Chapter 6: Model Simulation Results

In this chapter, we demonstrate the capabilities of the model that has been developed in the previous two chapters. First, we examine the key computational mechanisms in the model, focusing on the interaction between AIP and F5 as a grasp is selected and then executed. Second, we draw comparisons between single cell behavior observed in several monkey experiments and similar experiments performed with the model. Finally, a set of novel tasks is presented to the model to explore several important computational questions. Observations made at the single unit and population levels (in the model) serve as predictions for future monkey experiments.

Introduction

In the experiments that follow, the protocol is derived from that used by Sakata (Taira, et al., 1990). We first illustrate the general behavior of the model using this protocol, and then present a series of new questions that are designed to examine more subtle issues of object and grasp encoding at the neural level. These include:

1. *Grasping of two different objects with an identical grasp*. How are the objects and grasp represented in AIP and F5?
2. *Grasping of a single object in one of two ways (as determined by an instruction stimulus).* How are these two different grasps represented, and how is one ultimately selected for execution?

3. *Grasping of a set of objects of varying size*. How do F5 and AIP represent the size of grasp aperture?
4. *Visual/motor perturbation*. How does the system respond when there is a mismatch between the assessed parameters of the object (by the visual system) and the actual parameters?
5. *Object parameter representation in AIP*. Are all parameters of an object represented equally in AIP?

Unless otherwise stated, all experimental results presented in this Chapter are derived from the model.

Overview of Model Behavior

Figure 6.1 demonstrates the cascade of activity that flows though a set of simulated F5 cells as a precision grasp is executed. The trial starts with the presentation of a narrow cylinder to be grasped (t=0, event not shown). 700ms later, a *ready signal* is presented, indicating that the movement trial is about to begin. The system responds by preparing to execute a precision grasp (as encoded by a population of F5 set cells, of which one is shown in the figure). This preparation process provides positive support to F5 units that will be involved in the extension phase of movement (the support is provided indirectly through the BG). This state is maintained until the *Go signal* is given, at which point the F5 extension units begin to fire.

The activation of the extension-related units results in the following:

1. Via the basal ganglia, the set-related units are shunted (through negative support).
2. Also through the BG, the next phase of movement is prepared by providing positive support to those units that will be involved in the flexion phase.
3. The SII units that are responsible for detecting the maximum aperture for the current grasp are primed.

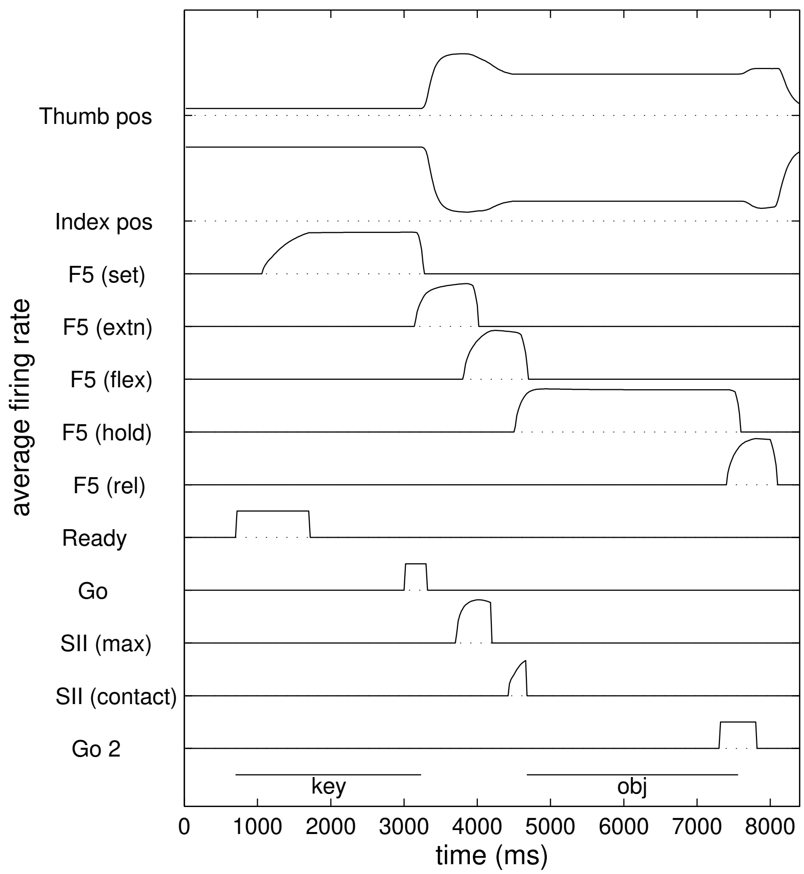

Figure 6.1: F5 activity during execution of a precision grasp. The top two traces show the position of the thumb and index finger, respectively. The next five traces represent the average firing rate of five F5 neurons (set-, extension-, flexion-, hold-, and release-related). The remaining five traces represent the various external (Ready, Go, Go2) and internal (SII) triggering signals.

4. The internal model of current hand state in area 7 is activated.

5. The index finger and thumb begin to open (as shown by the top two traces).

When the maximum aperture is achieved, as detected by either SII (labeled *SII max* in the figure) or by the internal model (not shown), a second phase transition is made. This phase transition is driven by the activation of the flexion-related units. The flexion units shunt the previous phase (extension), prepare for the next (hold) phase, prepare for detection of the contact event (by priming *SII contact*), and flex the thumb and index finger.

The contact event (establishment of the precision grasp in this case) forces a phase transition to the hold phase of the program. The network remains in this state until the secondary go stimulus is provided by the experimenter, at which point the release phase of the motor program is initiated, and the fingers are extended enough such that contact with the object is lost.

For comparison purposes, we also make use of the Sakata *key* and *object* phase notation (see figure). Here, the key phase is defined from the point at which the ready signal is received to the initiation of movement in response to the go signal. The object phase is defined by the point of contact with the object and the point at which the release movement is initiated.

Figure 6.2 demonstrates the relative time course of two AIP cells relative to F5 activity. The trial and displayed F5 unit traces are the same as in the previous figure. The two AIP cells activate after the F5 set unit becomes active, but both begin to fire within the key phase of the task. Once they begin to fire, note that they continue to do so until the task is completed. Note also that both AIP units show peak activity during the extension/flexion phases of the task, as is observed by (Taira, et al., 1990). We explain this property as follows:

Suppose that for a given F5/AIP pair of units that participate together in a grasping program, there is some constant probability that a projection will exist from the F5 unit to the AIP unit. It has been observed by (Rizzolatti, et al., 1988) that the number of F5 neurons which fire during the extension and flexion phases of the program is greater than the number that fire during either the set or hold phases of the program. Hence, on average, a single AIP unit will receive a larger number of recurrent projections from extension/flexion F5 units than any other type. As a result, its level of activity will usually be greater during these phases of the program.

Comparison to Observed Biological Data

F5 Phasic Responses

In the previous two figures (7.1, 7.2), the activation of each of the F5 cells was restricted to exactly one of the defined phases of movement. However, in Chapter 2 we observed that it is often the case that a single F5 cell in monkey is active over more than one of these phases. In the model, this property is achieved by connecting the basal ganglia phasing loop to the F5 unit in such a way that it is prepared to activate during one phase but then is not immediately shunted on the following phase. Rather, the circuitry allows for an arbitrary number of intervening phases to exist between these two events (see Figure 5.12 and accompanying discussion for more details). In addition, an F5 unit may exchange positive supporting connections with other F5 units that are active over disparate phases. The result is that from one phase to another, the F5 unit may receive different levels of positive support, causing a modulation of the overall activity of the column.

Both of these effects are demonstrated in Figure 6.3. F5 units A and B are set cells, although unit B continues to fire through the extension phase of movement. Units D and E are only active during the extension and flexion phases; the latter is qualitatively similar to the unit (recorded in monkey) that is shown in Figure 2.23. Unit G fires significantly during the enclose phase of the grasp, and then continues to fire at a moderate level through the hold phase of movement. This mimics the F5 cell behavior shown in Figure 2.24.

AIP Phasic Responses

Figure 6.4 demonstrates a range of AIP cell behaviors. Like the monkey cell of Figure 2.12, unit A becomes active almost immediately (before the ready signal is given). This feature indicates that unit A receives a significant degree of input from the visual areas. Units B and C only activate once the ready signal has been received. This is the case because these units rely

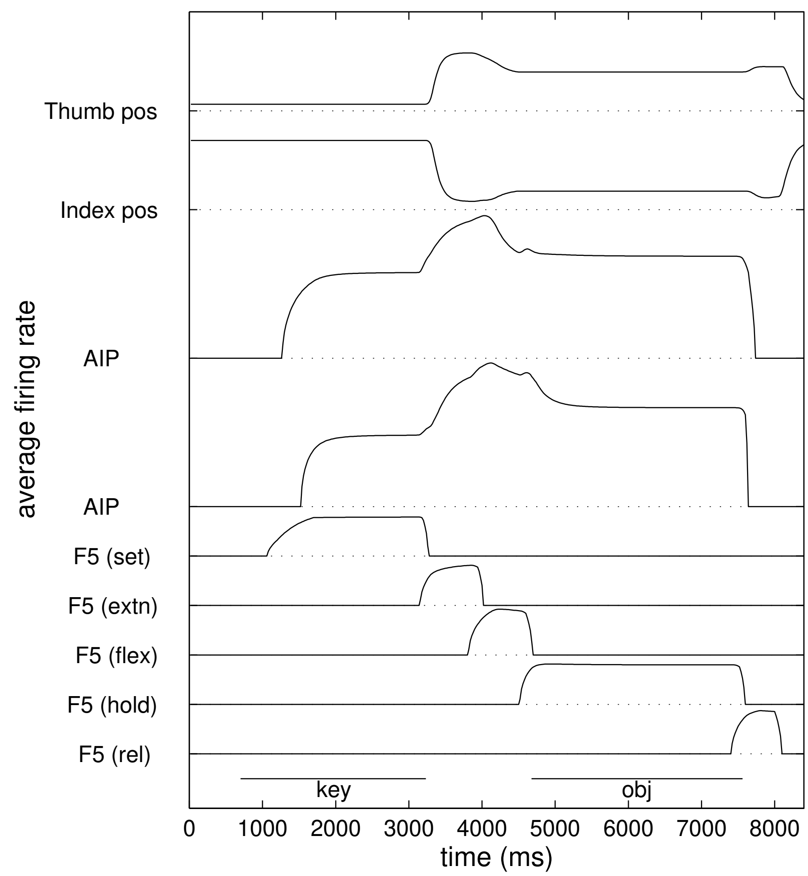

Figure 6.2: F5/AIP interaction during execution of the grasp.

on recurrent activity from the set cells in F5 in order to achieve a level above threshold. In regards to onset, these cells compare to the monkey cell of Figure 2.14. However, in this case, the activity ramps up continuously from onset until movement execution is complete. Units D and E do not receive any significant visual inputs, and thus rely completely on the recurrent F5 activity during the movement and the hold phases of the task (these compare in behavior to Figures 2.15 and 2.13, respectively).

AIP Visual/Motor Distinction

One classification that Sakata uses to describe AIP cell behaviors is that of visual- versus motor-related activity (Taira, et al., 1990). Neurons are tested in four different conditions: {grasping an object versus fixation of an object} x {performance in the light versus in the dark}.

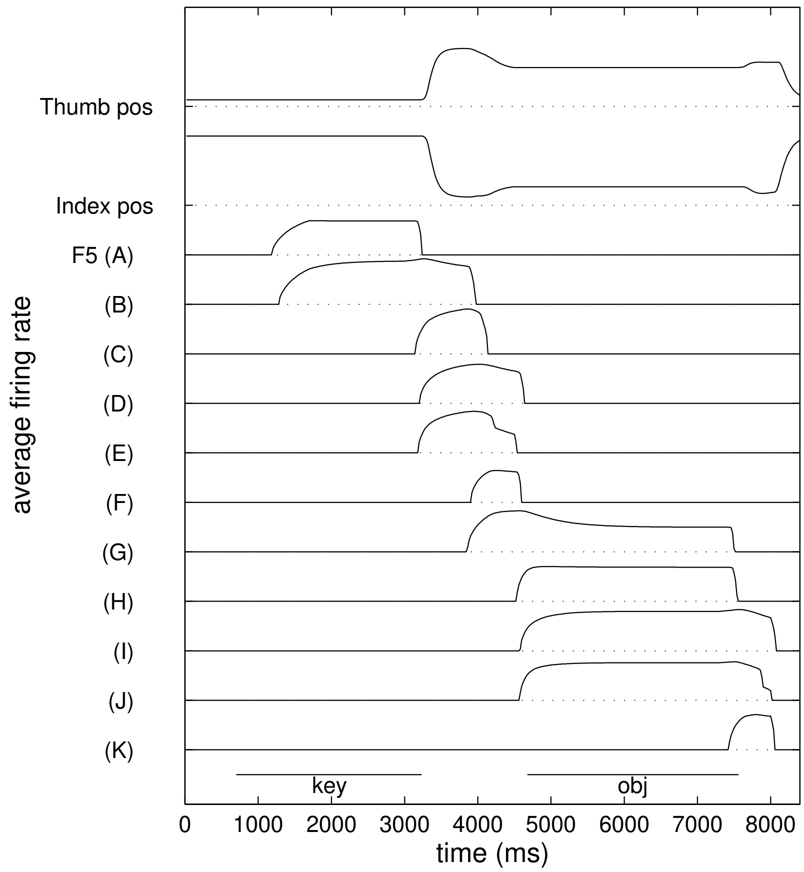

Figure 6.3: A set of F5 cells with a variety of temporal properties.

A neuron is considered to have a motor-related component when it is active during grasping, regardless of whether the task is performed in the light or dark. A neuron is visual-related if it is active for conditions in which the object is visible and is not active for the dark conditions.

In the model, we explain this distinction as a function of the relative strength of the inputs into the AIP columns from the visual areas as compared with F5 layer. Motor-related cells receive relatively large recurrent projections from F5; visual-related cells receive inputs primarily from PIP and IT.

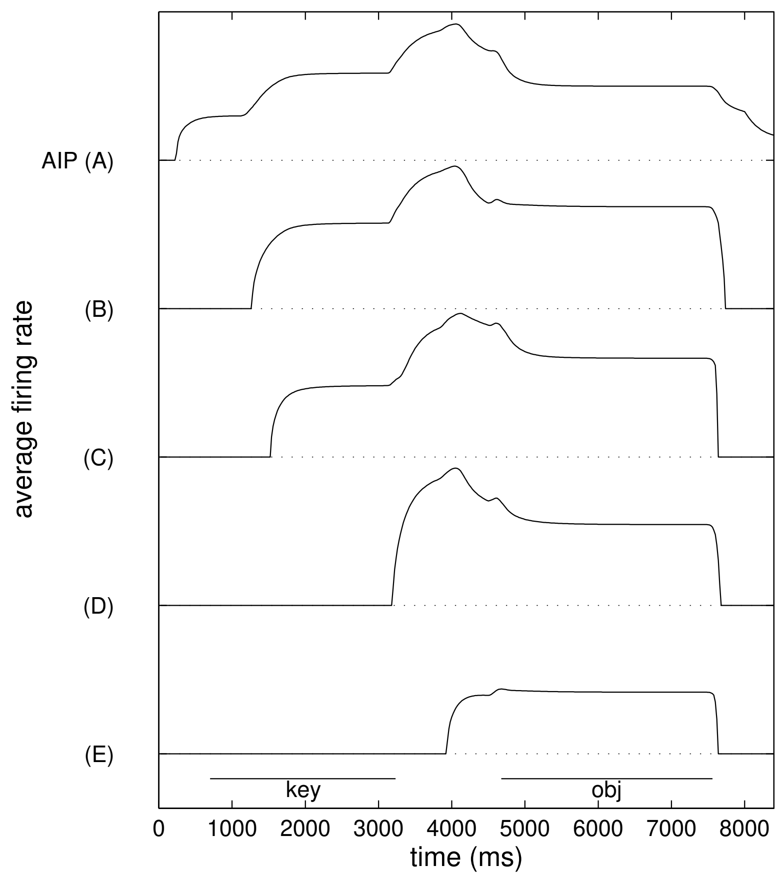

Figure 6.4: Variety of AIP cells.

Different experimental protocols are used to test the model in the three conditions of interest (here, we ignore the dark fixation task). For the movement in the light condition, the protocol is identical to that described above. In the fixation task, 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 is spent fixating on the object and no movements are generated. For the dark grasping task, the model first performs the task in the light (traces are not shown). During the performance of the task in the light, area 46 maintains a memory of those F5 cells that participate in the grasp. 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 positive support 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 is providing essentially a static description of the grasp that was recently executed. By *static* we mean that the temporal aspects of the grasp are not stored - only a memory of those units that were active at *some time* during the execution. How might this relate to the observation that area 46 is involved in human in the imagination of grasp execution (Decety, et al., 1994)? One possibility is that this region provides the static representation of the grasp to be executed. This information could then provide the context required by the mesial motor areas (F3 and F6) and F5 in order to simulate the grasp execution -- unfolding the static representation in time (Grafton, et al., 1996).

