

Are movement parameters recognizably coded in the activity of single neurons?

Eberhard E. Fetz

*Department of Physiology & Biophysics and Regional Primate Research
Center, University of Washington, Seattle, WA 98195*

Electronic mail: *fetz@locke.hs.washington.edu*

Abstract: To investigate neural mechanisms of movement, physiologists have analyzed the activity of task-related neurons in behaving animals. The relative onset latencies of neural activity have been scrutinized for evidence of a functional hierarchy of sequentially recruited centers, but experiments reveal that activity changes occur largely in parallel. Neurons whose activity covaries with movement parameters have been sought for evidence of explicit coding of parameters such as active force, limb displacement, and behavioral set. Neurons with recognizable relations to the task are typically selected from a larger population, ignoring those cells with complex relations to the task and unmodulated cells. Selective interpretations are also used to support the notion that different motor regions perform different motor functions; again, current evidence suggests that units with similar properties are distributed over widely different regions.

These coding issues are reexamined for premotoneuronal (PreM) cells, whose correlational links with motoneurons are revealed by spike-triggered averages. PreM cells are recruited over long times relative to their target muscles; they show diverse response patterns relative to the muscle force they produce; functionally disparate PreM cells such as afferent fibers and descending corticomotoneuronal and rubromotoneuronal cells can exhibit similar patterns. Neural mechanisms have been further elucidated by neural network simulations of sensorimotor behavior; the pre-output hidden units typically show diverse response patterns in relation to their target units. Thus, studies in which both the activity and the connectivity of the same units are known reveal that units with both simple and complex relations to the task contribute significantly to the output. This suggests that the search for explicit coding may be diverting us from understanding distributed neural mechanisms that operate without literal representations.

Keywords: chronic recording; motor cortex; movement parameters; neural coding; neural computation; neural networks; parallel distributed processing; premotoneuronal cells; representation; spike-triggered averages

1. Introduction

Many systems neurophysiologists record the activity of single units in behaving animals in the hope of understanding the neural mechanisms generating motor behavior. Such "chronic unit recording" experiments are typically designed to test a plausible hypothesis about the function of neurons at some recording site: The animal is trained to perform a behavioral task involving that function and the experimenter searches out relevant task-related cells. Over the last three decades this formula has generated numerous papers illustrating neurons whose activity appears to code (i.e., to covary with) various movement parameters or representations of higher-order sensorimotor functions. Initially, such studies seemed to provide supportive evidence for plausible notions, for example, that motor cortex cells code muscle force and that premotor cortex cells are related to programming movements. With an increasing number of more sophisticated studies it has become clear that the accumulating experimental evidence undermines many of our

simplistic notions about neural coding. Moreover, the search for neural correlates of motor parameters may actually distract us from recognizing the operation of radically different neural mechanisms of sensorimotor control.

This article begins with a review of experiments designed to show how various movement parameters may be represented in neural activity. This includes attempts to delineate a functional hierarchy of cells on the basis of their response latencies. We then consider studies of explicit coding of simple movement parameters such as active force and limb displacement and preparation to move. We discuss functional specialization in different cortical regions as well as the possibility that parameters are coded in populations of neurons. Since synaptic connections are an important determinant of the functional consequences of neural activity, we reexamine these coding questions for premotoneuronal cells, which have direct links with motoneurons. Finally, we reconsider these issues in light of results from neural network modeling studies.

2. Representation of movement parameters in neural activity

2.1. Relative timing of cell activity. To obtain evidence for a causal hierarchy of cells in different motor centers that mediate the programming and execution of movement, it first seemed reasonable to determine the sequential recruitment order of cells in different areas. A particularly useful behavioral paradigm for this purpose is the simple reaction-time response, in which an animal makes a repeatable movement in response to a stimulus such as light. The successive activation of neurons in different regions would then define a causal sequence of neurons mediating the transform between stimulus and response. For a visually triggered key release, for example, the sequence would begin with stimulation of retinal cells followed by propagation of activity to diverse cortical and subcortical centers, which might code the sensory aspects of the stimulus. The conversion of the stimulus-evoked activity into the preparation for movement might occur at intermediate times in cortical association areas. Finally, the neural activity involved in execution would converge in proper combination to activate agonist motoneurons that generate the movement. The peripheral links at the input and output stages of such a sequential scenario have been elucidated, but the central stages have consistently eluded temporal resolution.

The timing of motor cortex cells relative to movement was first studied by Evarts (1968), who showed that pyramidal tract neurons (PTNs) began to change their activity up to 100 msec before the onset of activity in agonist muscles. To determine the relative onset times of other cells that might precede activation of motor cortex neurons, Thach (1978) recorded neural activity in cerebellar nuclei, motor cortex, and muscles during the same responses. The onset times of activity changes of units in cerebellar nuclei were found to largely overlap those of precentral motor cortex cells (Figure 1). The onset times of different neurons in two cerebellar nuclei and motor cortex were distributed over hundreds of milliseconds, with a relatively slight difference in their mean onset times. Comparable overlap in recruitment times has been found in many subsequent experiments. Neurons in the supplementary motor area and primary motor cortex are recruited almost simultaneously in a reaction-time task (Chen et al. 1991) and during a step-tracking task (Alexander & Crutcher 1990c).

The basic problem in attempting to demonstrate serial activation of cells in different motor centers is that each region contains neurons that are recruited over diverse times. The extensive overlap in onset times makes it difficult to assign a sequential order of activation to different regions. Moreover, the duration of most movements as well as the duration of task-related activity greatly exceeds the conduction time between centers, so that recurrent loops could be "traversed" repeatedly during a single response. It is also relevant to note that the focus on the first change in neuronal activity puts undue emphasis on a subtle shift in firing rate that requires statistical determination. Functionally, the initial onset of a change has less to do with the cell's contribution to movement than its maximal activity. In any case, the appealing notion that initiation of movement involves the sequential activation of cells in hierarchically related


