North Holland.
- Matelli, M., Camarada, R., Glickstein, M., & Rizzolatti, G. (1986). Afferent and efferent projections of the inferior area 6 in the macaque monkey. *Journal of Comparative Neurology*, 251 (3), 281–298.
- Mitz A. R., Godshalk M., & Wise S. P. (1991). Learning-dependent neuronal activity in the premotor cortex. *Journal of Neuroscience*, 11 (6), 1855–1872.
- Muakkassa, K. F., & Strick, P. L. (1979). Frontal Lobe Inputs to Primate Motor Cortex. Evidence for Four Somatotopically Organized “Premotor” Areas. *Brain Research*, 177, 176–182.
- Murata, A., Gallese, V., Kaseda, M., Kunitomo, S. & Sakata, H. (1993). Hand-manipulation-related neurons of the parietal cortex of the monkey: further analysis of selectivity in shape, size, and orientation of objects for manipulation. *Proceedings of the Society of Neuroscience Annual Meeting*, Washington, DC.
- Paulignan Y., Jeannerod M., MacKenzie C., & Marteniuk R. (1991). Selective perturbation of visual input during prehension movements. 2. The effects of changing object size. *Experimental Brain Research*, 87, 407–420.
- Petit L., Orssaud C., Tzourio N., Crivello F., Berthoz A., & Mazoyer B. (1996). Functional-anatomy of a pre-learned sequence of horizontal saccades in humans. *Journal of Neuroscience*, 16 (11), 3714–3726.
- Preuss T. M., & Goldman-Rakic P. S. (1989). Connections of the ventral granular frontal cortex of macaques with perisylvian premotor and somatosensory areas: anatomical evidence for somatic representation in primate frontal association cortex. *Journal of Comparative Neurology*, 282 (2), 293–316.
- Quintana J., & Fuster J. M. (1993). Spatial and temporal factors in the role of prefrontal and parietal cortex in visuomotor integration. *Cerebral Cortex*, 3 (2), 122–132.
- Rizzolatti, G. (1987). Functional organization of inferior area 6. In *Motor Areas of the Cerebral Cortex (CIBA Foundation Symposium 132)* (pp. 171–186). Chichester: John Wiley and Sons.
- 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.
- Roland P. E., Larsen B., Lassen N. A., & Skinhøj E. (1980). Supplementary motor area and other cortical areas in organization of voluntary movements in Man. *Journal of Neurophysiology*, 43, 118–136.
- Sakata H., & Kusunoki M. (1992). Organization of space perception: neural representation of three-dimensional space in the posterior parietal cortex. *Current Opinions in Neurobiology*, 2 (2), 170–174.
- Sakata H., Taira M., Kusunoki M., Murata A., & Tanaka Y. (1997). The TINS lecture—the parietal association cortex in depth perception and visual control of hand action. *Trends in Neurosciences*, 20 (8), 350–357.
- 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.
- Tanji J., & Shima K. (1994). Role for supplementary motor area cells in planning several movements ahead. *Nature*, 371 (6496), 413–416.
- Ungerleider, L. G. & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale & R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior*. The MIT Press.