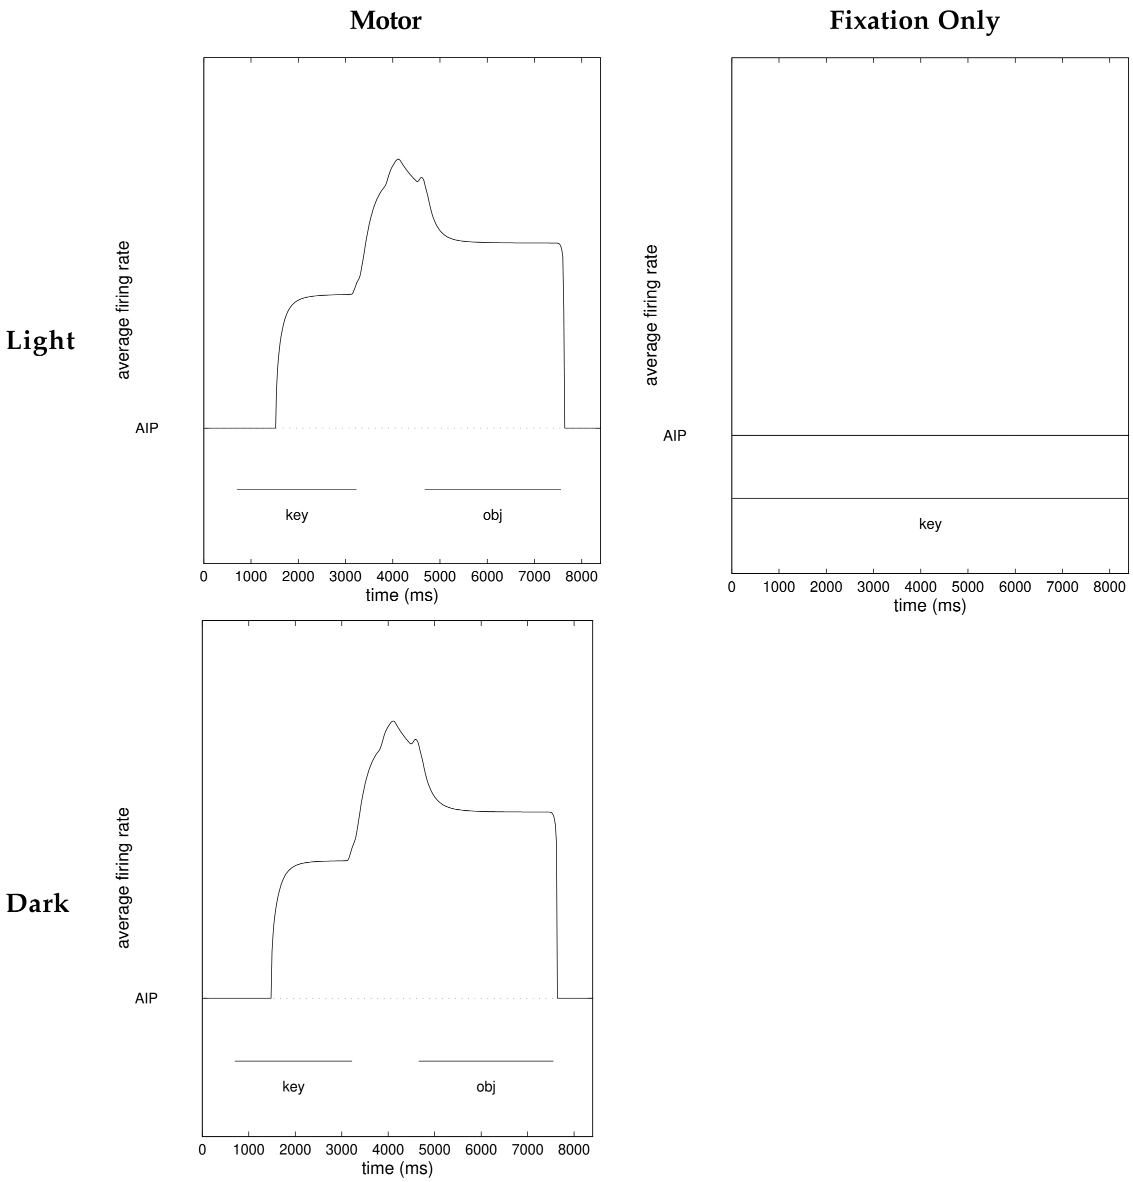

Figure 6.5: A pure motor-related AIP cell.

Figure 6.5 demonstrates one of the model's pure motor cells. During the fixation condition, the unit does not respond to the visible object. In addition, movement in either the light or the dark results in identical activity traces.

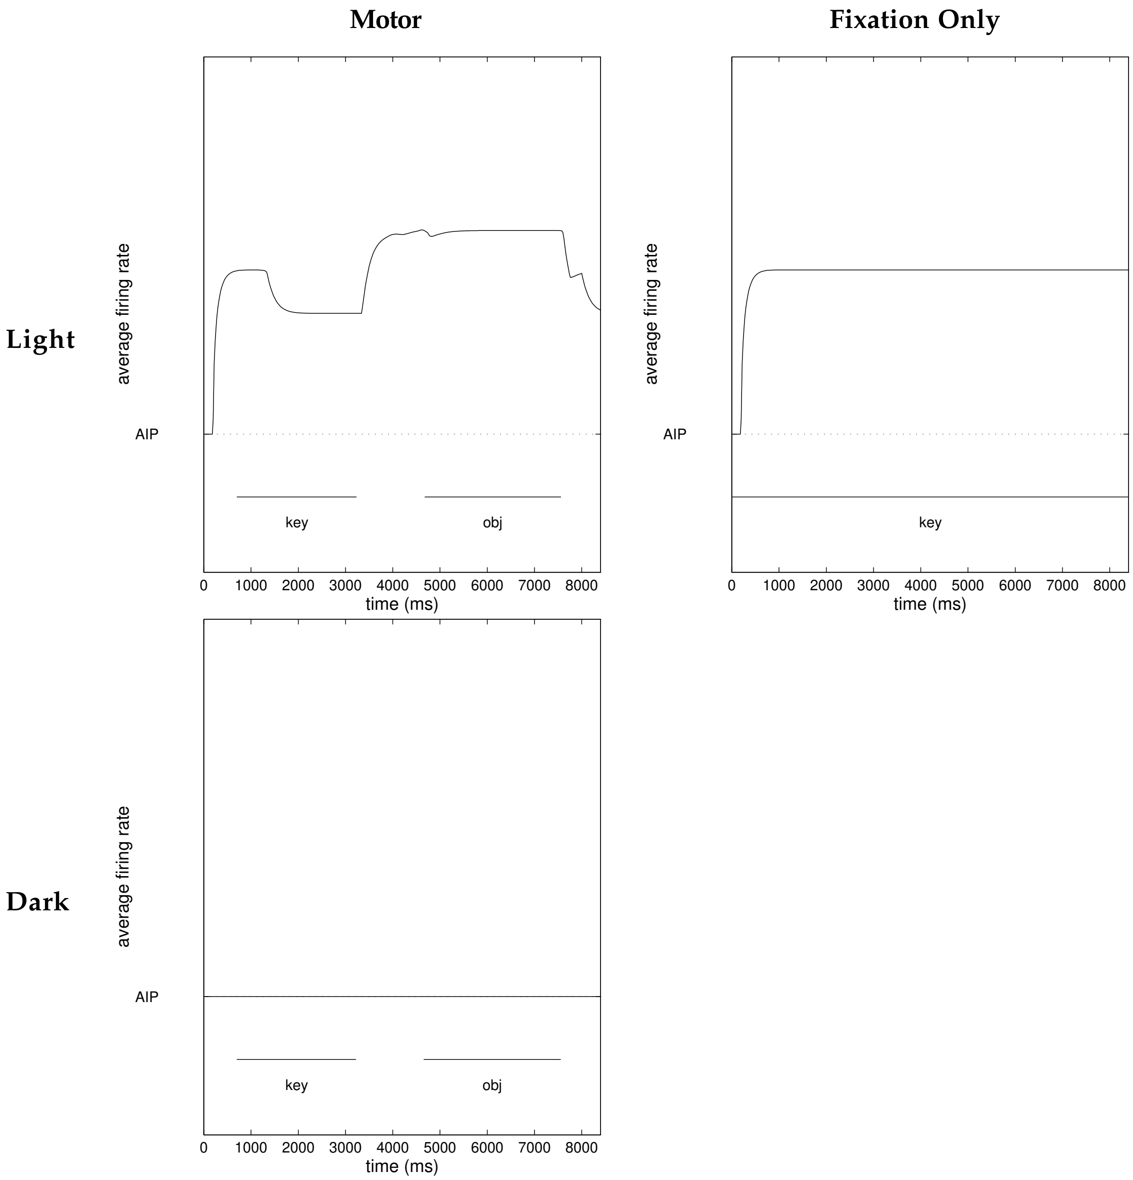

Figure 6.6: A pure visual AIP unit.

A pure visual-related cell is shown in Figure 6.6. Even though this unit does not receive direct projections from any F5 units, its activity level is modulated by the phase of movement. The initial increase in activity is due to the visual input from PIP. Shortly after the ready signal is received, a number of F5 set units begin to activate, which in turn activate motor-

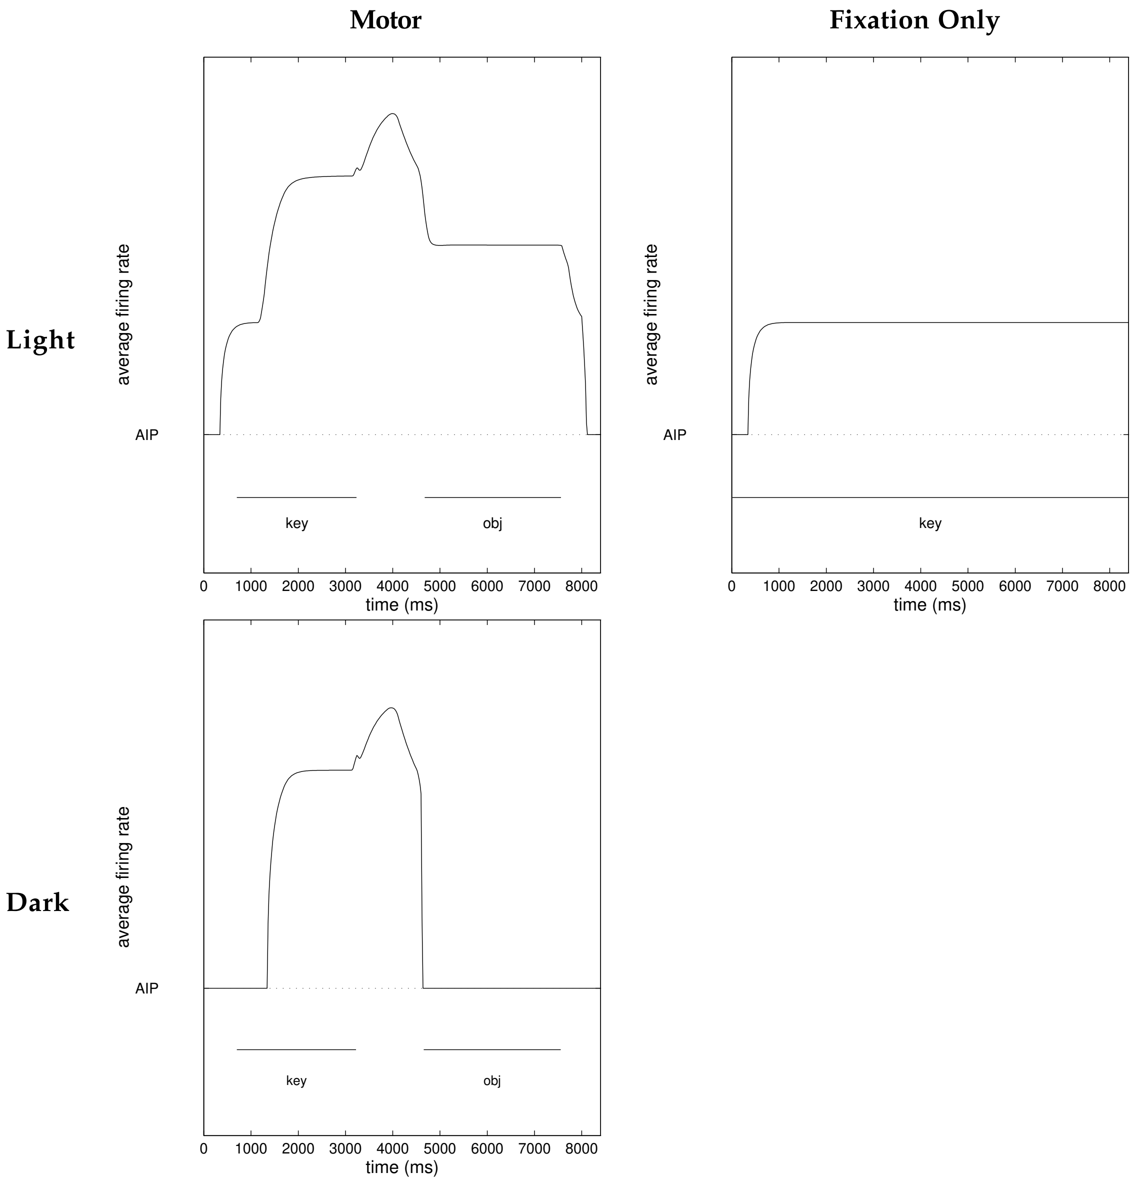

Figure 6.7: A visually-modulated AIP unit.

related units in AIP. This increase in overall activity in AIP results in an increase in the inhibitory signal through the basal ganglia loop, causing a reduction in the activity level of the unit shown in the figure.

The subsequent increase in activity after the key phase is also due to indirect mechanisms: surrounding motor-related units become highly active during movement, and thus provide a significant degree of positive support.

Finally, Figure 6.7 demonstrates the behavior of an AIP unit whose motor-related responses are modified by visual inputs. 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.

Population Analysis

Examining the behavior of individual cells allows for an understanding of how a cell contributes and is affected by the execution of a motor program. However, it is also desirable to look at how a set of cells work together in order to implement the program execution. The population analysis described here is designed to give a broad view of cell behavior across a large number of cells. The basic scheme is such: we measure the cells "participation" during the execution of a single trial (and possibly over a single phase of the trial). This measure is defined as the cell's average firing rate during the condition and phase of interest (note that much temporal information is lost in this averaging process). Given two (or more) such measures for different tasks or phases, we can compare the relative participation of the cell under these different conditions, and then examine how these relative measures vary over the population.

Pairs of measures can be plotted in a plane (see Figure 6.8) in order to construct a visual representation of these relationships. Panel A of Figure 6.8 demonstrates one such comparison between the normalized cell response during the set and extension phases of movement (during execution of a precision grasp). Each point in the plot corresponds to a single F5 module (the activity of the output unit of the column). By *normalized cell response*, we mean that we first compute the average firing rate of each cell during the tasks and phases of interest:

$$M_i^X = \frac{\int_{t_0^X}^{t_1^X} output_i^X(t)\, dt}{t_1^X - t_0^X} \tag{6.1}$$

where:

M_i^X is the average firing rate of unit i, under condition X.

t_0^X and t_1^X define the phase of interest for condition X.

$output_i^X(t)$ is the firing rate of the output unit of column i, at time t, under condition X.

The *normalized response* over a set of conditions {Y} and a population of cells is computed according to:

$$\hat{M}_i^X = \frac{M_i^X}{\underset{j \in cells,\, z \in \{Y\}}{\mathrm{Max}} \left\{ M_j^z \right\}} \tag{6.2}$$

Population Analysis of Phasic Behavior

The relative phasic responses of an entire F5 population during execution of a precision pinch are shown in Figure 6.8. The set and extension phases of the program (A) show relatively little overlap between the two phases, as indicated by the large number of units aligned along the axes (i.e., nearly a 0 component for one of the two measures). Comparing extension and flexion phases of movement, however, we see a significantly different result. Although there are still some neurons that are exclusively involved in one phase of the movement, a large number of units are shared between the two phases, and the degree of sharing varies significantly from one unit to another. Finally, comparing flexion versus the hold phase of movement, we again have a significant separation between the two populations. Also note that the flexion phase of movement utilizes a significantly larger number of cells than does the hold phase.

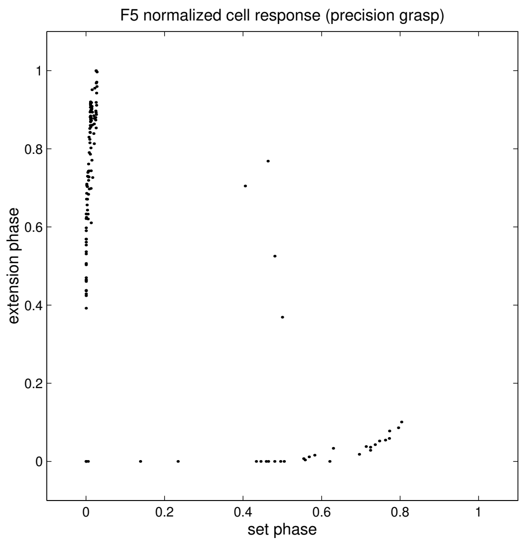

(A)

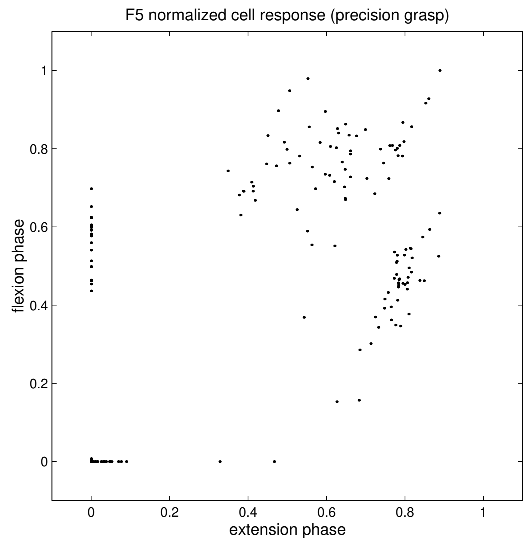

(B)

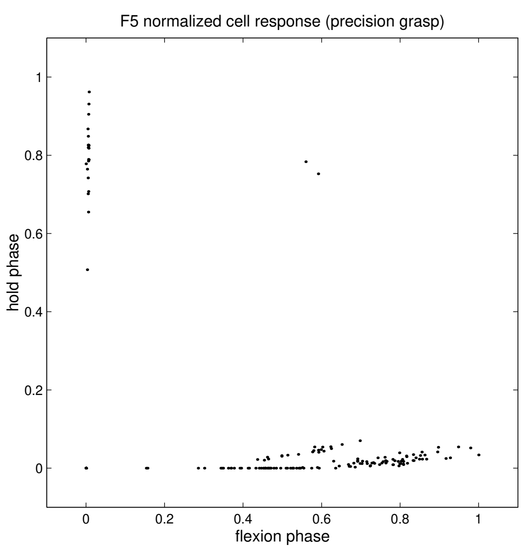

(C)

Figure 6.8: Population analysis of F5 phasic behavior during a precision grasp. (A) set versus extension phase; (B) extension versus flexion; and (C) flexion versus hold.

The AIP population, however, demonstrates a significant amount of sharing between all phases (Figure 6.9). This is the case because once many AIP neurons become active, they remain active over the course of the motor program (through the end of the hold phase). The set and extension phases share a reasonable number of cells, although an equal number only begin to fire as the extension phase begins to execute (A). Note also that the overall response of the cells is greater during the extension phase as compared with the set phase. This is due to the larger amount of recurrent feedback being received from F5 during the movement phase of the program.

Comparing extension phase with flexion, we see a large number of cells that are shared nearly equally (B). However, a few cells demonstrate a very late onset (left-hand column of points). During the hold phase of movement, a number of neurons decrease their activity significantly (the line of points in the lower, right-hand corner of panel C). Otherwise, the two phases of movement share the remainder of the cells equally.