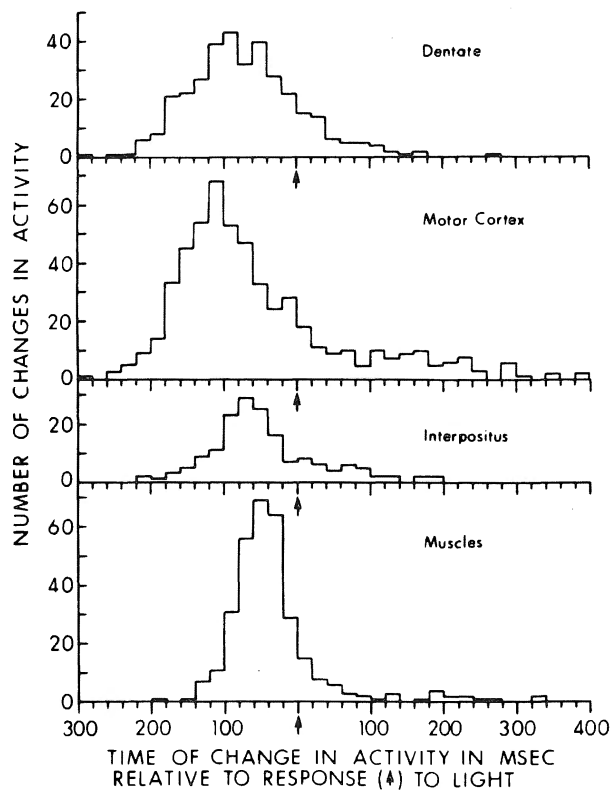


Figure 1. *Top*: Distribution of times of change in neural activity relative to light-triggered wrist movements (from Thach 1978; task is illustrated in Figure 2). *Bottom*: Distribution of onset times of CM cells relative to onset of activity in their target muscles (from Cheney & Fetz 1980).

centers is quite difficult to prove experimentally. In fact, the experimental results suggest that cells in diverse regions are activated more or less in parallel.

2.2. Coding of movement parameters. Although a cell's onset time provides equivocal evidence for its role in a causal hierarchy, its discharge pattern could provide a more robust indication of its contribution to movement. The hypothesis that parameters of movement are recognizably coded in the activity of motor system cells seems so reasonable that many experiments have been launched on the basis of this assumption. Neural coding, in the sense of covariation, has been amply investigated for a variety of movement parameters (reviewed in Evarts 1981; Fetz 1981; Fuster 1985). Since muscles are ultimately the generators of active force, it seems plausible that central cells controlling muscles could be coding the

force exerted during a movement. On the other hand, since we normally think in terms of moving a limb to particular target positions, it also seems reasonable that cortical cells could code the *displacement* or position of the limb. Evarts's first experiments to determine whether motor cortex neurons code force or displacement provide an excellent example of a behavioral paradigm designed to dissociate these variables. Evarts trained monkeys to make the same movements against different loads and, in some cases, to generate isometric activity without any displacement. In these studies, the activity of selected PTNs was related more to the active force or to the change of force than to displacement (Evarts 1968).

Yet a third variable to which cells could be related was found in monkeys prepared to make a movement: Some cortical cells changed activity long before an intended movement, suggesting that these cells may be involved in the *preparation* to make a movement, as contrasted with its execution, that is, with a behavioral *set*. Experiments designed to reveal set-related activity typically involve behavioral trials beginning with a sensory cue that indicates the correct movement, followed by a delay period and then a go signal to execute the cued motor response. During the delay between the cue and go signal, the monkey is prepared to initiate the movement and neurons in many cortical and subcortical regions exhibit associated changes in discharge.

Numerous other movement parameters have been suggested to be coded in neural activity, such as limb velocity (Gibson et al. 1985), direction of movement (Fortier et al. 1989; Georgopoulos et al. 1984; Schmidt et al. 1975), and target position (Alexander & Crutcher 1990c; Martin & Ghez 1985). The issue of neural coding can be discussed in relation to three parameters: active force, displacement, and behavioral set. Thach (1978) was the first to investigate all three variables systematically for motor cortex and cerebellar cells. He used a task (shown schematically in Figure 2, top) that involved each of these three parameters: The monkey moved the wrist through a sequence of successive hold positions against different loads. The lower trajectories schematize the expected activity patterns of cells primarily related to patterns of muscle force (MPAT), position of the wrist joint (JPOS), and preparation or set for the next direction of movement (DSET). Thach calculated the degree of correlation between these idealized patterns and the recorded activity of cells in motor cortex, cerebellum, and muscles and found that the degree of correlation for different cells was continuously distributed from weak to strong. As expected, many motor cortex cells showed the best correlation with the MPAT sequence. Many other cells in the motor cortex correlated more strongly with joint position and still others correlated with set. Perhaps the most remarkable finding was the relative numbers of cells in each category; as described by Thach, "In summary of the rather astonishing results on neural discharge patterns in motor cortex during holding, all the types of neuron that were looked for were found, in nearly equal numbers" (1978, p. 665). Proponents of coding of movement parameters can only interpret this result as indicating that the motor cortex contains a variety of cells, each coding a different parameter of the movement.

However, such a conclusion would have to be tempered by another remarkable finding in this study: A

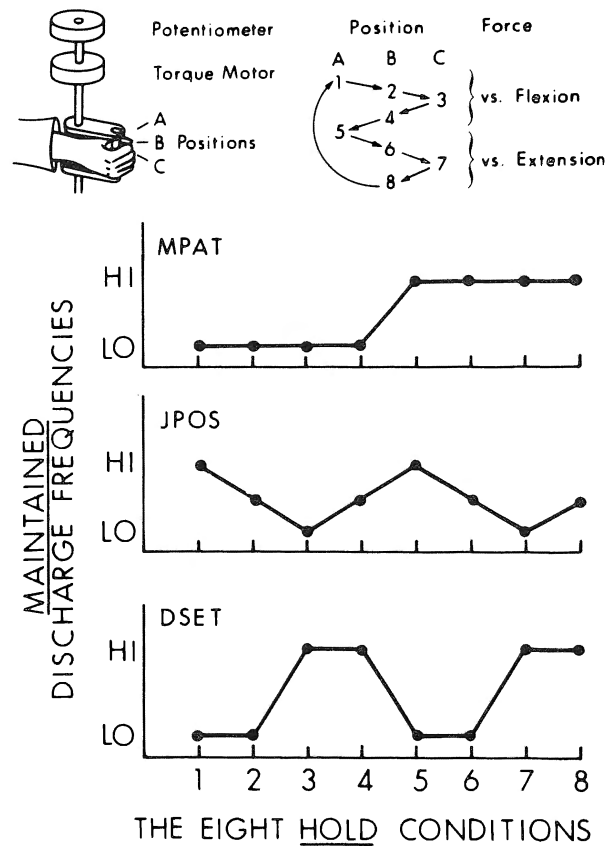


Figure 2. *Top*: Representation of a behavioral task used to test relation of cells to three parameters: pattern of muscle activity (MPAT), joint position (JPOS), and direction of intended next movement (DSET). *Bottom*: Trajectories show discharge frequencies of neurons or muscles having optimal relation to each parameter, plotted as a function of the hold positions illustrated at top (from Thach 1978).

slight change in response conditions could change the parameter that correlated best with a particular neuron. For example, the motor cortex cell illustrated in Figure 17 of Thach's paper had a strong relation to muscle patterns under condition of external load but was better related to joint position when the load was removed. Consistent relations between the activity of motor cortex cells and isometric muscle force have also been dissociated by changing the rewarded response patterns (Fetz & Finocchio 1975). Diehard proponents of coding would have to conclude from such flexible relationships that the same cells can "code" different parameters under different response conditions.