Population Analysis of Visual/Motor Properties

The population analysis technique not only can be used to compare average activity during different phases of the task, but can also be used to compare cell activity under different conditions. This idea is illustrated in the comparison between the lighted fixation and dark movement task (still a precision grasp is being executed in the latter case; Figure 6.10). The most striking effect in the figure is the fact that a large number of cells are only active during dark movement, and are not at all active during fixation. However, a smaller population of cells exhibits both visual and motor responses (and only a small number are exclusively visual-related).

Panel (B) demonstrates the relative activity of the AIP population during the execution of the dark movement task as compared with the light movement task. A significant number of these cells do not show any modulation between the two tasks. However, a few cells are further facilitated during movement in the light. Finally, as a comparison to coding in F5, panel (C) demonstrates that the two grasps that are executed (precision grasp, in the light and dark) are coded using the same set of cells.

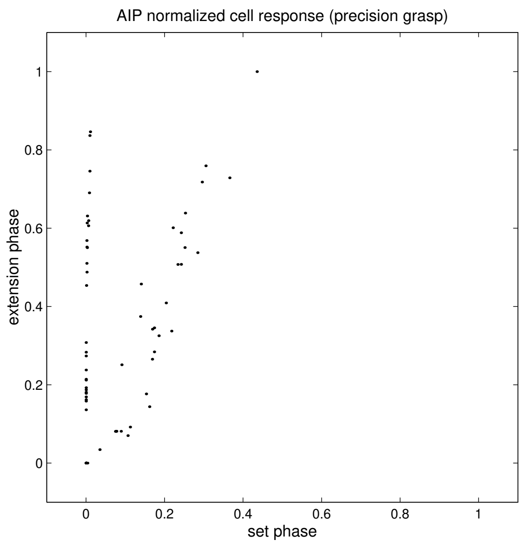

(A)

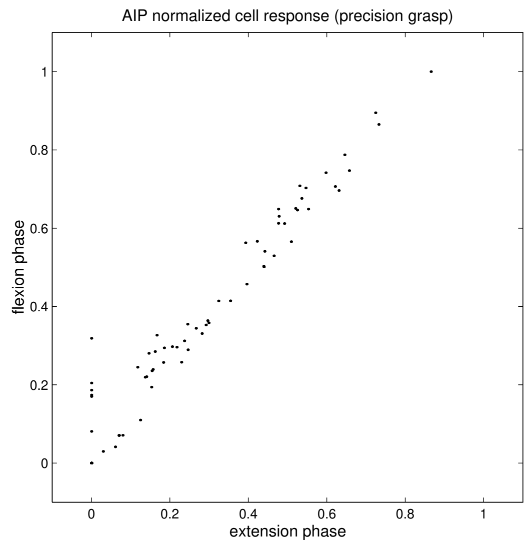

(B)

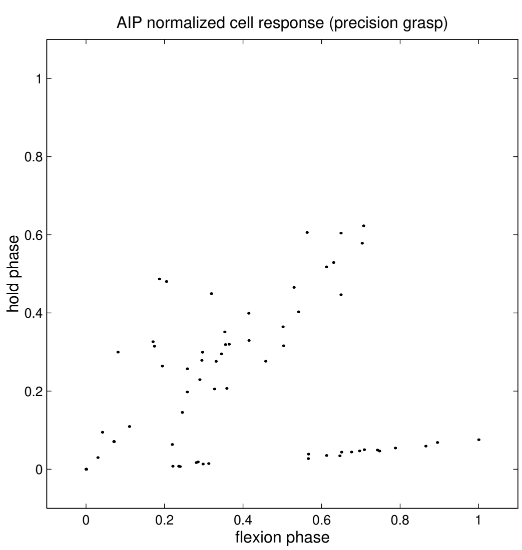

(C)

Figure 6.9: Normalized AIP responses over the entire population; (A) set versus extension phase; (B) extension versus flexion; and (C) flexion versus hold phase.

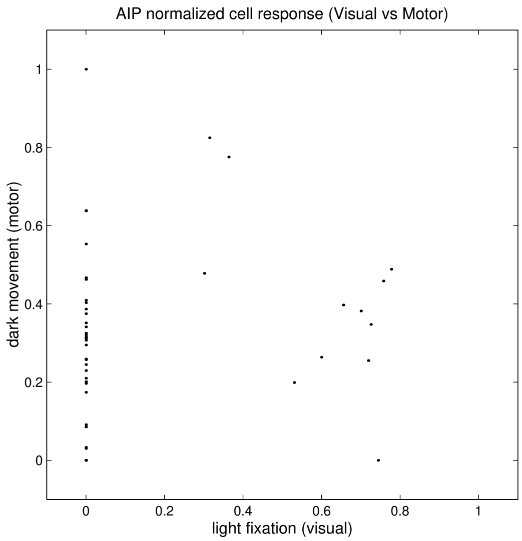

(A)

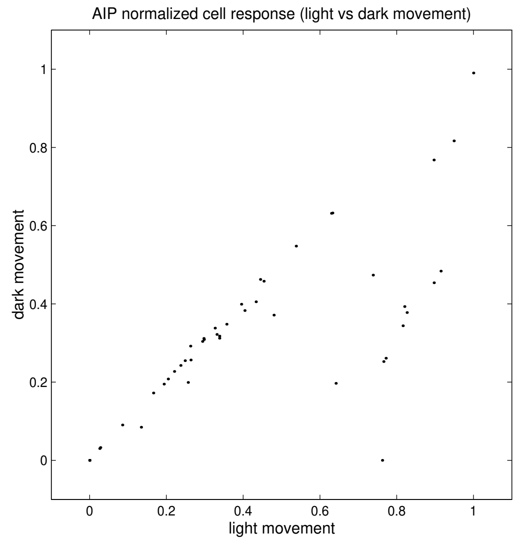

(B)

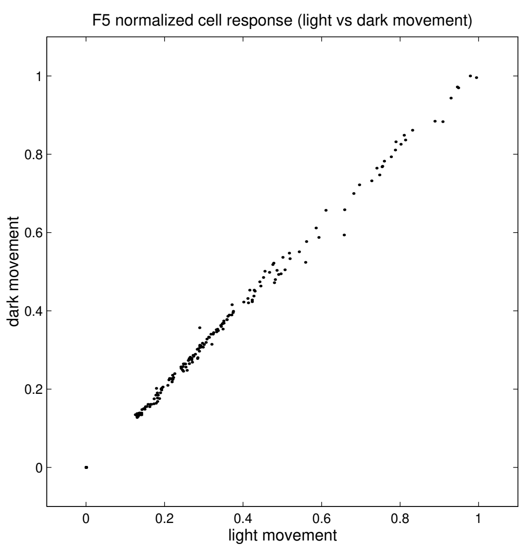

(C)

Figure 6.10: Visual/Motor coding in AIP: comparisons between (A) fixation/dark movement, and (B) dark and light movement conditions; and (C) comparison between light and dark movement coding in F5.

Model Predictions

In this section, we present a series of experiments that are designed to ask further questions about 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 monkey.

The first task examines one aspect of how objects are encoded in AIP (prediction #1). We then turn to a task in which a single object may be manipulated in more than one way and look at how AIP and F5 respond when a delayed instruction stimulus is used to inform the model as to which grasp to use (prediction #2). Third, we examine at how fine parameters of the grasp (in this case the grasp aperture) might be coded in both F5 and AIP (prediction #3), and how these two regions respond when there is a mismatch between the expected object size and the actual object size (prediction #4). Finally, we ask how object parameters are coded in AIP relative to how they are used in the programming of the grasp (prediction #5).

Prediction #1: Object Coding in AIP

In this experiment, the model is presented with two different objects that are roughly the same size (Figure 6.11). 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.

Figure 6.12 shows the responses of two visual-related cells under the two experimental conditions. Unit A is active only for the condition in which the system is grasping the plate. Unit B, on the other hand, is somewhat active during the set and movement phases of the cylinder condition, but is most heavily activated when grasping the plate.

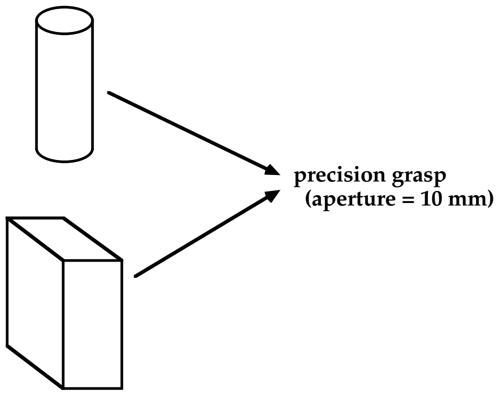

Figure 6.11: Two objects that map to the identical grasp.

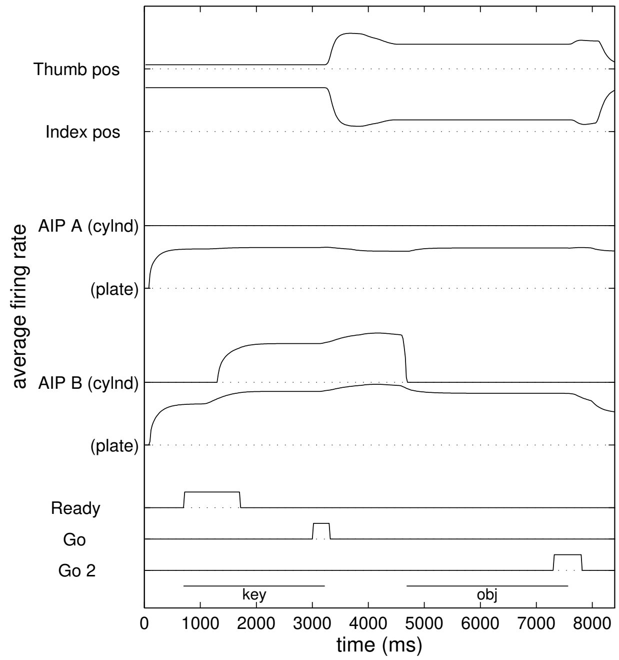

Figure 6.12: Two visual AIP cells that show significant object-specific modulation.

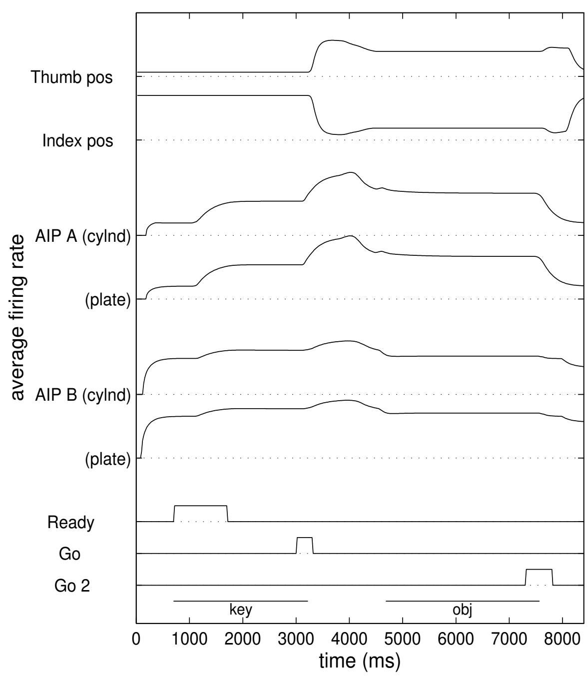

Figure 6.13: Two visual-related (with some motor) AIP cells that demonstrate only slight modulation as a function of the type of object that is presented.

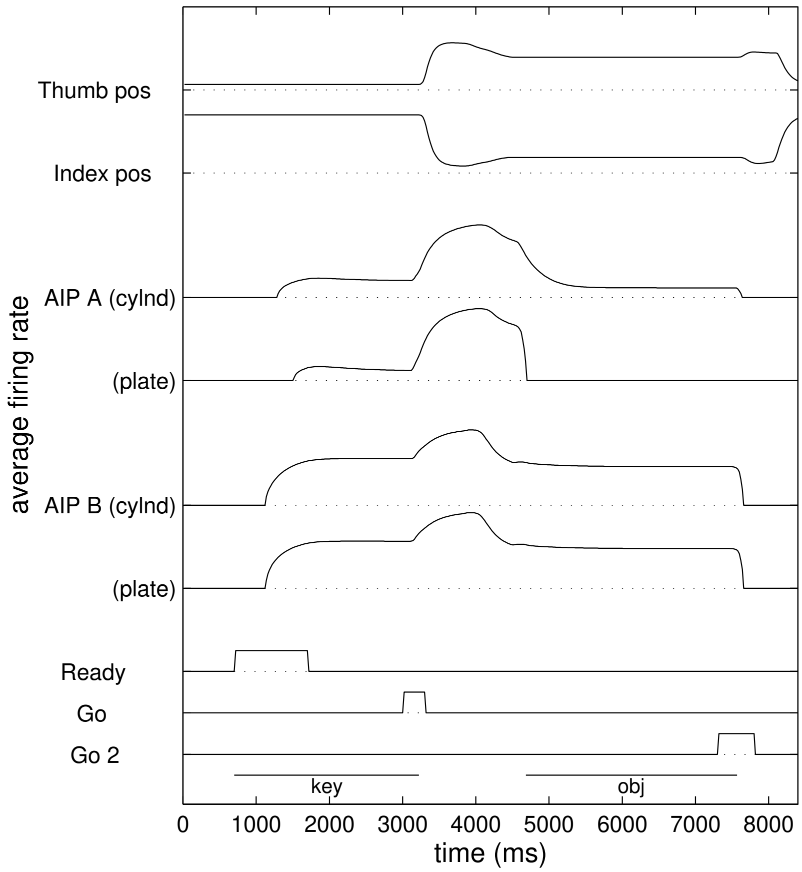

Figure 6.14: Two motor-related (set phase) AIP cells. Only cell A is modulated the type of object that is presented.

Two cells with both visual- and (somewhat) motor-related components are shown in Figure 6.13. In the case of cell A, the primary difference in response level is observable only when the object is first presented to the model (it responds more quickly to the cylinder). Cell B, on the other hand, demonstrates a small overall advantage for the plate condition.

Motor-related cells are less sensitive to the type of object that is presented. Figure 6.14 shows two units that begin to fire after the ready signal is presented. Only cell A demonstrates object-specific modulation, shutting off after the flexion phase of movement is completed. Cell B, along with cells A and B of Figure 6.15 does not change in activity level as a function of the presented object.

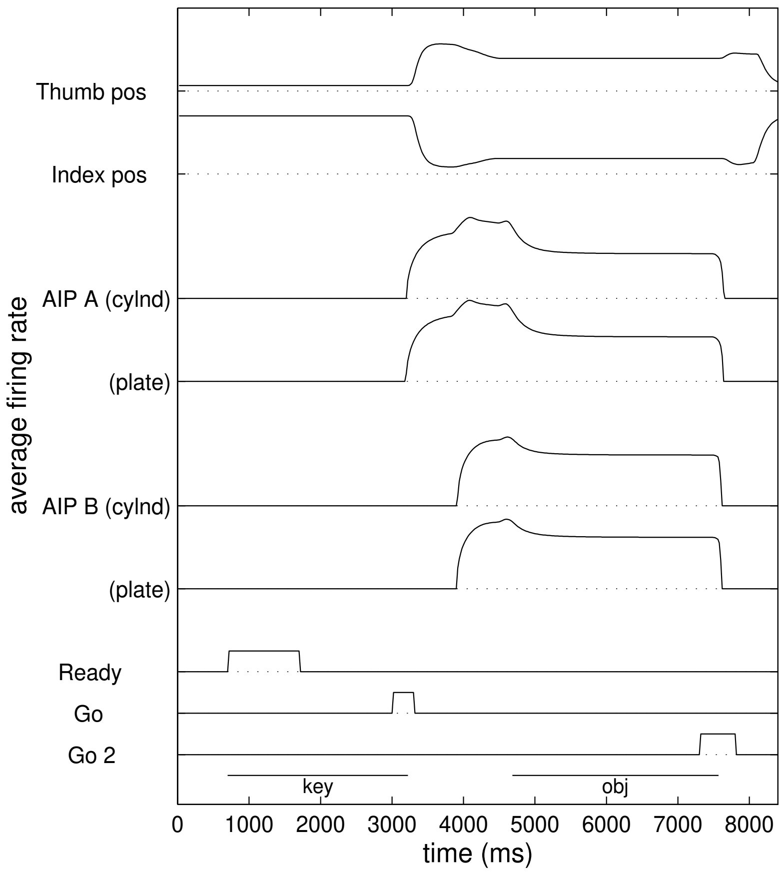

Figure 6.15: Two motor-related AIP cells that are not modulated by the type of object.

Population Analysis of Object-Related Responses

The visual/motor distinction in AIP becomes more clear at a population level of analysis (Figure 6.16). Under the lighted movement condition, the AIP population demonstrates a range of object responses (A). Although a few cells demonstrate significant differences in response to the two objects, many are equally active for both. As a comparison, almost every F5 cell is equally responsive to both objects (B).

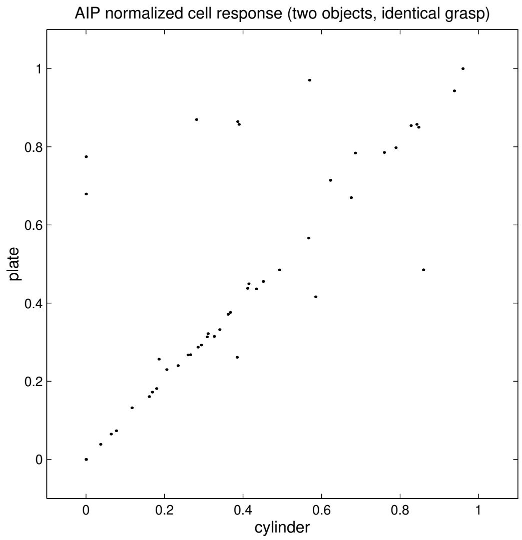

(A)

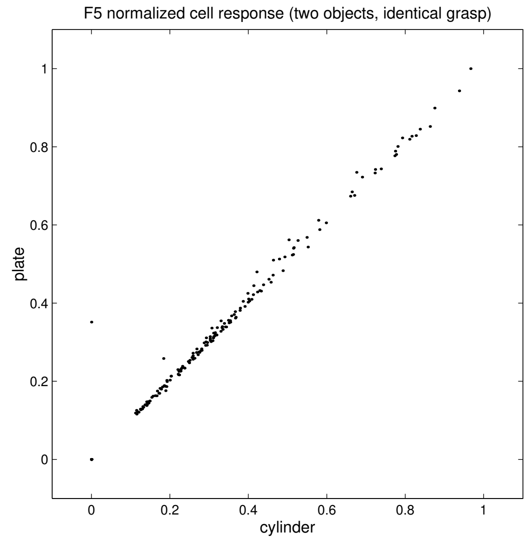

(B)

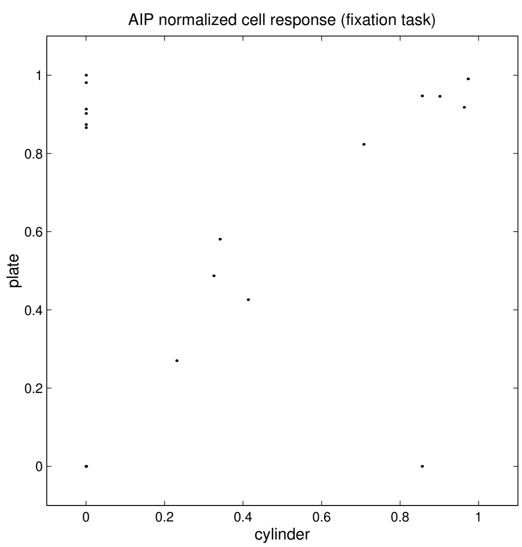

(C)

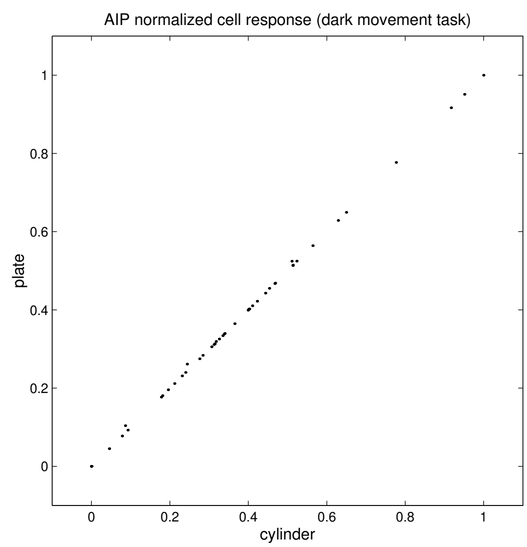

(D)

Figure 6.16: 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.

In order to examine the visual/motor responses of the AIP cells, we compare the cell responses under the fixation (C) and dark movement (D) tasks. The fixation responses are a direct result of connections from visual areas into AIP, and show significant differences in behavior between the two objects. However, under the dark movement task, the active AIP cells do not show object specificity.

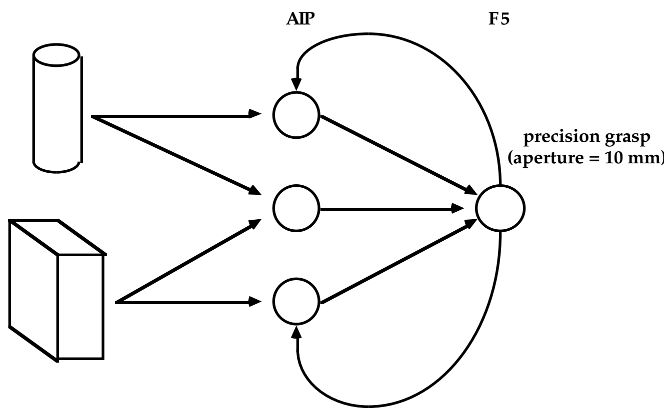

Figure 6.17: Visual-related AIP receive object-specific inputs; motor-related cells receive recurrent inputs from F5, which do not demonstrate object-specific activity.

In general, units that receive a significant amount of visual input (and thus are classified as visual-related cells) are more likely to display differences in how they represent the two different objects, as compared to the motor-related units. As discussed in Appendix A, when a mapping is established from an object representation in PIP (or IT) to a set of appropriate AIP units, the decision as to whether or not to make a connection to an individual AIP unit is determined probabilistically. In this experiment, because the two objects are represented in PIP using non-overlapping populations of neurons, only a small subset of the possible AIP units will respond to both objects (Figure 6.17).

For the most part, however, the object distinction is lost at the level of F5, which is the source of the motor-related activity. 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 AIP and F5. As the activity makes successive cycles through the loop, more and more motor-related cells (from both AIP and F5) are recruited, until all or 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 AIP cells that are classified as more motor-related.

Prediction #2: When an Object Affords Multiple Grasps

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 is captured in the F5/AIP circuitry? And might we be able to see evidence of this mapping in a monkey recording study?

In the following experiment, the model is presented with a single object (a small cylinder), and asked to perform one of three tasks. For each task, the model repeatedly grasps the cylinder a fixed number of times. The three different tasks are:

1. Grasp the cylinder using a precision pinch.
2. Grasp the cylinder using a side opposition.
3. As a function of an instruction stimulus (e.g., the color of a light), grasp the cylinder using either a precision pinch or a side opposition.

When the grasp is known ahead of time (tasks 1 and 2), it is assumed that some higher level planning region 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 arbitrary stimuli with the appropriate grasp. As discussed in the previous chapter, we implicate the dorsal premotor cortex (F2) in this association process (Fagg & Arbib, 1992; Mitz, et al., 1991).

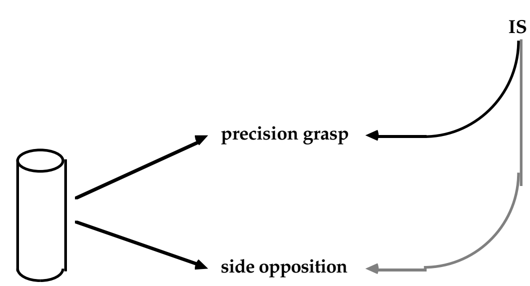

Figure 6.18: A single object mapping to two possible grasps. Before execution, one grasp must be selected based upon the current context (e.g., based upon an Instruction Stimulus).

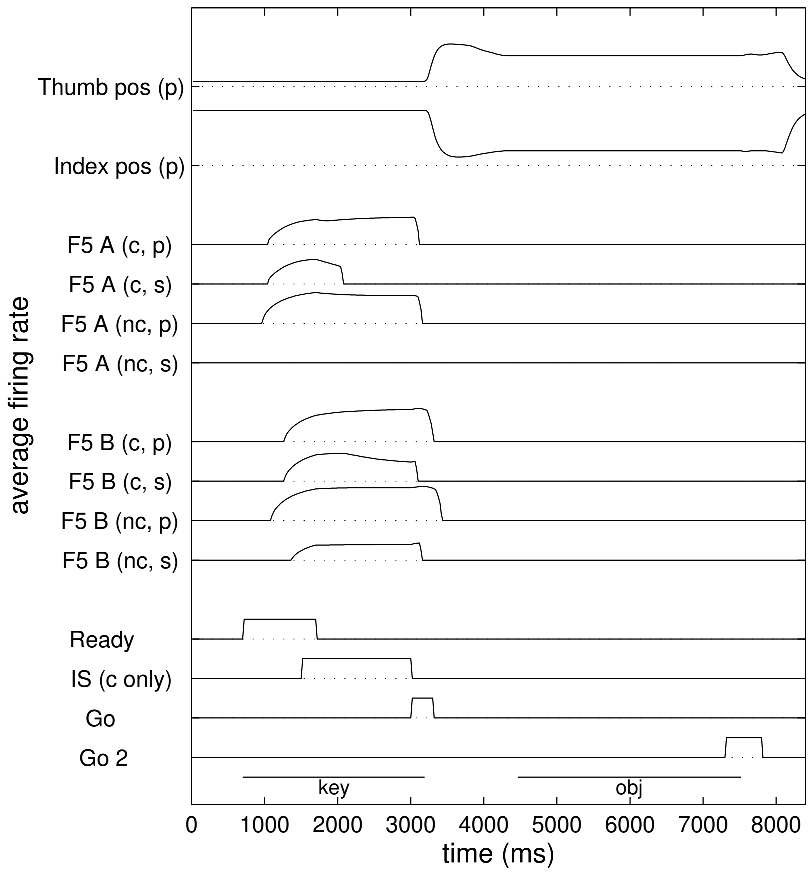

Figure 6.19: Two F5 units in response to the four conditions: (c,p), (c,s), (nc,p), and (nc,s). c = conditional; nc = non-conditional; p = precision grasp; s = side opposition.

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 800ms after the ready signal.

In the following set of figures, we examine in detail the various classes of cell behavior observed in F5 and AIP when presented with these four different conditions. We refer to task 1 as non-conditional/precision, and task 2 as non-conditional/side. Task 3 consists of two separate conditions: conditional/precision and conditional/side.

Figure 6.19 shows the behavior of two set-related cells in F5. For both the precision and side conditional cases, the cell activity is identical up until the IS is presented (as would be

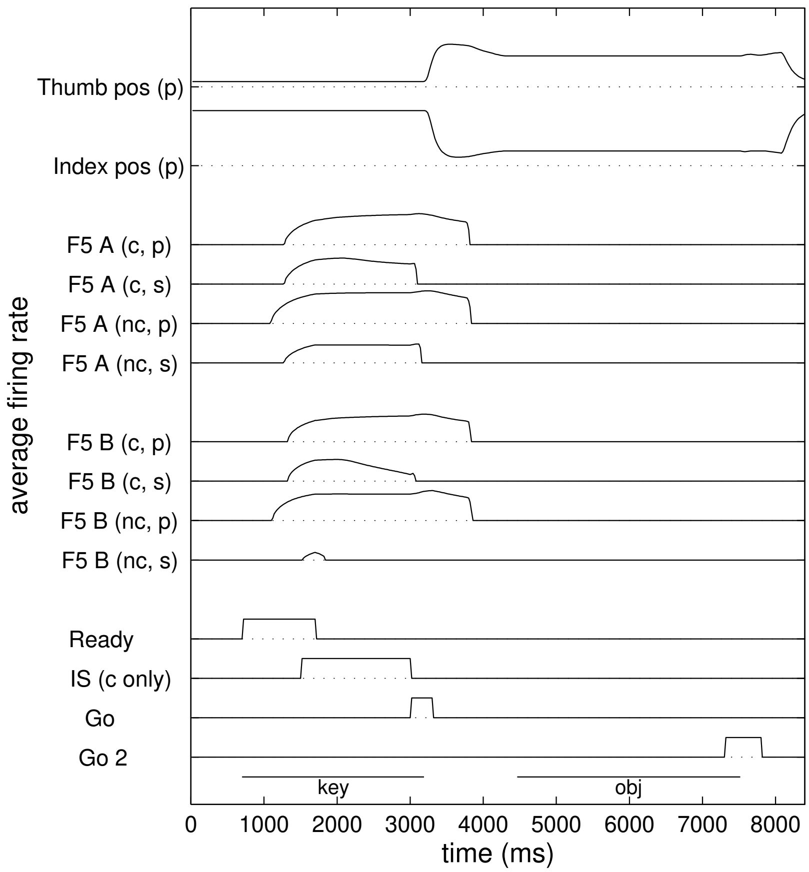

Figure 6.20: A pair of set/extension cells in F5.

expected for all cells). 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 side condition (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 side opposition case, this cell is significantly inhibited by other F5 units and thus cannot become 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). Note, however, that for the precision grasp case, B is much more active than the side opposition case.

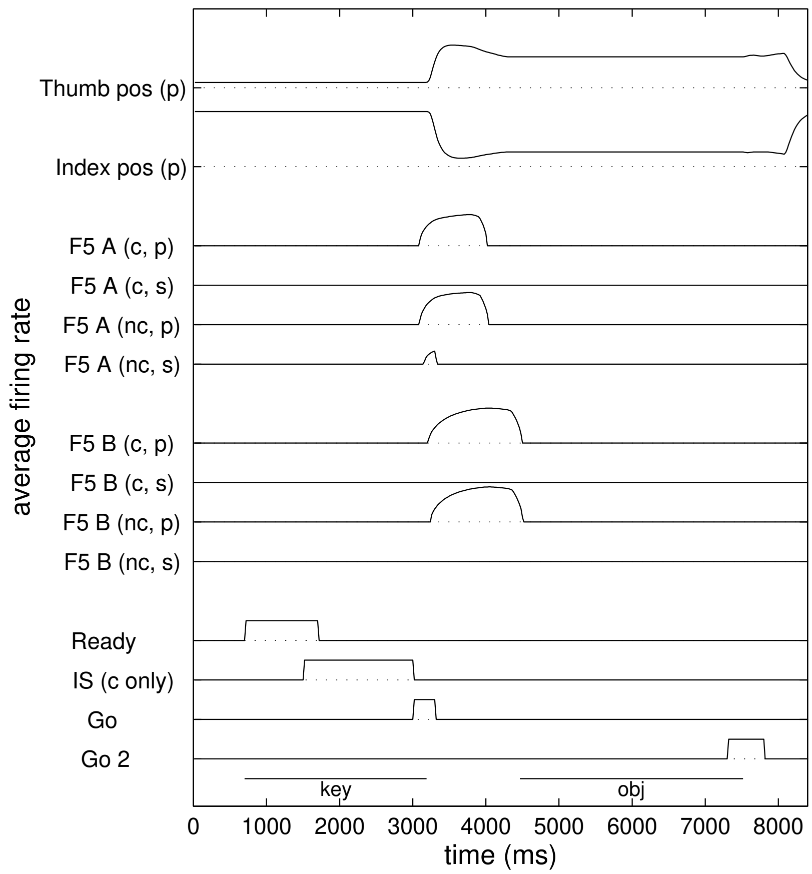

Figure 6.21: A pair of extension-related F5 cells.

Two combination set and extension-related F5 units are illustrated in Figure 6.20. The key behavioral feature of these units is the fact that both cells continue to fire into the extension phase of movement only when a precision grasp is being executed. When the side opposition is executed, these cells are shunted by the time the movement is initiated.

This shunting of competitor grasps just prior to or during the extension phase of movement is a feature that naturally comes out of the structure of the model. Because there are many more extension units in the F5 region than there are set-related units (Rizzolatti, et al., 1988), when phase transition takes place in response to the go signal, we see a rapid increase in the number of active cells in the region. In turn, there is a rapid increase in the level of inhibition through

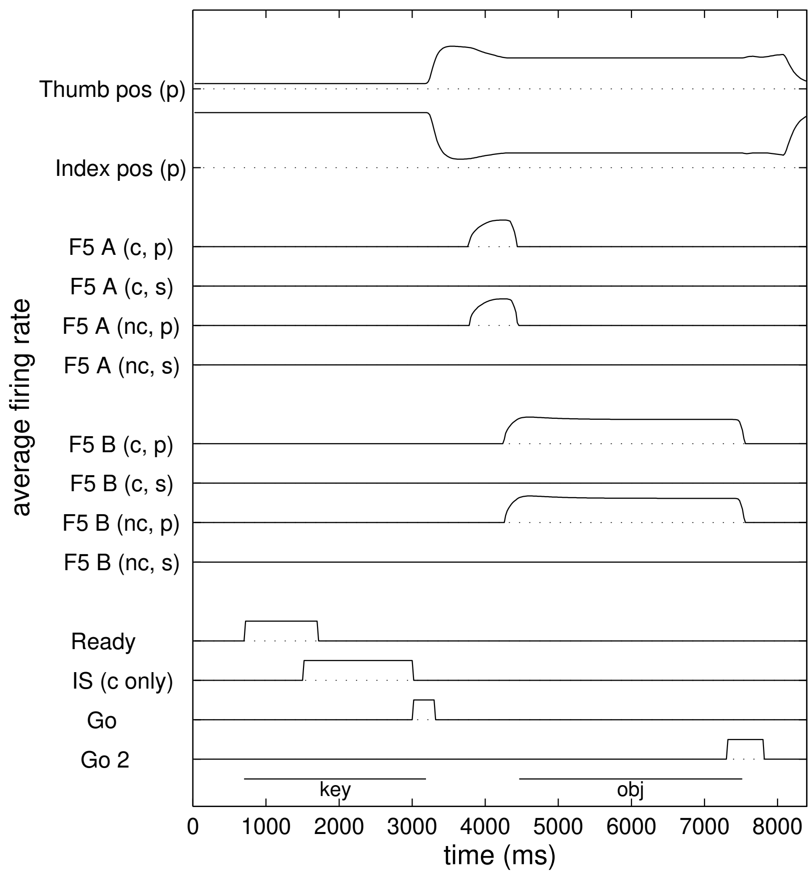

Figure 6.22: F5 flexion and hold units.

the basal ganglia inhibitory loop. The result is a sudden sharpening of the motor program that is about to be executed.

We see this effect in Figure 6.21, which shows a pair of extension-related F5 units. In only one competing case (execution of the side opposition), does cell A become active. However, it only does so for a short period of time. Once the program is sharpened, units participating in the flexion (Figure 6.22, cell A) and hold (cell B) phases of movement are generally compatible with the program being executed.

Because the number of hold-related cells is rather small in the F5 region, the motor program sharpening disappears to some degree during the release phase of movement (Figure 6.23, cell B).

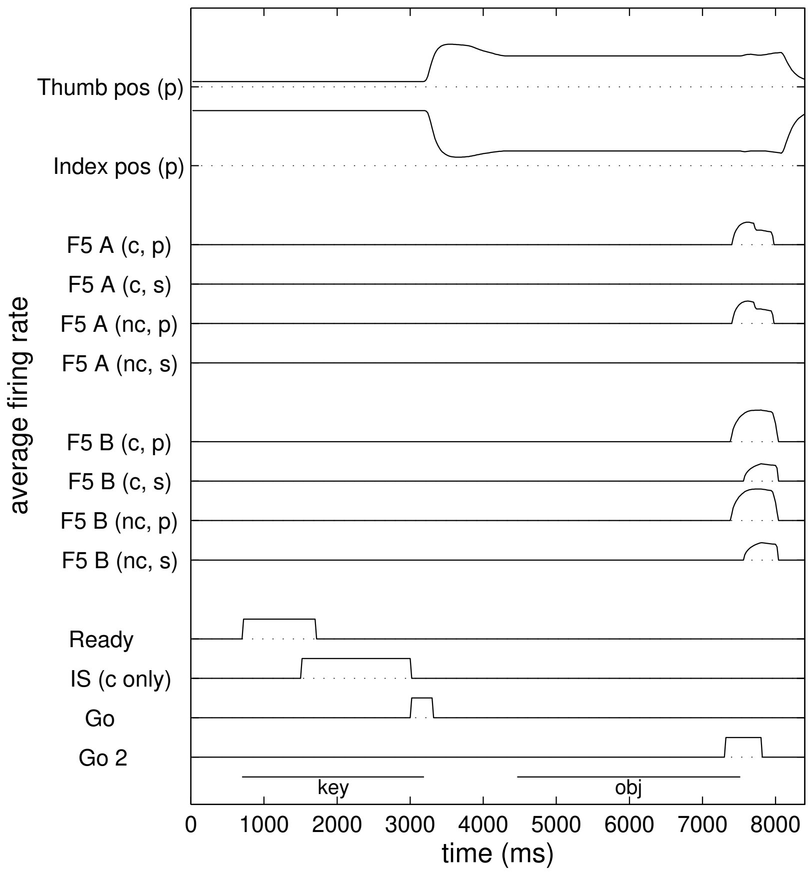

Figure 6.23: A pair of release phase F5 units.