One basic problem with many attempts to relate neural activity to movement parameters is that the data are usually skewed by two types of experimental bias. One is the "task-induced bias" introduced by recording neural activity under particular behavioral conditions. In these experiments animals are performing a specific task designed to test the experimenter's hypothesis, and the activity of the modulated cells is interpreted in relation to that task. The data are further skewed by what could be called a "selection bias." The analysis of neural activity is typically confined to the subclass of neurons that show an interpretable relation to the task. Those cells that best

support the hypothesis are illustrated in the figures, and those that are statistically consistent are tabulated. However, in such studies two other groups of cells are invariably encountered: Many neurons are modulated with the task, but in complex ways that seem paradoxical or uninterpretable. In addition, many more neurons are simply unmodulated with the task. Although the latter two classes outnumber the interpretable task-modulated group, they are typically ignored. The paradoxical cells are rarely illustrated in papers, since they would detract from the main hypothesis, provoking the reviewers and confusing the readers. Instead, the uninterpretable and unmodulated cells are usually neglected in the final account of neural coding. Of course, many cells would in fact be marginally relevant to the ongoing task; however, by ignoring all the neurons with complex patterns we risk misunderstanding the real neural mechanisms in favor of dealing only with idealized and simplistic correlates.

This strong selection bias clearly undermines the contention that the functionally interpretable group of task-related cells provides convincing evidence for coding. Observations consistent with a given hypothesis can always be selected from a sufficiently large random data set. A rarely acknowledged fact of life in the neurophysiology laboratory is that neurons in many regions, including the motor cortex, exhibit an enormous variety of responses, a fact that provides an opportunity to find cells related to any given functional hypothesis. Thach (1978) eliminated this bias by objectively correlating the activity of the same population with three different, dissociable parameters. His finding that cells related to all three functions existed in nearly equal numbers suggests that something else may be going on besides preferential coding of particular movement parameters.

2.3. Localization of function. A common notion that is closely related to coding and also turns out to be simplistic in retrospect is the idea that different cortical areas are devoted to computing different motor functions. For example, it is commonly thought that the role of the motor cortex is to execute movement whereas motor association areas, such as the premotor and supplementary motor areas, are supposed to be concerned with motor programming or preparation to move under particular circumstances. Experiments designed to record neural activity in these regions under the appropriate behavioral conditions did indeed discover cells with the appropriate sorts of relationships. However, experiments in which neural activity in different regions was obtained under similar behavioral conditions have revealed that cells of the same type are found widely distributed over many areas. For example, neurons related to activation of muscles are found not only in the motor cortex but in the supplementary motor area (Chen et al. 1991; Crutcher & Alexander 1990c), premotor cortex (Godschalk et al. 1985; Weinrich & Wise 1982), prefrontal cortex (Fuster 1985; Niki & Watanabe 1976), and posterior parietal cortex (Mountcastle et al. 1975). Similarly, experiments involving delayed movements reveal set-related activity in the motor cortex (Tanji & Evarts 1976), premotor areas (Godschalk et al. 1985; Weinrich & Wise 1982), and prefrontal cortex (Fuster 1985; Niki & Watanabe 1976), as well as in the thalamus (Alexander & Fuster 1973) and basal ganglia (Alexander & Crutcher 1990c). Taken to-

gether, these studies suggest a very broadly distributed representation of these motor functions.

An extreme example of cortical specialization, which still remains almost axiomatic, is the presumed dichotomy between the functions of somatosensory and motor cortex. In this view, all precentral cells are thought to be involved in the execution of movement, whereas all postcentral cortex cells are assumed to be involved in somatosensory function. This view was challenged by Woolsey (1958), who noted that the maps of sensory input and motor output are similar and overlapping, in both precentral and postcentral gyri. Chronic unit recordings under active and passive conditions show that cells with similar response types can be found in both areas (e.g., Fetz et al. 1980; Soso & Fetz 1980). If this functional dichotomy is considered to be absolute rather than relative, identical response properties of single units must be interpreted in totally different functional terms. The responses of postcentral cells to passive stimulation are naturally interpreted as subserving somatic sensation, but the equally clear responses of precentral cells to passive joint movement and cutaneous stimulation are thought to subserve unconscious reflex functions. Similarly, the early responses of precentral cells preceding active limb movement are naturally thought to be involved in generating the movement, whereas identical early responses in postcentral cells are interpreted as subserving some sensory "corollary discharge."

The rationale for these diverse interpretations of identical response properties rests on functional presumptions derived in part from the effects of cortical stimulation. Stimulation thresholds for evoking movements are clearly lower in precentral than postcentral cortex (Woolsey 1958). And in conscious humans, cortical stimulation evokes somatic sensations from a larger number of postcentral than precentral sites (Penfield & Boldrey 1937); however, these differences are a matter of degree rather than absolute. In fact, similar effects can be evoked from both gyri, albeit at different thresholds. Nevertheless, the conceptual dichotomy between "sensory" and "motor" cortex is again preserved by applying a double standard to this experimental evidence. The somatic sensations evoked by stimulating precentral "motor" cortex in conscious humans are ascribed to a spread of activity to postcentral sites. The movements evoked by stimulating postcentral cortex are similarly ascribed to mediation via precentral cortex; reports that movements can be evoked from postcentral sites after ablating precentral cortex (Woolsey 1958) are even taken as evidence that the lesions were incomplete.

The assumed functional dichotomy of sensory and motor cortex is further based on their differing output projections. In the macaque the corticospinal axons from postcentral cortex terminate more dorsally in the spinal cord than axons of the precentral PTNs, although there is a good deal of overlap (Coulter & Jones 1977). The postcentral PTNs are undoubtedly more likely than precentral PTNs to affect afferently driven spinal cells, but their target region also contains cells involved in reflex circuitry, as well as dendrites of motoneurons. Perhaps more relevant to the function of individual cortical neurons than the output projections of the descending cells are the strong interconnections between pre- and postcentral cortex. These massive corticocortical connections

allow the cells in each region to participate in the functions of the other; indeed, these reciprocal connections would explain the similar response properties of neurons found in these areas.