Finally, Figure 6.24 illustrates a pair of set-related F5 cells whose temporal properties are significantly modulated by whether the system is performing a conditional side opposition or a non-conditional side opposition. In the conditional case, both cells delay their activation by at least 400ms. This happens because the cells are below threshold unless they receive both visual input via AIP and positive supporting inputs from either F2 or F6. Neither cell is significantly active for either precision grasp case.

The AIP cells in this conditional/non-conditional task reflect similar behavior as compared with the F5 units. The visual/motor cell shown in Figure 6.25 activates in response to the presentation of the object. In the cases when a side opposition is executed, this cell is

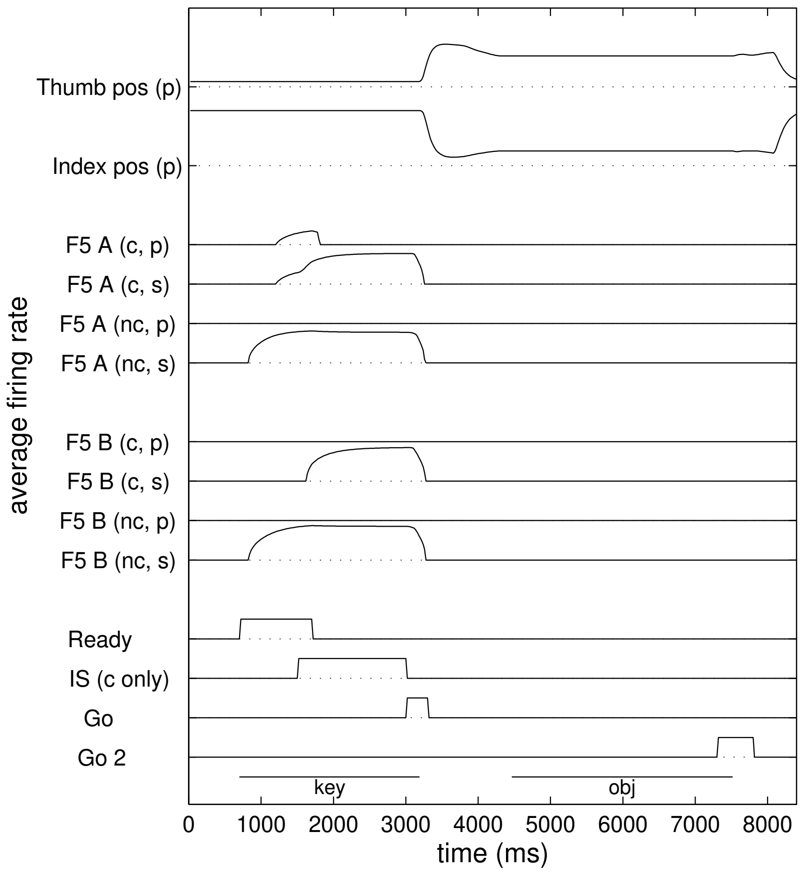

Figure 6.24: A pair of set-related F5 units. Note that in both cases, their temporal behavior is modulated by whether a conditional or a non-conditional task is being performed.

shunted during the set phase. The shunting is a result of competition with other AIP units via the inhibitory basal ganglia loop.

The visual/motor AIP unit illustrated in Figure 6.26 receives a more significant visual projection than the previous cell (Figure 6.25). This is evident in the fact that even in the side opposition cases, the cell is not completely inhibited (although it is significantly depressed).

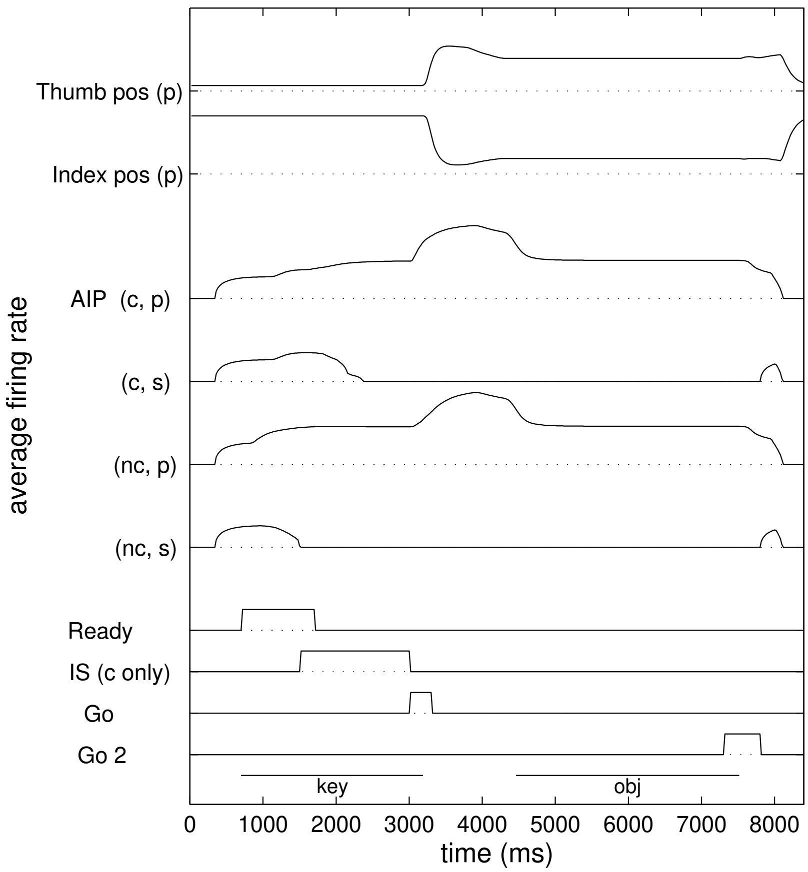

Figure 6.25: A visual/motor AIP cell (more motor-related). The notation is the same as in the previous figures.

Finally, the cells shown in Figures 6.27 and 6.28 are both purely motor-related. The former is a set-related unit, and demonstrates some level activity for all four conditions. In the non-compatible conditions, this cell is shunted by the time movement begins. The latter cell on the other hand (Figure 6.28), is never active in the non-compatible conditions. This is the case because it does not receive recurrent input from F5 set cells. By the time the movement is initiated, any possible F5 inputs into this cell will have been shunted due to the sharpening process that occurs at the set/extension phase boundary.

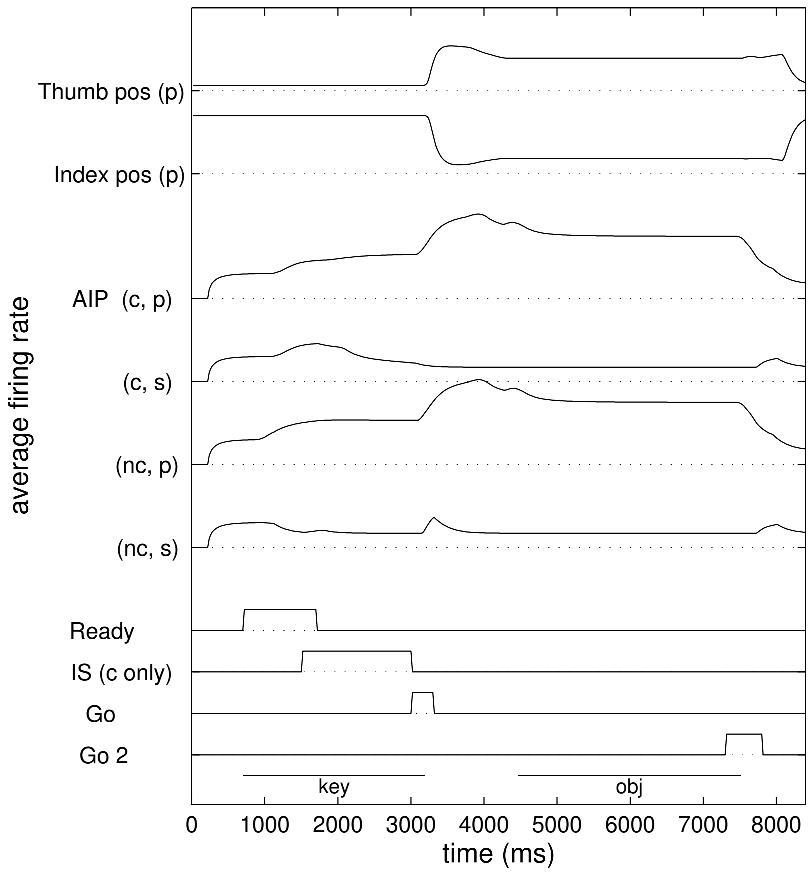

Figure 6.26: A visual/motor AIP cell (primarily visual), which participates in the execution of the precision grasp. During non-compatible (side opposition) conditions this cell is shunted.

Population Analysis Comparing Conditional and Non-conditional Tasks

The population responses of the cells in F5 and AIP are summarized in Figures 6.29 and 6.30, respectively. In order to explicitly illustrate the sharpening process, normalized cell responses are computed for both the set and movement phases of the task. Many F5 set cells, under the conditional task, are active to varying degrees in both the precision grasp and side opposition cases (6.29.A). However when the grasp is known a priori (non-conditional task), relatively few set cells are active for both cases (panel B). In this task, when the ready signal is received, set cells from both grasps begin to respond. Because of the grasp-specific bias from F6, however, those cells that correspond to the specified grasp are the ones that first achieve threshold.

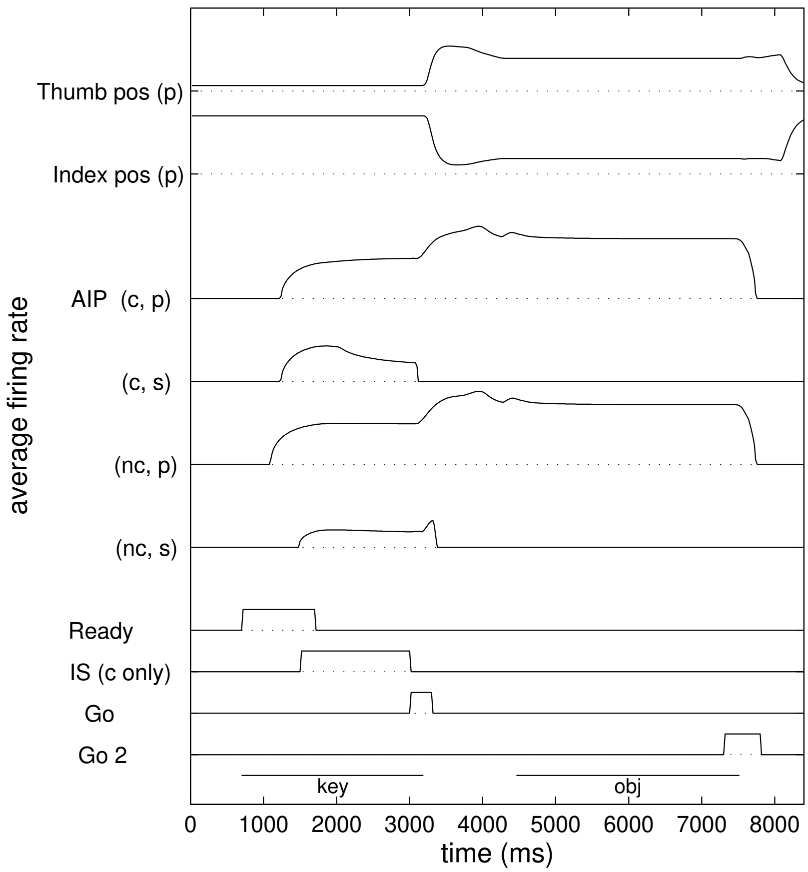

Figure 6.27: A primarily motor-related AIP cell, which is completely inhibited during execution of the side opposition.

Via the inhibitory BG loop, the other set cells (of the non-specified grasp) are then inhibited (in most cases, before they reach threshold).

Comparing the set responses to the movement-related responses (panels C and D), we see that cells participate essentially in one grasp or the other. This demonstrates that as a population, by the time movement is initiated, the grasp decision is completely made.

The AIP cell population reflects a similar behavior as the F5 population when comparing the conditional and non-conditional tasks (Figure 6.30). The primary difference is that a larger number cells demonstrate reasonable levels activity for both side and precision grasps. This is the case because many of the cells receive the non-discriminant visual input (even after the grasp is selected), and because the competition between affordances is not as significant as

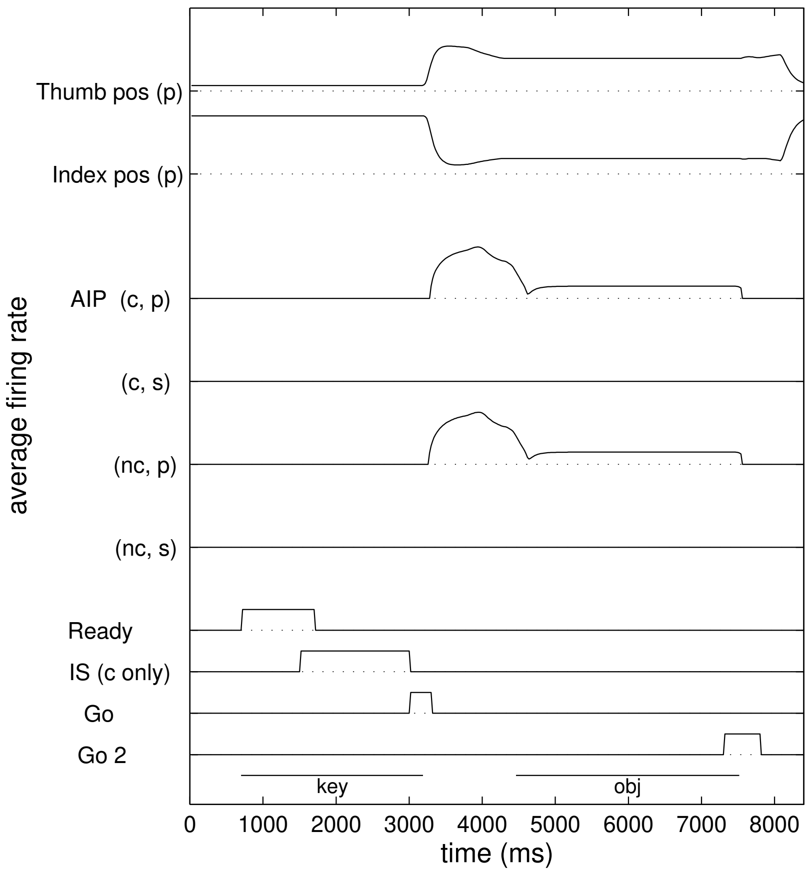

Figure 6.28: A pure motor-related AIP cell, which is active only under the precision-pinch condition.

is the competition that takes place within F5 between the different grasp alternatives. However, note that despite the fact that many cells are active for both grasps, most still show a significant orientation towards one of the two possibilities.

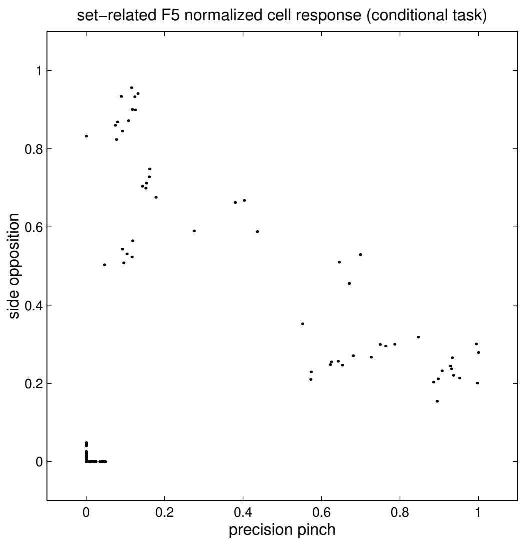

(A)

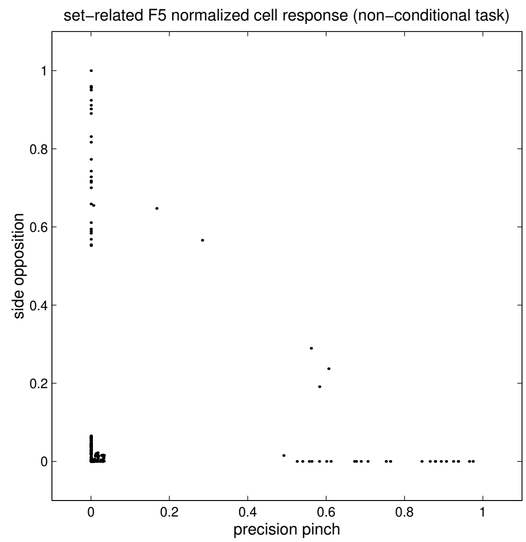

(B)

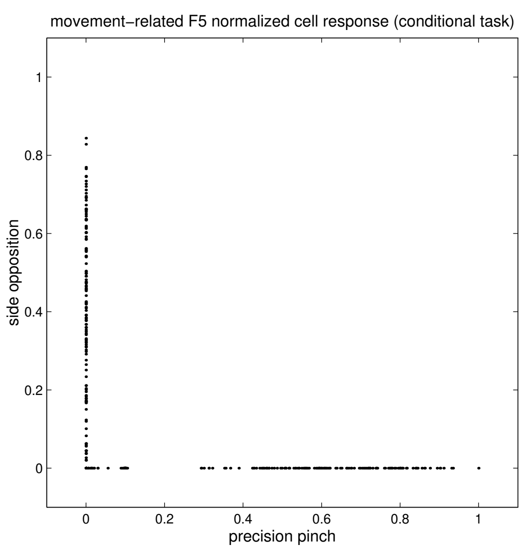

(C)

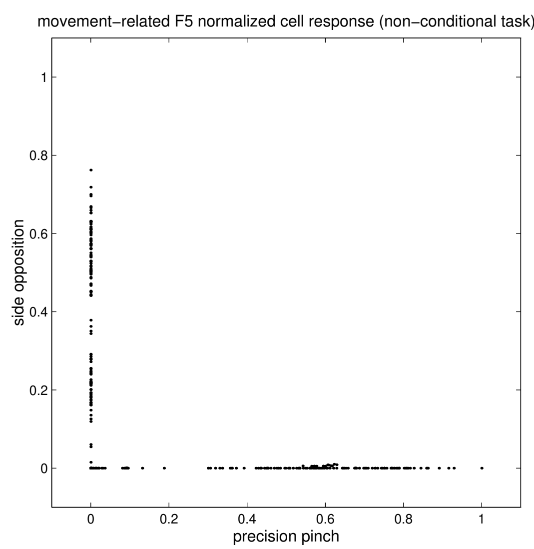

(D)

Figure 6.29: Comparison of set-related F5 responses towards two different grasps under conditional (A) and non-conditional (B) tasks; and movement-related responses under conditional (C) and non-conditional (D) tasks.