Thus, the notion that cortical functions are segregated into different cortical areas can be preserved only by imposing different interpretations on similar experimental evidence. Units with the same response properties are imagined to code either sensory or motor parameters, depending on the presumed function of their recording sites. A plausible alternative is to consider the similar response properties of cells in different cortical regions as evidence that they are involved in performing similar functions; the neural substrate for these functions is then distributed correspondingly. This means that a given cortical region would be involved in diverse functions, consistent with the diverse cell types observed. This view provides a basis for distributed interactions between the functionally related sets of cells and helps explain the recovery of function after lesions. Note that this view does not claim equal involvement of all regions in all functions, since cortical areas are undoubtedly specialized. The point is that a region's specialized function need not be its only function, and certainly should not be the only standard for interpreting what each of the cells in this region is coding.

2.4. Population coding. Investigators have recently found that the activities of populations of cells can provide functions that match movement parameters more closely than the firing pattern of any single neuron. The fact that movements are ultimately produced by activity in large ensembles of neurons provides a clear rationale for population coding. Humphrey et al. (1970) first showed that the activities of multiple motor cortex neurons could be added together in the right proportion to match different parameters of wrist movement in an isotonic task. Weighted sums of the cells' firing rates could match the force trajectories and the wrist displacements, as well as their temporal derivatives, if the weighting factors for each cell could be optimally chosen for each parameter. Moreover, the match between the cells' weighted activities and the mechanical parameters improved with the number of cells included. The ability to freely optimize the weighting coefficients, of course, helped to ensure convergence on the movement trajectories; closer matches were obtained with larger populations because each additional nonredundant cell could serve to further reduce the remaining difference.

More recently, Georgopoulos et al. (1984) showed that populations of motor cortex cells could be used to match the direction of limb movement by invoking the "vector hypothesis" to sum the activity of directionally tuned cells. For a given movement direction, each cell was assumed to make a vector contribution pointing in the direction of its maximal activity, and by an amount proportional to the change in its overall mean rate during the given movement. The vector sum of all the unit vectors then approximated the direction of hand displacement. Again, the match improved as more cells with diverse directional preferences were included. This match with movement direction could be taken to suggest that the direction of hand displacement by the arm rather than muscle force is coded in motor cortex populations. The

direct match between the population function and arm displacement is appealing because it conveniently avoids the intervening complexities of synaptic connections and limb mechanics, which present formidable obstacles to a causal explanation. Moreover, the vector hypothesis is virtually guaranteed to work, given a sufficient distribution of cells. For a particular movement the cells whose best direction coincides with the movement will make the largest direct contribution; cells whose vectors point in the opposite direction will make a negative vector contribution, since average rates are subtracted, and therefore also contribute positively to the movement direction. The other cells have off-axis vector components that would tend to cancel with a sufficiently large population. Thus, the vector hypothesis will produce a match with movement direction whether the directionally "tuned" cells have any output effects on muscles or not. The same sorts of matches have been demonstrated for populations of posterior parietal area 5 neurons (Kalaska et al. 1983) as well as cerebellar cortical and nuclear cells (Fortier et al. 1989) and globus pallidus (Turner 1991).

Mussa-Ivaldi (1988) showed that the findings of Georgopoulos et al. (1984) would also result from a population of cells that code muscle shortening, and thus reconciled the apparent coding of limb displacement with the fact that many precentral cells do have effects on muscles. Recent studies by Kalaska et al. (1989) have shown that when the required force is varied independently of movement direction, the population vector of certain motor cortex cells shifts in the direction of active force. This result is consistent with a role in activating the agonist muscles. However, there are other motor cortex neurons whose population vector remains in the direction of movement, independent of force, much like posterior parietal cells (Kalaska et al. 1983). In this case, a key ingredient in making the matches with force or displacement is the ability to select the appropriate cells for each population.

Although one can find good descriptive matches between functions derived from the activity of neuronal populations and particular movement parameters, this correspondence is no proof of neural coding in the causal sense. To demonstrate that the candidate cells actually make a causal contribution requires additional evidence that they have appropriate output effects. Such a direct link is obviously difficult, and often impossible, for many central neurons. Still, a coding theory based merely on a descriptive match with a parameter provides no further basis for dealing with the neural interactions that would mediate the control of that parameter. A useful coding theory should provide some framework for understanding how the observed activity could contribute to the movement. For example, it would be helpful to know how the activity of the population whose "vector" points in the direction of movement is actually transformed into the movement. Descriptive correlates alone do not provide a causal framework for dealing with the underlying neural computation.

2.5. The coding problem. In retrospect, experiments designed to demonstrate coding of movement parameters have provided data that can be interpreted in either of two ways. Proponents of neural coding can point to the slight differences in mean onset latencies in different

regions as evidence of a sequential hierarchy of cells; they can point to examples of covariation of neural discharge with movement parameters as evidence of coding and they can ignore the complex and unrelated cells as unlikely to be involved; finally, they can consider different proportions of cell types in different areas as evidence of functional segregation. Alternatively, one could now argue that the accumulating experimental results have largely undermined these simplistic notions. The extensive overlap of activation times in different regions speaks more for parallel than for serial activation. Neural correlates of movement parameters in a particular task can always be selected from what is invariably a much larger variety of response types, but cells with more complex discharge could be just as involved in generating the movement, albeit in more complex ways. The distribution of similar cell types over diverse cortical fields speaks more for distributed representation than for functional segregation.

These issues cannot be resolved by more chronic unit recording data, because observing the activity of single and even multiple units is inherently insufficient to determine the mechanisms that generate movements. These studies usually lack another essential ingredient required to make causal inferences, namely, the connectivity between cells. In addition to the activation patterns generated during task performance, one must also know the output connections of the recorded cells in order to determine the consequences of that activity. The possible output connections are often inferred from independent anatomical evidence on major projections. However, such indirect inference is misleading for many neurons, since the cells encountered at a given recording site typically have diverse projections. If the cells' response properties are correlated with their projections, the functional distinctions described above could have been blurred by lumping them all together.

3. Response coding in premotoneuronal cells

To determine whether the variety of relationships observed in previous studies could be reduced by dealing with cells that directly affect motoneurons, some investigators have focused on those cells that have correlational linkages to motoneurons, as determined by spike-triggered averaging. These premotoneuronal (PreM) cells include the so-called corticomotoneuronal (CM) cells in precentral motor cortex (Buys et al. 1986; Cheney & Fetz 1980; Fetz and Cheney 1980; Lemon et al. 1986), the rubromotoneuronal (RM) cells in red nucleus (Cheney et al. 1988; Mewes 1988), and peripheral afferent fibers recorded in cervical dorsal root ganglia (DRG) (Flament et al. 1992). These PreM cells all produce short-latency postspike facilitation of EMG activity and have been documented in relation to comparable ramp-and-hold wrist movements – a response designed to elucidate the relation of cellular activity to changes in force and sustained force. It is interesting to consider the properties of PreM cells in the context of the four issues discussed above with regard to neural coding.