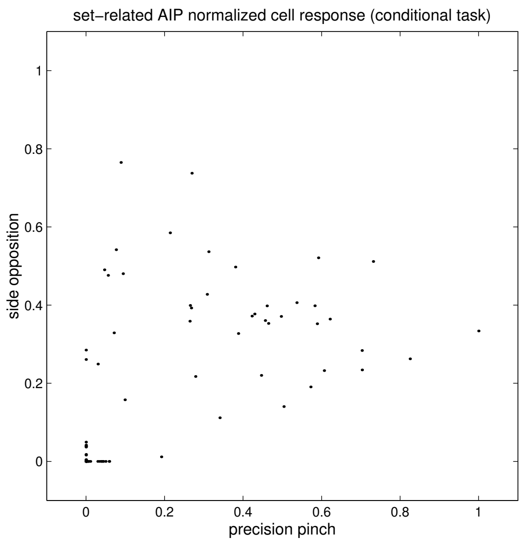

(A)

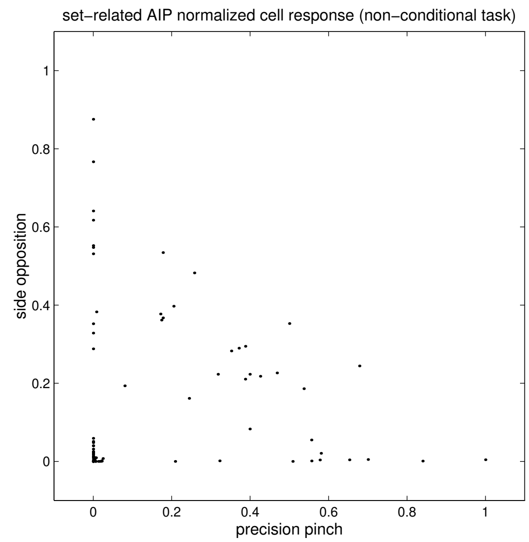

(B)

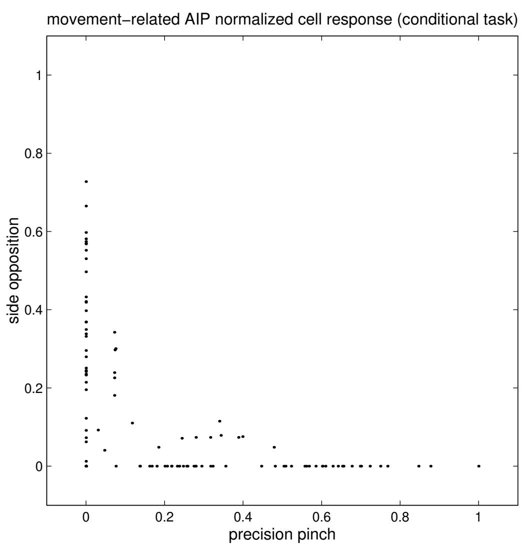

(C)

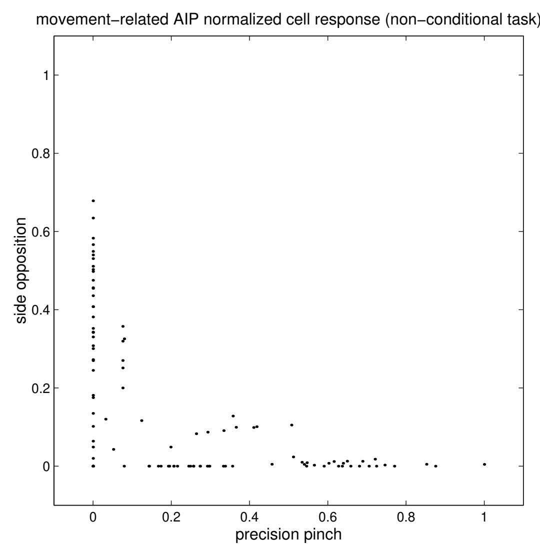

(D)

Figure 6.30: Comparison of set-related AIP responses towards two different grasps under conditional (A) and non-conditional (B) tasks; and movement-related responses under conditional (C) and non-conditional (D) tasks.

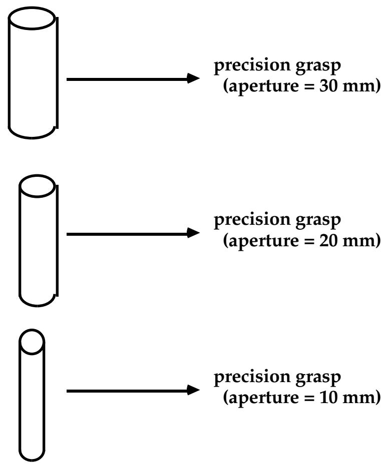

Figure 6.31: Cylinders of different widths map to a precision grasp of varying aperture size.

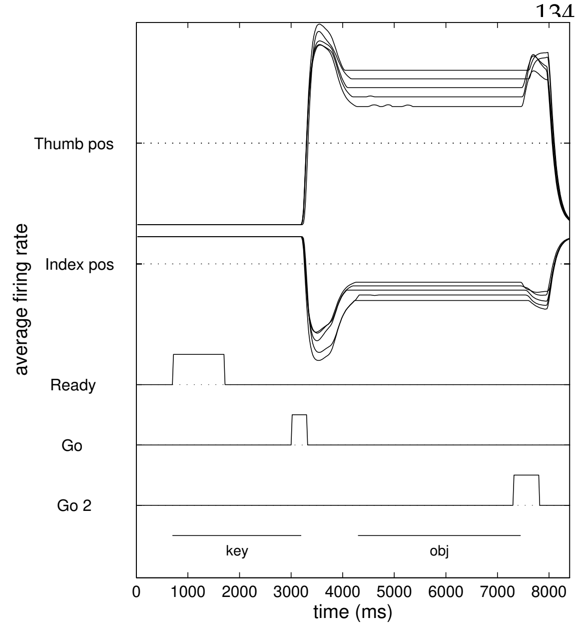

Figure 6.32: Thumb and index finger temporal behavior as a function of cylinder size.

Prediction #3: Object Size Coding

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., 1991a). How is this information encoded in the various grasping motor programs? (Murata, et al., 1993) has reported that within AIP, some cells are modulated by the size of the object that is presented to the monkey. As discussed in chapter 5, we extrapolate this observation to F5, assuming that sub-populations of grasp-specific units encode the aperture of grasp (using a population encoding scheme).

In this experiment, we present the model with a set of cylinders which range in diameter from 10 to 40mm (Figure 6.31), and examine the resulting responses of a set of AIP and F5 units. Figure 6.32 compares the behavior of the thumb and index finger under these different conditions. As the cylinder diameter increases, both the maximum aperture and point of contact anticipated by the model increase accordingly.

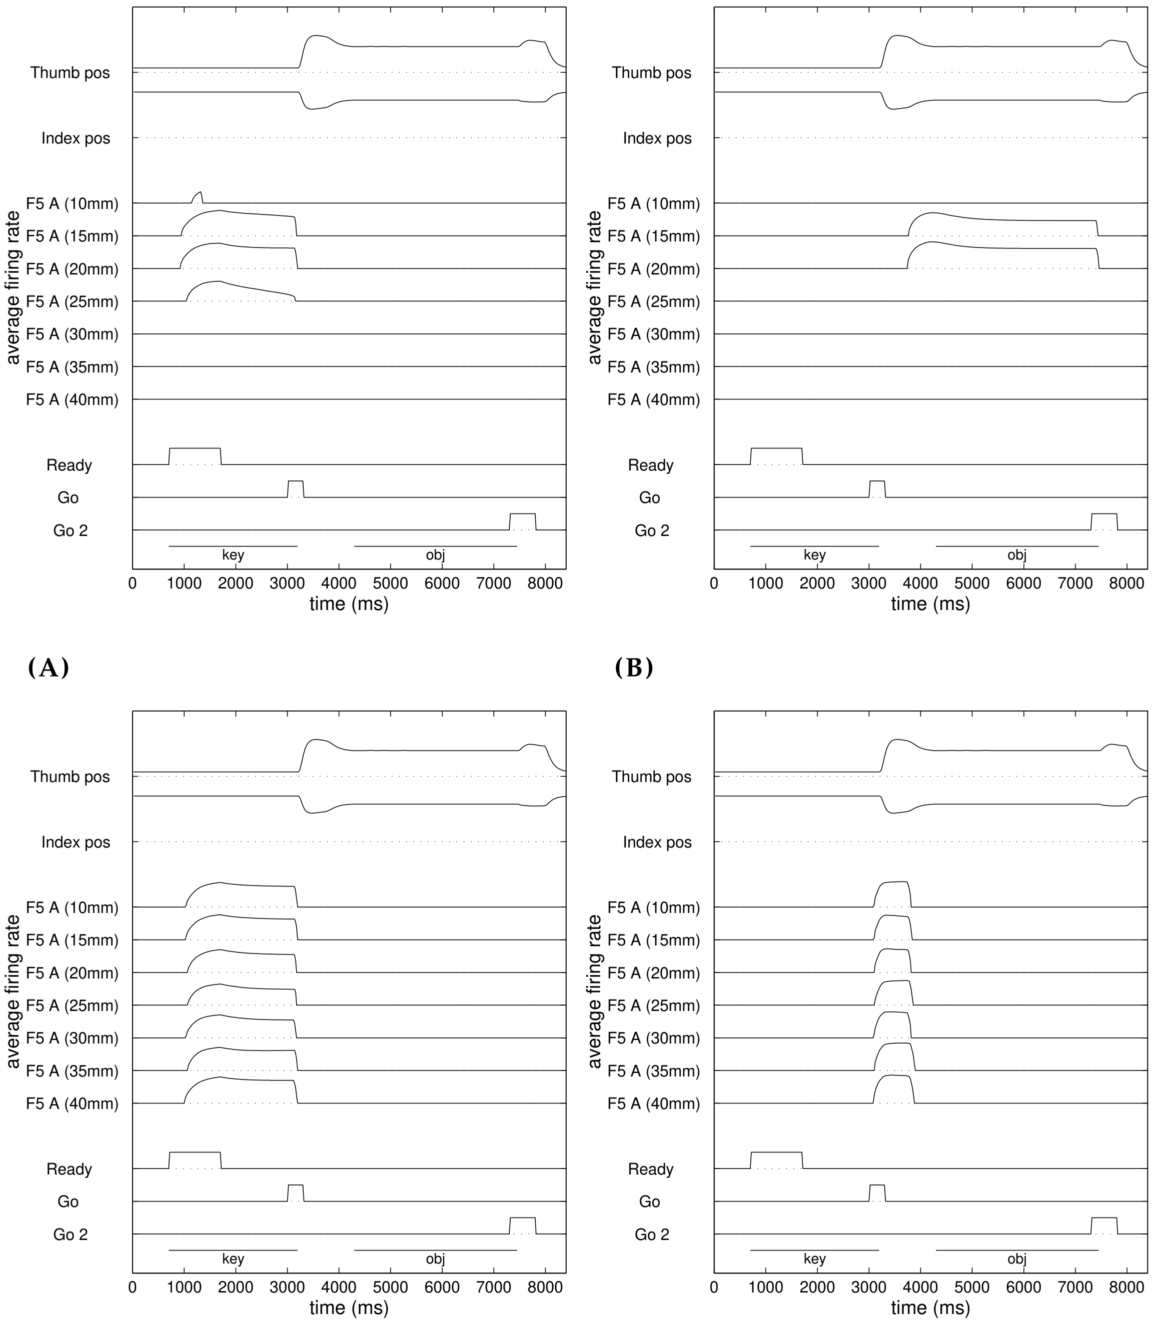

Figure 6.33: F5 cell responses during precision grasps of different apertures. Aperture-specific set (A) and hold (B) units; and aperture non-specific set (C) and movement-related (D) units.

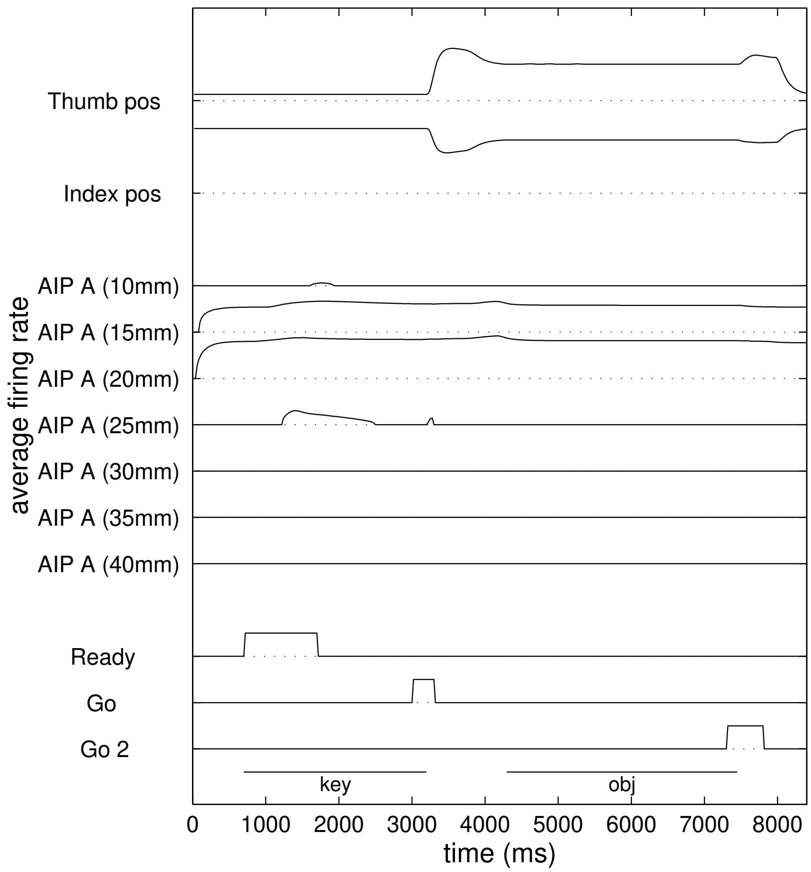

Figure 6.34: Visual-related AIP cell that demonstrates object size specificity (preferred aperture ~= 20mm).

F5 cells exhibiting both aperture-specific and aperture insensitive properties are illustrated in Figure 6.33A-D. Note that both aperture sensitive cells (A and B) are responsive over a range of grasp apertures. In addition, each cell responds maximally for one or two object sizes, and as the objects become larger or smaller than these *preferred sizes*, the response level drops off rapidly. In contrast, cells C and D are not significantly modulated by the size of the object being grasped.

AIP cells also demonstrate a range of behaviors relative to the size of object presented to the system. Sub-populations of both visual- (Figure 6.34) and motor-related (Figure 6.35) units encode object (or grasp aperture) size. Other sub-populations of units encode the precision grasp in general and are only modulated to a small degree by object size (Figures 6.36 and 6.37). Note

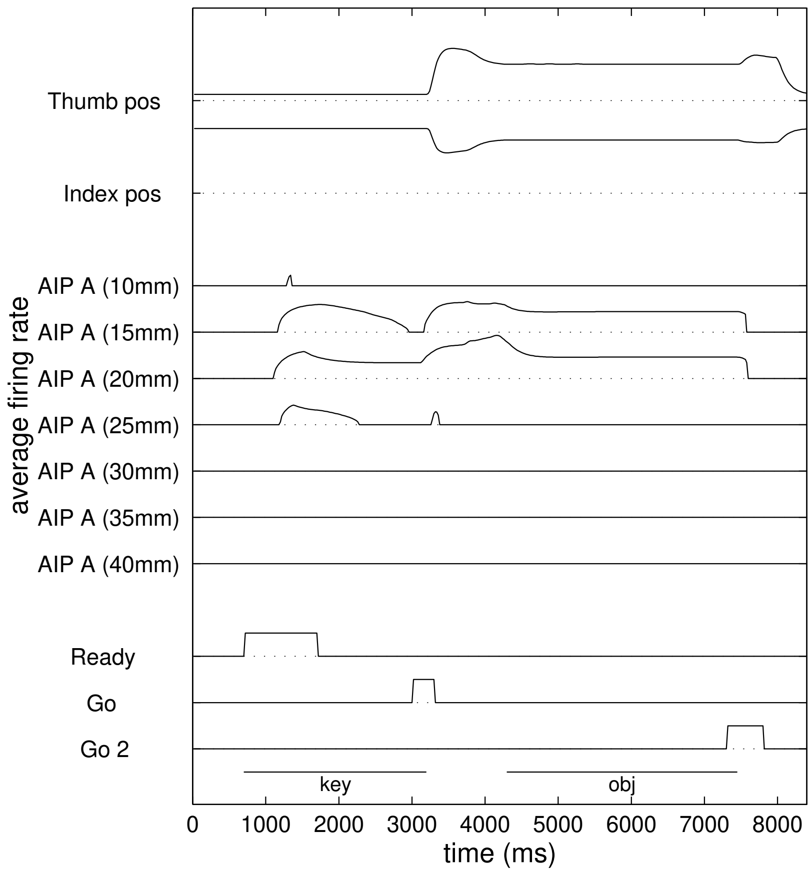

Figure 6.35: Motor-related AIP cell (preferred aperture ~= 20mm).

that in the latter case, the visual/motor cell is shunted during the hold phase of the task for mid-range cylinder sizes, but is active for small and large sizes of cylinders. In addition, the cell exhibits a late onset for only the cylinder of width 35mm. These effects are due to a combination of the specific group of F5 units from which this cell receives recurrent projections, and to the general competition that takes place within the AIP layer (via the BG inhibitory loop).

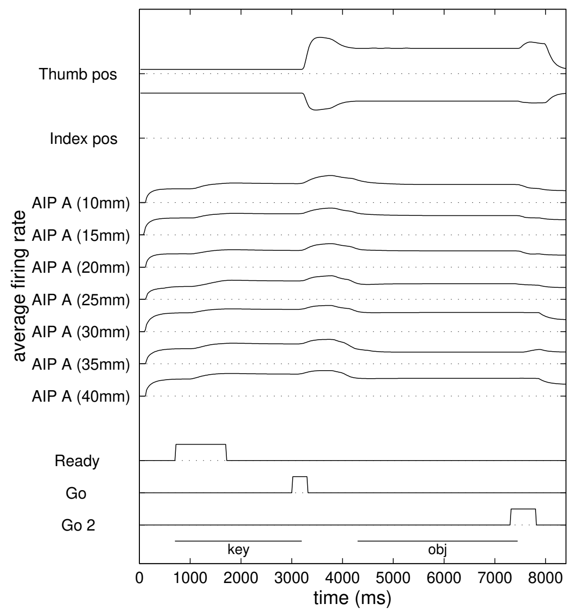

Figure 6.36: Visual-related AIP cell that demonstrates little modulation due to object size.