3.1. Timing. Although previous studies had shown that unidentified motor cortex cells exhibit a broad range of onset times relative to movement, it seemed possible that

CM cells would show a more restricted range of recruitment times relative to onset of activity in their target muscles. This turned out to be only partly true, as shown in Figure 1. The surprising result was that CM cells began to fire up to several hundred milliseconds before the onset of activity in their target muscles. Similar broad distributions of onset latencies have been observed for RM cells (Mewes 1988) and for afferent fibers in DRG (Flament et al. 1992). Since these PreM cells produce postspike facilitation of their target muscles in about 10 msec, the much earlier onset times are presumably related to bringing the motoneurons to threshold; those cells with reciprocal inhibitory linkages to antagonists of their target muscles would also contribute to turning the antagonist muscles off.

The inescapable conclusion is that even connected PreM-motoneuronal pairs are recruited relative to each other over times that straddle hundreds of milliseconds. Thus, connectivity is not a critical factor in restricting relative recruitment times; instead, there may be other relevant variables (if indeed there are any systematic explanations). Even within the same motoneuron pool, motoneurons are recruited in sequential order over extended periods of time. One variable that may be more relevant to recruitment order than the spatial location or the output connections of a neuron is its relative size. An increasing synaptic drive on motoneurons recruits the smallest motoneurons first and then successively larger ones with higher thresholds. Similar size relations may explain the timing of early and late recruited cells in the PreM population, a subject for future investigation. In any case, the distribution of onset times of PreM cells relative to their target muscles is almost two orders of magnitude broader than the latency of their postspike effects.

3.2. Coding of muscle force. The activity of PreM cells clearly has a direct output effect in facilitating their target muscles, which in turn generate active force. In this sense the PreM cells can be said to causally affect force. The relation of CM and RM cell discharge to active force has been confirmed by having monkeys generate different levels of isometric force. The tonic discharge rates of these cells during the static hold period are indeed proportional to active force over a range of torques, as shown in Figure 3 (Cheney & Fetz 1980; Cheney et al. 1985; Mewes 1988). In addition, many of these PreM cells show a phasic discharge at the onset of movement, which is preferentially related to change of force.

These observations pertain to the major subsets of the PreM cells, namely, those that show phasic-tonic or tonic discharge patterns during the ramp-and-hold movement. Figure 4 illustrates the basic response patterns of the three groups of PreM cells and single motor units during the ramp-and-hold movements. Of these patterns, only the tonic pattern is strictly proportional to the ramp-and-hold force trajectory. Other PreM cells show patterns that differ significantly from the active force and from the activation patterns of their target muscles. For example, the phasic-ramp CM cells show a strong burst of discharge at the onset of movement and a gradually increasing discharge during the static hold period. This pattern is totally different from the discharge of its target muscles.

There are also many PreM cells that show more com-

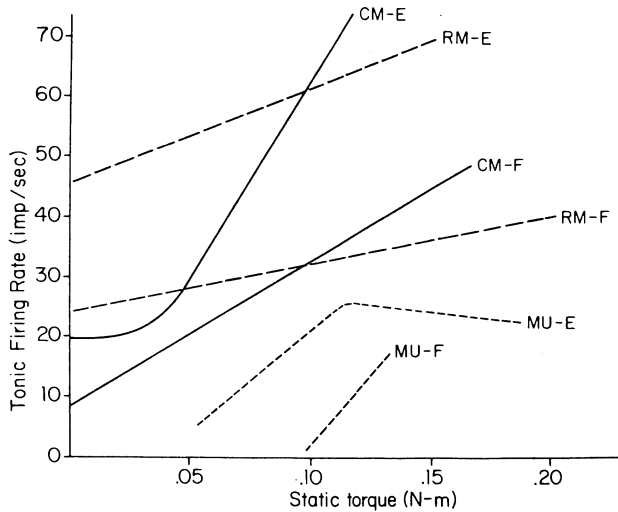


Figure 3. Tonic firing rate as a function of net static force about the wrist for representative CM and RM cells and single motor units. Rates are shown for flexor- and extensor-related cells (F and E, respectively; from Fetz et al. 1989).

plex and even counterintuitive response patterns. Many RM cells and some DRG cells exhibit bidirectional responses, firing during both flexion and extension, even though they facilitate only one set of agonist muscles. In motor cortex, some cells appear to have quite paradoxical relations to muscles: They covary with muscles in which they produce postspike suppression (see Figure 5 in Cheney et al. 1985).

Still another remarkable class of PreM cells discovered in the red nucleus are the unmodulated RM cells (Cheney et al. 1988; Mewes 1988). These increase their activity during wrist movement but are not modulated with alternating flexion and extension movements. Thus, they represent cells that essentially bias their target muscles during movement in both directions. These unmodulated cells, one should note, *are* causally involved in the active movement, as confirmed by their postspike facilitation of agonist muscles.




Figure 4. Response patterns of PreM cells and motor units during generation of ramp-and-hold wrist force (middle records). The proportion of units showing each pattern is tabulated for CM cells (from Cheney & Fetz 1980; Fetz et al. 1989), RM cells (from Mewes 1988), DRG cells producing postspike facilitation (from Flament et al. 1992), and single motor units (from Palmer & Fetz 1985).

Thus, PreM cells show three different relationships to their target muscles (as well as to force): Some PreM cells are simply coactivated with target muscles, others exhibit more complex and counterintuitive patterns, and still others are unmodulated. Despite these various relationships in discharge patterns, the correlational evidence confirms that they are all causally involved in activating their target muscles.

Furthermore, just as central cells can change their discharge patterns in relation to motor parameters under different conditions, single CM cells also can fire differently relative to their target muscles and force for different types of movements. When a monkey performed a finely controlled ramp-and-hold tracking task, the CM cells were strongly modulated with their agonist target muscles. When the monkey made rapidly alternating ballistic movements, however, the same cells were relatively inactive, even though their target muscles were more strongly activated (Cheney & Fetz 1980). Similarly, Muir and Lemon (1983) found that CM cells were preferentially active during a precision grip of a force transducer between thumb and forefinger, but the same CM cells were paradoxically inactive during a power grip, which involved even more intense activity in their target muscles. These results indicate an unexpected variability in the relation between even PreM cells and their target muscles under different movement conditions.