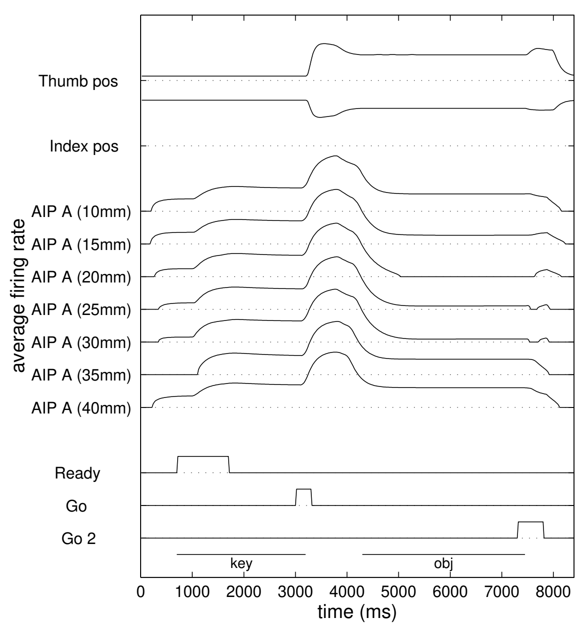

Figure 6.37: Visual/motor AIP cell that is inhibited during the hold phase of movement for mid-range cylinders.

Prediction #4: Object Size Perturbation

One perspective on AIP's role in the grasping process discussed earlier is the idea that AIP is maintaining an *active memory* of the grasp that is about to be or is currently being executed. If this is the case, then we expect that any changes in the executed grasp (from the planned grasp) 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, this sort of perturbation is sensed by hyperfeature detectors in SII at the time of contact with the object. Because the hyperfeature units are sensitive to very specific combinations of tactile and proprioceptive (population-coded) inputs, 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). This shift in SII contact triggering input into the F5 layer causes a different set of F5 (hold) units to be activated than would normally be the case. This new set of F5 neurons comprises units that are tuned for holding an object that is the same size as the object actually encountered.

This shifting of activity from one grasping program to another is demonstrated in Figure 6.38. Average firing rate traces for two units (A and B) are shown for four different conditions (grasping of a 20mm cylinder, a 30mm cylinder, perturbation of 20 to 30mm and perturbation from 30 to 20mm). 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 seeing the 20mm cylinder. Unit B, on the other hand, is a hold-related cell. Although its preferred aperture is 20mm, it also activates for a short period of time when a 30mm 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

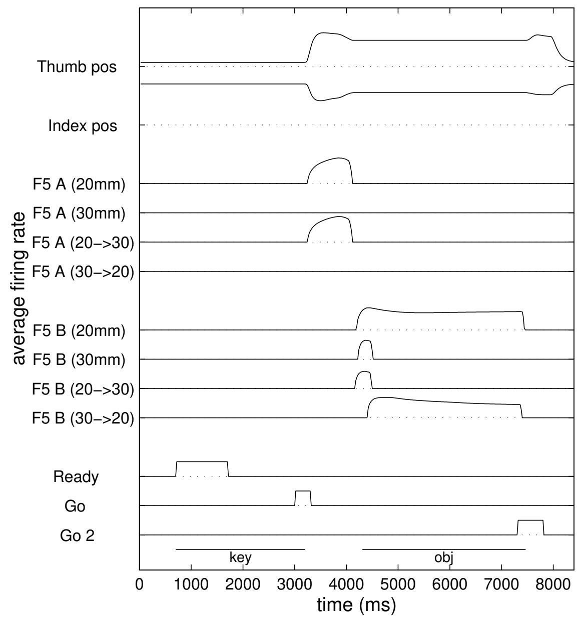

Figure 6.38: F5 movement-related cell (A) and a hold-related (B) cell during the perturbation experiment. 20mm/30mm traces correspond to presentation and grasping of a 20mm and a 30mm cylinder, respectively; traces labeled 20->30 and 30->20 indicate perturbation trials, in which a 20mm cylinder is switched for a 30mm cylinder, and a 30mm cylinder for a 20mm one, respectively.

20mm and 30->20 conditions) is delayed by about 200ms. This delay is due to a combination of 1) the priming of the SII unit in anticipation of contacting a 20mm 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 30mm grasping program in F5.

Examining the population responses for F5 cells, we see that not all F5 cells shift in response to a perturbation. In Figure 6.39, we plot the hold-related activity of F5 units that have a preferred aperture between 20 and 25mm. Comparing the responses for 20 and 30mm grasps, we see that the two grasps share some subset of cells (panel A), due to the population coding of grasp aperture. In general, however, these cells are more oriented towards coding a precision grasp of aperture 20mm. When a perturbation occurs (panel B), a subset of these cells react by becoming even more active, thus shifting the currently executing program from the 30mm precision grasp to the 20mm grasp.

This shift in F5 activity subsequently causes a change in the AIP representation of the ongoing grasp. Here we illustrate the main classes of AIP responses. The visual-related cell demonstrated in Figure 6.40 responds during the set phase to both objects nearly at the same level. However, during movement in the 20mm object (perceived size) conditions, this cell is shunted; the activity returns when a grasp is established with an object of size 20mm.

The motor related unit of Figure 6.41 demonstrates a clear movement-related response that is oriented towards objects of perceived size of 30mm. However, the cell maintains a hold response only for 20mm (actual size) objects. In contrast the visual/motor cell of Figure 6.42 is highly active in all four conditions, but maintains a significant level of activity during the hold phase only when an object of 30mm is being held.

Finally, some cells demonstrate a specific behavior for only one of the four tested conditions (Figures 6.43 and 6.44). In the first case, the unit activates during the set phase in 20mm conditions, but only maintains this state during the hold phase when the object size matches the perceived size. In the latter case, the cell is shunted during the hold phase of the task only when an object of 30mm is both perceived and grasped.

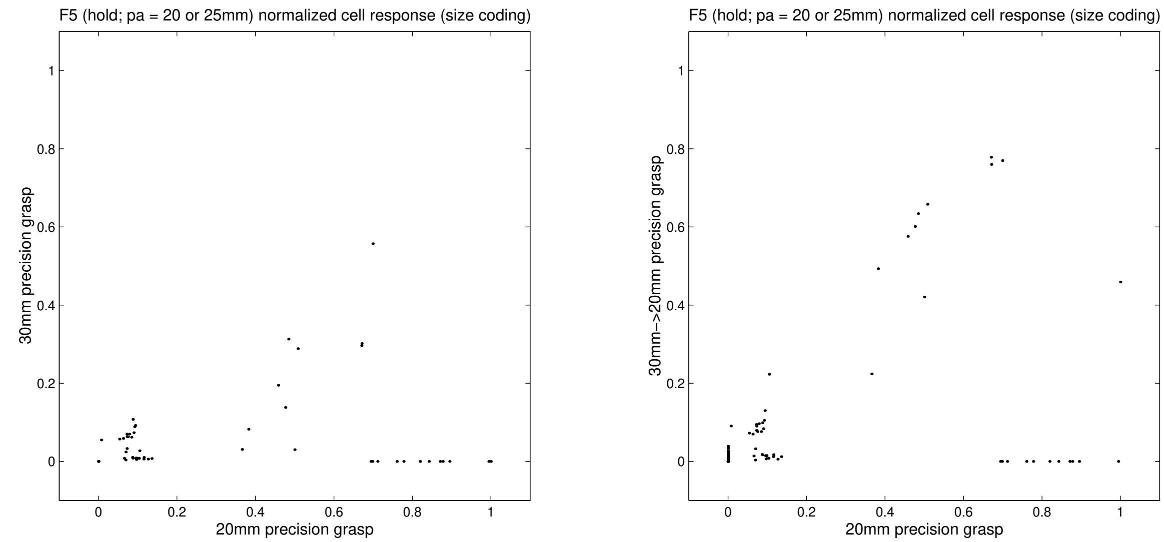

(A) **(B)**

Figure 6.39: F5 cell responses (of cells with preferred aperture between 20 and 25mm only) during execution of the hold phase of a 20mm precision grasp, compared with both the execution of a 30mm precision grasp (A) and a perturbed grasp (B).

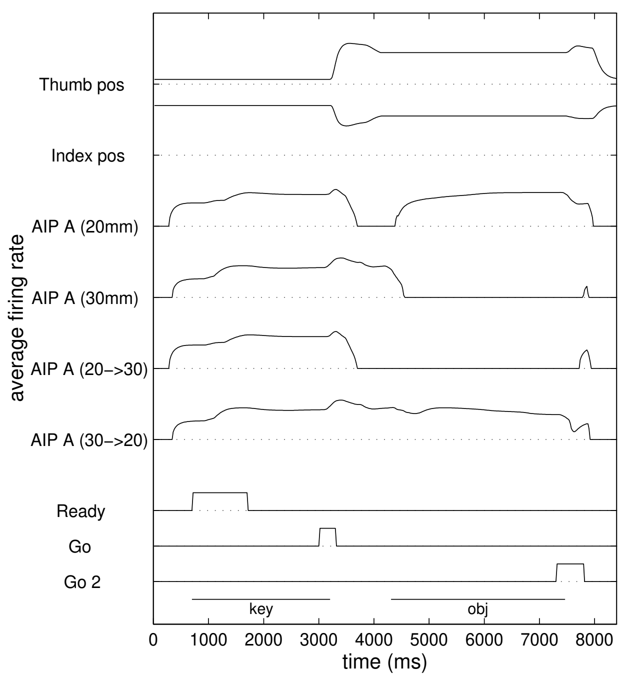

Figure 6.40: AIP visual cell that is active for all four conditions, but is shunted during the hold phase of a 30mm object.

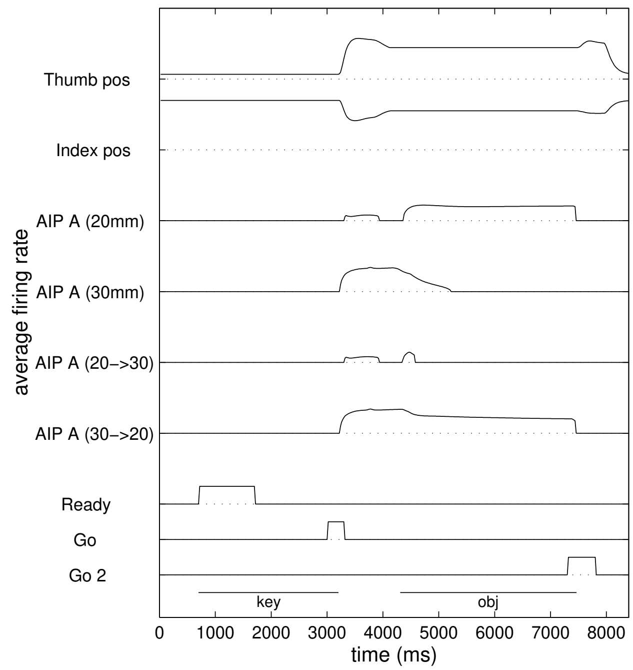

Figure 6.41: AIP motor-related cell that is most active during the movement phase of 30mm aperture grasps, but only continues to be active through the hold phase of 20mm grasps.

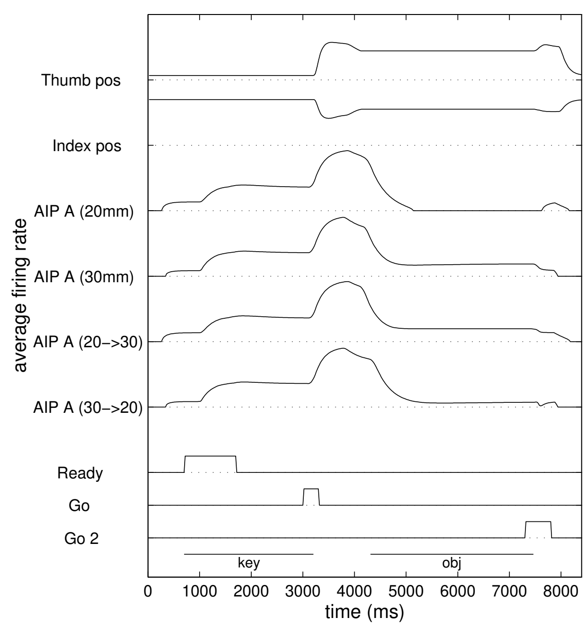

Figure 6.42: AIP visual/motor cell that is active in all conditions. Hold phase activity is most significant in cases in which a 30mm grasp is executed.

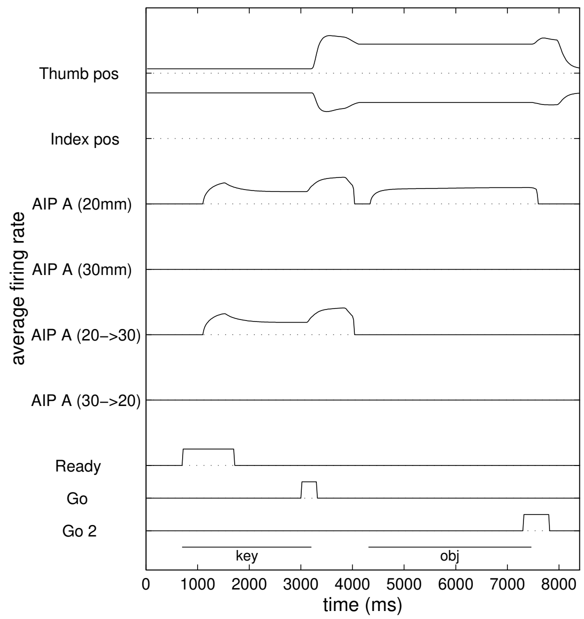

Figure 6.43: AIP motor-related cell that initially responds to the presentation of a 20mm stimulus, but only continues to be active through the hold phase in the unperturbed case.

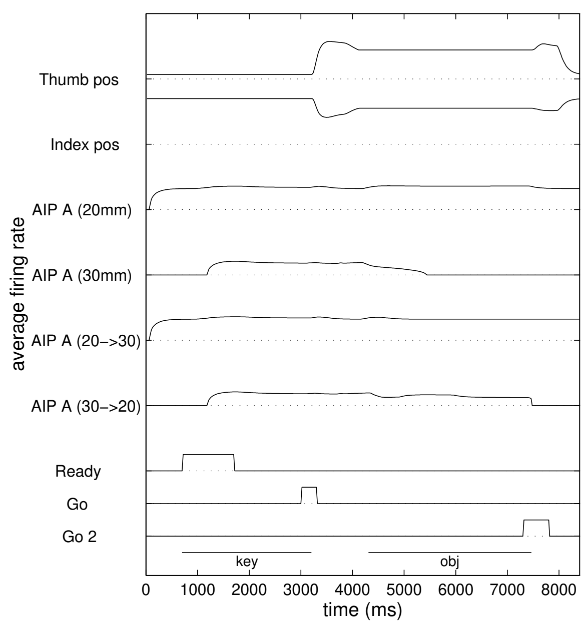

Figure 6.44: AIP (primarily) visual cell that is shunted during the hold phase only when a 30mm object is presented and manipulated.

Prediction #5: Object Parameter Coding

In looking carefully at object size-related activity described by (Murata, et al., 1993), one observes that during the experiments, when objects of different sizes were tested, the objects were scaled simultaneously in all dimensions. If we believe that AIP is really computing something like the affordances of the attended object, then we would expect that only the properties of the object that are relevant for physically interacting with it should be captured in AIP. For example, when objects are grasped in particular ways, certain dimensions of the objects are not as relevant as others to selecting the grasp aperture.

In this 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 (Figure 6.45). The two boxes of the leftmost column are to be grasped using a precision pinch of aperture 10mm, although they differ significantly in size in the other two dimensions. Likewise, the boxes of the rightmost column must be grasped using a precision pinch of aperture 20mm. The key question to be asked is how these four objects are coded by the AIP cell visual and motor responses.

Figure 6.46 A compares the visual-related population responses for the two objects in the leftmost column. All active AIP units demonstrate a significant level of activity for both objects, with little variation between the two. This is the case even though the two objects match in size along only the width dimension. However, comparing the responses towards the two objects in the bottom row of Figure 6.45, we see a much wider distribution of responses (with some cells responding exclusively to one object or the other) (panel B). These differences in response properties are due to the fact that in the model, objects represented in the

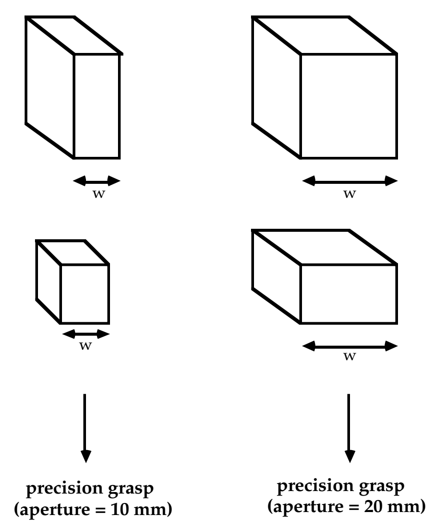

Figure 6.45: Four boxes of different dimensions.