3.3. Localization of function. To determine whether supraspinal cells in motor cortex and red nucleus may have functional specializations that are different from those of afferent cells providing feedback from the periphery, one can compare the CM and RM populations with PreM cells recorded in the DRG. Surprisingly, the response patterns of the DRG units fell into the same categories as the most common supraspinal cells (tonic, phasic-tonic, and phasic; see Figure 4). Moreover, the relative onset times of many afferent cells also preceded the onset of their target muscle activity. This suggests that many PreM cells in radically different locations are recruited in similar ways. A similar result was obtained by Schieber and Thach (1985): During a slow ramp-and-hold tracking task they found similar classes of cells in DRG, motor cortex, and cerebellum.

In addition to PreM cells with similar response properties in all three locations, the supraspinal populations each included some unique types. As indicated in Figure 4, the ramp cells were observed only among cortical cells and the unmodulated neurons were found only in the RM group. This suggests that the three groups of PreM cells are not entirely equivalent but contain subsets of cells whose unique properties suggest some functional distinctions.

3.4. Population coding. Although individual PreM cells exhibited a variety of distinct discharge patterns, the net contribution of all the PreM cells to a target motoneuron would be more relevant to assessing their total effect. The response patterns of the PreM cells can be synthesized into a population average (Fetz et al. 1989). Since the cells were recorded under similar behavioral conditions, the average activity of the population can be obtained by summing the response histograms of individual cells (as well as their target muscles) aligned with the onsets of the

movements. This was done in stages, by first compiling subaverages for each response type and then adding these in proportion to the number of cells of each type. The resulting net ensemble averages of the discharge patterns of both the CM and RM population exhibited a phasic-tonic pattern. However, the motor cortex population showed a greater difference in the depth of modulation between opposite directions of active wrist force than the rubral cells, whose population histogram showed tonic activity during both directions of movement.

The net synaptic drive of the PreM cells on their target motoneurons would be proportional to these population histograms. The population histograms could also be used to infer the quantitative effect of the cells on their target muscles; the population activity can be multiplied by the correlational consequences of the postsynaptic potentials evoked from cortex and red nucleus (Fetz et al. 1989). The results provide a causal picture of the population influence on target motoneurons that is based on physiological measures of the synaptic linkages.

It is interesting to note that coding of muscle force in motor units requires a population average. Under normal conditions, single motor units code net muscle force in a highly nonlinear manner, since motor unit firing rates are limited at the lower end by their recruitment threshold and at the upper end by saturation (Figure 3 and Palmer & Fetz 1985). Moreover, the net force generated by the twitch tensions of a motor unit is a nonlinear function of its firing rate. This nonlinear behavior of the individual motor units is resolved by the population sum, which includes the successive recruitment of motor units with larger twitch tensions.

3.5. Implications of PreM cell properties for neural coding.

The properties of PreM cells have significant implications for the coding issue, insofar as their activity is causally related to generating muscle force, but this activity comes in a remarkable variety of discharge patterns. The connectivity of PreM cells to motoneurons is confirmed by cross-correlation methods, yet the response patterns of these PreM cells include all three types of relation observed between central cell activity and movement parameters. Many PreM cells clearly covary with the muscles that they facilitate, as one would intuitively expect. Others, such as the phasic-ramp CM cells and the bidirectionally activated RM cells, show counterintuitive discharge patterns that are distinctly different from the activity of their facilitated target muscles. Moreover, some cortical cells are paradoxically coactivated with arm muscles that they inhibit. In addition, a large group of unmodulated RM cells is tonically active during both phases of movement. This would indicate that the response patterns of neurons alone are not a reliable guide to their causal role in the task and that neural interactions between connected cells involve some highly nonlinear relationships. If the activities of connected PreM neurons and their target motoneurons can show such diverse relations, the chance of finding meaningful correlates of movement parameters would seem even more remote.

The same considerations apply to the relation of PreM cell discharge and the mechanical parameter of force, on which they clearly have a causal effect. Only the activity of the tonic PreM cells is directly proportional to force in this task. Indeed, the entire population of cortical and

rubral PreM cells exhibits a net phasic-tonic pattern, suggesting that force is coded in a nonlinear way even in the output cells that generate this force.

4. Computation of movement in neural networks

4.1. Holographic coding mechanisms. The basic reason that movement parameters need not be explicitly coded in the activity of single neurons is that movement is the consequence of large populations of interacting cells, which can generate an output without requiring any one cell to fire in proportion to the resultant movement parameters. Instead, the activity that is appropriate for a given cell is determined largely by its connections with the rest of the network rather than by any need to code an output parameter explicitly. This point can be illustrated by an apt analogy: the storage of images by holographic mechanisms. Holographic storage is based on a distributed representation of phase relations between wavefronts rather than a literal representation of the stored image. Recall that a holographic plate is constructed by exposing a photographic plate to the interference patterns between two coherent light beams – a reference beam obtained directly from the coherent source and an object beam reflected from the object whose image is to be stored. The spots on a holographic plate record the points of constructive interference, where the two light beams are in phase. These spots are distributed in a pattern that has no recognizable relation to the image. However, when the plate is illuminated with the reference beam, this distribution of spots forms a diffraction grating that reconstructs the object wavefront by the interference patterns in the transmitted beams.

The idea that neural networks may store and process information through holographic mechanisms is based on many salient analogies between the two systems. A small lesion in a holographic plate does not destroy any specific portion of the image but rather degrades the overall image quality; similarly, small lesions in the nervous system typically produce subtle behavioral deficits at most. The association between images of two spatially adjacent objects can be readily demonstrated by creating a hologram from the light reflected from the two objects; illuminating the developed plate with the light reflected from only one of the objects will reproduce a ghost image of the missing one. In this case the light from the remaining object essentially acts as the reference beam for reconstructing the other. This mechanism provides an analogue of associative memory and a model of content-addressable memory (Hinton & Anderson 1981; Pribram et al. 1974). Such a mechanism is likely to be involved in perceptual processes such as figure completion. The ability to execute a skilled movement sequence in a particular context may well involve similar associations between changing sensory inputs and central programs.

The basis for these analogous properties is the distributed representation of the information, using constructive interference between activity propagated in parallel pathways. Activity in a neural network is also propagated by the coincident arrival of sufficient synaptic input to activate the relaying neurons. With regard to coding mechanisms, the relevant point is that the spots on the hologram do not form a literal pictorial representation

of the image; instead, the “meaning” of each spot depends on its relation to the rest of the hologram – each point diffracts light in such a way that the net interaction with adjacent points reconstructs the wavefront of the stored image. Similarly, in the nervous system, the activity of a cell need not form a literal representation of a movement parameter; instead, its contribution to movement depends on its diverse connections and interaction with the rest of the network.

Optical holograms clearly represent highly simplified examples of this type of distributed nonliteral coding mechanism, insofar as they store only static images. Nevertheless, the same principles apply to storage and retrieval of dynamic information in neural networks (see papers in Hinton & Anderson 1981). Such analogies between neural and holographic mechanisms have been largely speculative until now. With the advent of neural network modeling, it has become possible to demonstrate these same properties in simulated populations of cells.

4.2. Neural network models. Model networks can be used to simulate the mechanisms operating in populations of cells; they also have unique heuristic value in elucidating the principles of neural computation. The behavior of ensembles of neurons is difficult if not impossible to synthesize by “bottom-up” inferences from single-unit recordings alone, mainly because the relevant connections between the recorded neurons are typically unknown. In contrast, model networks that simulate a particular behavior can be obtained by “top-down” derivations based on examples of the behavior, using training algorithms such as back-propagated error correction (Rumelhart et al. 1986b) or trial-and-error learning (Kuperstein 1988). The resulting dynamic networks can simulate motor activity without explicitly representing movement parameters in the activity of particular units.

For example, to determine what sort of neural network might be able to transform the step change in target position that a monkey sees into the response patterns generated by his agonist motor units, we used back propagation to derive the appropriate dynamic recurrent networks (Fetz & Shupe 1990; Fetz et al. 1990). The input and output layers were connected to intervening excitatory and inhibitory hidden units, as shown by the schematic diagram in Figure 5; an example of a specific weight matrix is shown in Figure 6. Initially, the synaptic weights were assigned randomly; presenting the input produced an initial output that deviated drastically from the desired target output. The difference between the actual output and the desired output was used to change the weights appropriately to reduce the error between actual and target outputs. Successive training iterations produced a network that transformed the temporal input patterns (step and transient inputs for flexion and extension) to the desired output patterns (eight motor unit patterns: tonic, phasic-tonic, decrementing, and phasic, for both flexion and extension). One resulting network is shown in Figure 6; the size of each square represents the strength of connection from the unit identified at the left to the unit at the top. The activation patterns for a flexion-extension cycle are illustrated for each unit. This result represents a complete neural network solution for this simplified sensorimotor transform in that the activity patterns of the




Figure 5. Schematic of dynamic recurrent network that simulates the step tracking task. Input functions represent step and transient signals for flexion and extension; target output patterns represent the discharge patterns of flexor and extensor motor units. The types of recurrent connections implemented are shown schematically. To derive a network that performs the transform, the output activations of an incompletely trained network $[N(t)]$ are subtracted from the target outputs $[T(t)]$ and the weights between units $[w_{ij}]$ are modified to reduce this error.

sensorimotor transform in that the activity patterns of the intervening hidden cells, as well as their connectivity, are completely specified.

These network solutions can be analyzed systematically to determine how the output patterns are derived. Relevant to the issue of “coding” we can examine how the response patterns of the output units are represented in the activity of the hidden units. For example, to see how the network in Figure 6 produced activity of the phasic flexor output unit (fp), one can examine its synaptic inputs (represented by the vertical column of weights under fp). The strongest weights indicate that phasic flexion was derived by two different means. As proponents of explicit coding might predict, the phasic output cell had strong excitatory connections from phasically active hidden units (e.g., a1, a8). A second contribution, however, came from excitatory units with tonic activity (a11), in conjunction with a delayed tonic input from inhibitory hidden units (b1). The difference between these two also contributes to the phasic output. Yet a third mechanism has been observed in other network simulations, which were allowed to have tonic biases on the cells. In those cases, the phasic output could also be derived from the sum of excitatory input from a phasic-tonic hidden unit in combination with a negative bias that essentially subtracted the tonic component. Thus, a pertinent lesson from these simulations is that many combinations of hidden unit activity can and do contribute to the same output response pattern.

It is interesting to note that many properties of the hidden units in these networks are analogous to those found in cells in the nervous system. For example, a given hidden unit (e.g., a11) may have divergent excitatory connections to many different types of output cells, just as CM cells facilitate motor units of different response types. Conversely, a given output unit typically receives convergent input from many hidden units, with different activations. Nevertheless, the connections are not equally distributed; this simulation produced preferen-




Figure 6. Neural network transforming step and transient inputs to firing patterns of motor units at output. Unit activations during a flexion-extension cycle are shown at left and along the top. The weight matrix gives the strength of connections from the row to the column units (weight scale at top). The rows represent, from top to bottom, the bias (which was eliminated for this simulation), the inputs (fs to eb), the excitatory hidden units (a's), inhibitory hidden units (b's), and the output flexor and extensor motor units (ft to ep). The target output patterns for both flexion and extension are tonic (ft & et), phasic-tonic (fpt & ept), decrementing (fd & ed), and phasic (fp & ep). To better visualize the relationships between units, the hidden units were sorted in order of the strength of their contribution to the phasic-tonic output units (from Fetz & Shupe 1990).

tially strong connections within the sets of units with sustained and transient activity. For example, the first two flexor hidden units exhibiting tonic activity (a11 & a2) are strongly interconnected; they receive potent input from the flexion step and connect strongly to the tonic output unit. Similarly, the brief flexion input (fb) is most strongly connected to the phasic hidden units (a1 & a8), which are strongly interconnected and which also have strong reciprocal connections with the phasic flexion output unit (fp). Although units with similar response patterns tend to be more strongly interconnected, there are also significant connections between units with quite dissimilar responses; this is even more pronounced in other simulations.

Relevant to the coding issue, one significant result of these simulations is the demonstration that a large number of network solutions can produce the same transform. Starting a given network architecture with different initial weights will usually produce solutions with a different set of final weights and activations. Even within the same network, one can discern a variety of solutions: In Figure 6, the flexion and extension phases involve essentially identical inputs and outputs, but the network utilizes different types of hidden units devoted to each. There is every reason to suspect that biological nervous systems can also utilize a variety of stratagems to perform a given behavior. The variance of experimental data from one animal to the next may be due to the fact that different animals could develop different neural computations to generate the same behavior. In recording data from multiple animals, experimenters typically assume that each animal performs the task using the same network solution; indeed, this assumption is a prerequisite for pooling data recorded from different animals. Network simulations suggest that many different neural network solutions can mediate the same behavior.