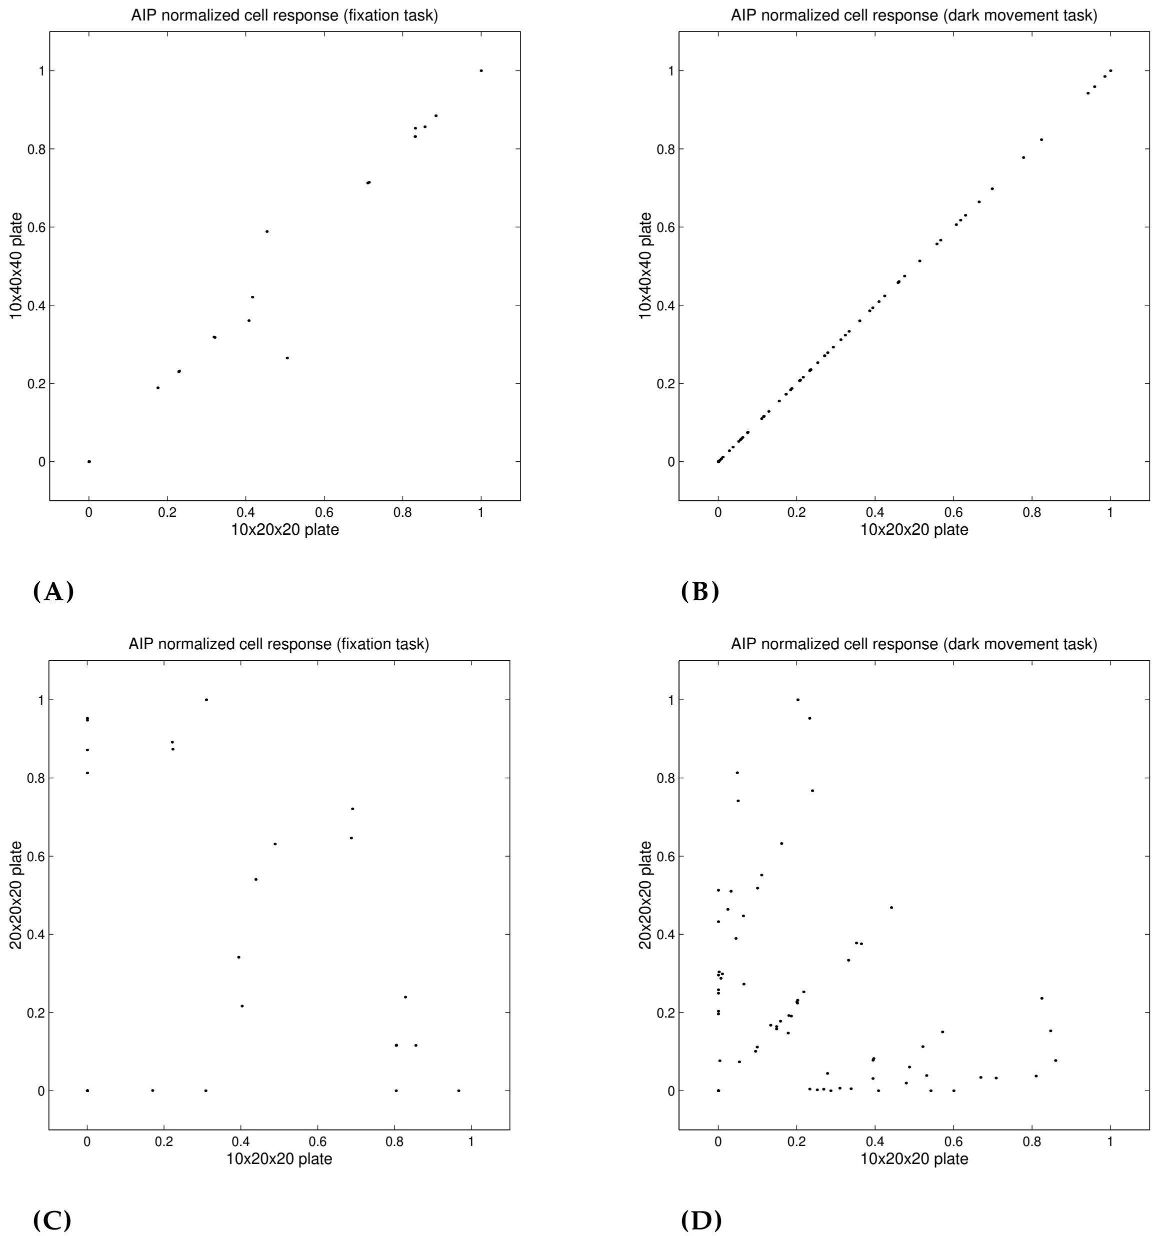

Figure 6.46: Comparison of AIP visual responses for objects of the same (A) and different (B) widths; and AIP motor-related responses (dark movement condition) for objects of same (C) and different (D) widths.

visual areas map to AIP cells that capture the critical parameter for physical interaction (the grasp aperture in this case).

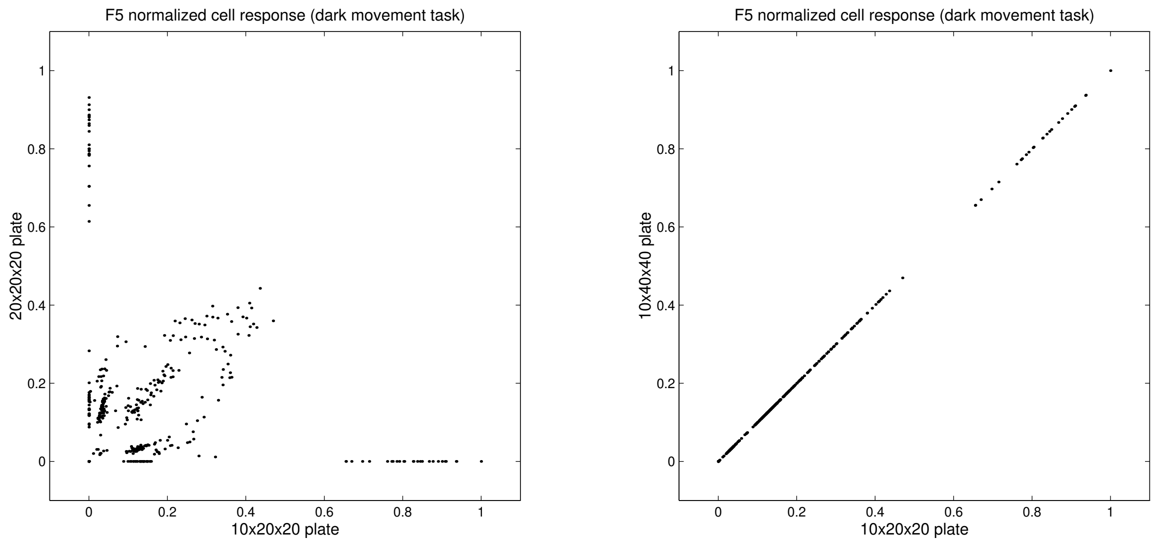

(A) **(B)**

Figure 6.47: F5 population responses for two objects of the same width as measured along the grasp axis (A), and two objects of different widths (B).

Figures 6.46 C, D demonstrate the motor-related AIP population responses for the same pair of objects. As was seen in an earlier section (Prediction #1), 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 (C); objects mapping to different grasps demonstrate grasp-specific modulation (D). In this case, however, many AIP motor-related cells are shared between the two objects. 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 due to the population code used to capture the aperture parameter. These latter responses compare closely with the F5 population responses of Figure 6.47.

Object versus Grasp Coding in the Monkey AIP

The monkey experiment reported in (Taira, et al., 1990) looked at AIP responses to a variety of objects and conditions. In one case, a pull-knob was presented to the monkey; the knob was mounted such that it protruded from a flat plane (fourth row in Figure 6.48). In general, the monkey preferred to grasp the knob using a side opposition. In a second case, an identical knob was presented, but was inset into a vertically cut slot (third row). In this case, the slot forced the monkey to adopt a precision pinch in order to successfully grasp the knob.

This task compares to some degree to the conditional experiment described earlier that was performed with the model (Prediction #2). In the monkey experiment, an object (the knob) affords two different grasps (the side opposition and the precision pinch). The selection of one of the two possibilities, however, is not done as a function of an abstract instruction stimulus, but rather as a function of task constraints (in this case, the reachability of the knob). Because the visual description of the knob is very similar in the two cases (discounting the representation of the slot), the model predicts that visual-related responses in AIP should be the same. However, because two different motor responses are generated, the model predicts that movement-related AIP firing patterns should reflect the type of grasp, and hence the condition that is presented to the monkey.

The AIP movement-related cell shown in Figure 6.48, which was recorded during the monkey experiments, demonstrates a selectivity (in the movement conditions) towards the pull-knob. However, the cell is very significantly active when a precision pinch is elicited, and is only slightly active during execution of the side opposition. On the other end of the spectrum, the visual-related cell of Figure 6.49 demonstrates only a small difference in firing rate during object fixation of the two pull-knobs. But, when movement is generated in the light, a significant difference emerges in the two cases.

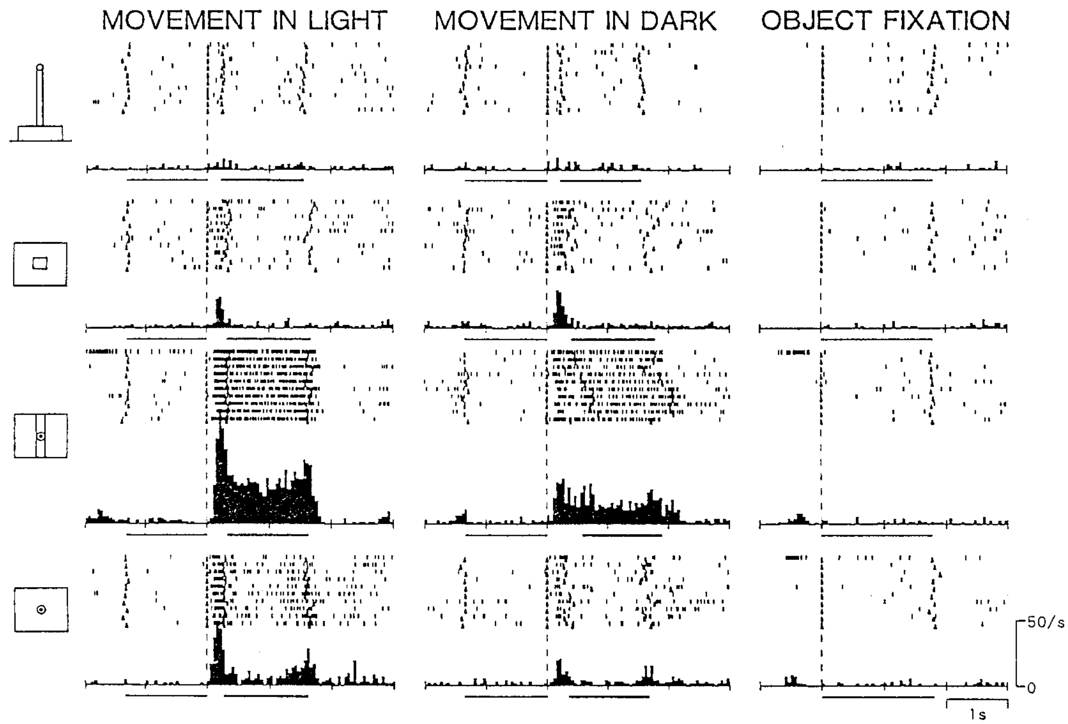

Figure 6.48: Response of a single monkey AIP cell to four objects (joystick, push-button, pull-knob in a slot, and a pull knob) and three conditions. The cell demonstrates motor-related behavior, and is most responsive when a precision pinch is used to grasp a knob. The experimental details are described in (Taira, et al., 1990); actual data from (Sakata, 1994).

Finally, the cell shown in Figure 6.50 is active in both fixation and dark movement conditions (and hence is considered a visual/motor cell). Although the cell demonstrates some difference in the firing rate for the fixation condition depending upon which knob is presented, the difference becomes much more significant when a movement is executed (either in the light or the dark).

In a more recent recording experiment, (Murata, et al., 1993) has looked at AIP responses to a set of different geometric primitives. In some cases, different primitives (e.g., a cylinder, and a rectangular block), which were approximately of the same dimensions, elicited the same type of grasp from the monkey. Under these conditions, the model predicts (Prediction #1) that visual-related responses can show alignment towards one or the other object. However, because

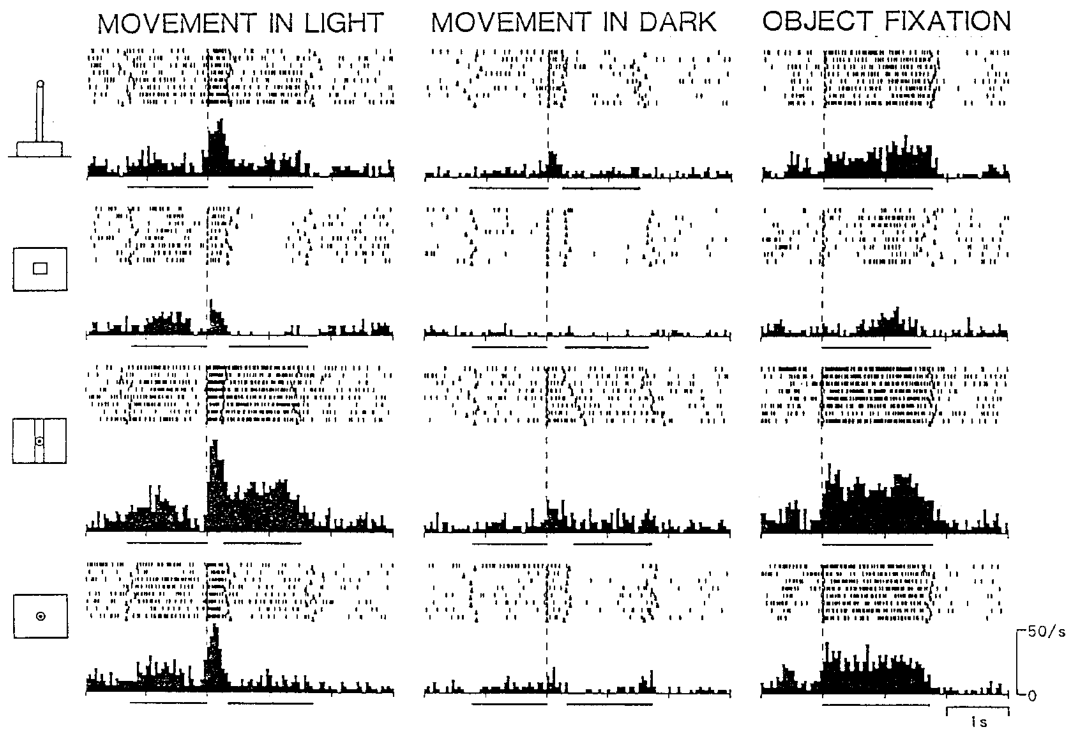

Figure 6.49: Response of a primarily visual-related AIP cell. Data from (Sakata, 1994).

the generated movement is identical in the two conditions, movement-related responses will lose the object specificity. Although the full analysis of this effect in the monkey data has not been performed, the few cells that have been studied in detail demonstrate the predicted behavior (Murata, 1995).

Summary

In this chapter, we have demonstrated the neural-level behavior of the model in a variety of experiments. We first illustrated the general behavior of the model and then made comparisons to single unit recordings performed in monkey F5 and AIP. Finally, we performed several novel experiments, making explicit predictions about expected single unit and population behavior. These experiments addressed the following key questions:

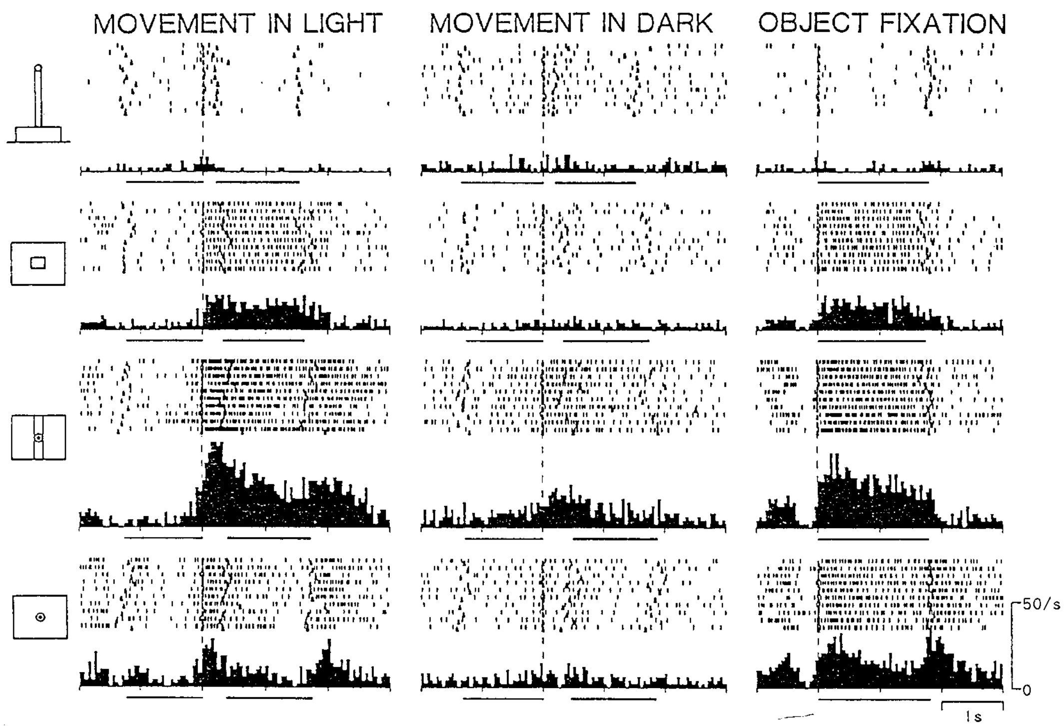

Figure 6.50: Response of a visual/motor AIP cell. From (Sakata, 1994).

1. What is the neural code in AIP that is used to represent different objects that map to the same grasp? The model predicts that an AIP cell's orientation (either object- or grasp-centered) is directly related to its visual/motor classification (Taira, et al., 1990). This arises from the fact that in the model visual-related responses are due to direct projections from the visual areas (PIP or IT), in which different objects can be coded using significantly different patterns of activity. The motor-related responses rely on recurrent projections from F5, which do not carry non-relevant object properties (from the point of view of the grasp).

2. What is the representation and conflict resolution process necessary for dealing with objects that can be grasped in more than one way? When the model is presented with an object of this type, and an instruction stimulus (IS) is used to determine which of the grasps should be used, then a) prior to the presentation of the IS, activity patterns in F5 and AIP reflect all

grasp/affordance possibilities; and b) after presentation of the IS, all but the selected grasp/affordance are shunted.

3. How are object/grasp parameters encoded in AIP and F5 (specifically object width and grasp aperture)? Aperture size is encoded in F5 and AIP (in the model) using sub-populations of cells (belonging to populations that are specific to the grasp/affordance pairs). The sub-populations encode the aperture parameter using a spatial code, in which each cell has a "preferred" aperture. Note that this effect is built directly into the model.

4. What is the behavior of the system when the object size is perturbed? When the apparent size of the object does not match the actual size, we see a shift in activity at the time of contact (first in F5, then in AIP) from the pattern that represents the grasp program for the expected object to the pattern of activity that represents the program for the object that was actually grasped.

5. Are all object parameters encoded with equal weight in AIP? A natural consequence of the wiring scheme is that when object parameters are explicitly relevant for programming the movement, the model predicts a higher concentration of cells that are involved in encoding that parameter. Parameters that are not relevant for physical interaction with the object will tend to not be captured in AIP.