In fact, modeling experiments indicate that even in the same network, several different strategies for generating a response pattern can be implemented simultaneously, as discussed above for the phasic output. In this light, the commonly observed variance in a neuron's discharge pattern from trial to trial could well represent a variance in the degree to which different solutions are implemented in each trial. Thus, pooling data from different animals, and perhaps even from different trials in the same animal, would yield some average hybrid of different specific solutions. The common criticism that neural network models do not provide the same solution as biological networks is predicated on the debatable assumption that there is only one "real" biological solution (and, moreover, that it can be found by present experimental methods).

The coding issues that have been discussed above in relation to single-unit recording studies can be reexamined in light of the neural network simulations. In many dynamic network solutions, the *relative timing* of onsets of hidden unit activations can be widely distributed with regard to a given output response (Fetz et al. 1990). This is true even for hidden units that contact the output unit directly. This staggered timing is related to the build-up of recurrent activity in interconnected units; it does not represent sequential, hierarchical stages of processing.

The *response patterns* of hidden units that contribute to the output also show all three types of relation to the output. Some hidden units simply covary with the output unit that they excite, as might be expected intuitively (e.g., a1 and fp). Many other units show activity patterns that differ significantly from their target outputs (e.g., a11 & a10 compared with fp). In some simulations of reciprocal movement, many hidden units have bidirectional responses during both flexion and extension; the inappropriate portion of their activity is simply eliminated by inhibitory units. In addition, the inhibitory hidden units frequently show counterintuitive coactivation with cells they inhibit (e.g., b10), as has been seen in some cortical neurons. Finally, the tonic bias units used by many of these networks are clearly analogous to the unmodulated activity seen in many RM cells. Thus, the activity of output units is not necessarily coded recognizably in the activity of hidden units, even those that provide direct input. Not only do the network simulations reveal all three types of relations between hidden units and output units, but experiments with network lesions (Fetz & Shupe 1990) confirm that each type makes a significant contribution to the output.

The issue of *localization of function* can also be seen in a new light with network simulations. The functional consequence of the activity of any particular hidden unit is determined by its connectivity in the network; its physical location would be entirely arbitrary. Thus, if the hidden units were physically implemented, they could be reorganized in space without affecting the network computation, so long as their connectivity remained intact.

Relevant to functional localization, a common property of representation in cortical fields is the tendency to form topographic maps of the peripheral receptors or muscles. This feature of cortical organization has been simulated in neural network models; using local Hebbian rules to change synaptic strengths will lead to topologically organized feature maps (Kohonen 1982). This type of topographical organization within a cortical field should be distinguished from the segregation of functional computation among different fields. As demonstrated by network models, topographic organization can result from local synaptic interactions; in biological networks this may also have some wiring convenience. In contrast, functional segregation in the form of explicit separation of computational stages does not appear in network simulations.

The issue of *population coding* is also illuminated by these network simulations. The response pattern of any particular output unit is simply derived from the computed sum of all its inputs, weighted by the connection strengths. There is no need for explicit coding of any other sort. One could imagine taking the activity of a population of hidden units and matching some movement parameter by an appropriately weighted sum of their activities. Despite the success of such a mathematical exercise, the weights that are actually significant for the neural calculations are the synaptic links between units, not the mathematical coefficients required to calculate an optimal match. Put another way, the ability to obtain a population function that matches a parameter is quite irrelevant to the neural mechanisms that generate the output.

Clearly, these initial network simulations are still too simplistic in their connectivity and cell properties to be

taken as realistic models; nevertheless, they serve to illustrate some of the mechanisms at work in large populations of units interacting in ways analogous to neuronal interaction. Thus, network models provide a useful heuristic tool for investigating network mechanisms and can help to bridge the impasse between single-unit data and behavior. In the future, these network simulations can be improved to provide more realistic models of biological networks by incorporating the activity of more cells recorded in behavioral experiments and by making the connections more appropriate.

5. Concluding comments

We have taken the devil's advocate position on the notion that movement parameters are explicitly "coded" in neural activity. If "coding" is defined simply as covariation with movement parameters, the nervous system will provide ample opportunities to search out cells whose activity correlates with this or that parameter. Given the variety of neural discharge patterns and the ability to select the best examples, one can anticipate further examples of cells that could code some hypothesized variable. Like reading tea leaves, this approach can be used to create an impression, by projecting conceptual schemes onto suggestive patterns. This selective approach ignores two major groups of neurons: those with a complex or paradoxical relation to the task and those that are not modulated. It seems significant that studies in which both the activity and the connectivity of the same neurons are

known – namely, physiological studies using spike-triggered averaging and modeling studies with neural network simulations – reveal that all three classes of units can and do contribute significantly to the output. Thus, the search for explicit coding may actually be misleading, and may divert our understanding of distributed neural mechanisms that operate without literal representations.

If virtually any neuron can potentially contribute to generation of movement, how can we ever hope to understand the underlying mechanisms? Ultimately, systems neurophysiologists can profitably use a combination of single-unit recording techniques and neural modeling to investigate the network mechanisms generating motor behavior. Unit recordings can provide important constraints on the activity of related neurons, but the network models can provide working examples of complete solutions to sensorimotor behavior. To the extent that models can incorporate anatomical and physiological constraints, they can provide plausible explanations of the mechanisms of neural computation.

ACKNOWLEDGMENTS

I thank Paul Cordo for organizing the "Controversies in Neuroscience: I" meeting, which provided a unique opportunity to present provocative perspectives, and the NSF for supporting this meeting. I thank Kate Elias for editorial help and various colleagues for helpful suggestions. This work was supported in part by NIH grants RR00166 and NS12542 and by ONR contract N00018-89-J-1240.

6. References

- Alexander, G. E. & Crutcher, M. D. (1990) Neural representations of the target (goal) of visually guided arm movements in three motor areas of the monkey. *Journal of Neurophysiology* 64:164-178.
- Alexander, G. E. & Fuster, J. (1973) Effects of coding prefrontal cortex on cell firing in the nucleus medialis dorsalis. *Brain Research* 61:93-105.
- Buys, E. R., Lemon, R. N., Mantel, G. W. B. & Muir, R. B. (1986) Selective facilitation of different hand muscles by single corticospinal neurones in the conscious monkey. *Journal of Physiology (London)* 381:529-549.
- Chen, D.-f., Hyland, B., Maier, V., Palmeri, A. & Wiesendanger, M. (1991) Comparison of neural activity in the supplementary motor area and the primary motor cortex in monkeys. *Somatosensory and Motor Research* 8: 27-44.
- Cheney, P. D. & Fetz E. E. (1980) Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of Neurophysiology* 44:773-791.
- Cheney, P. D., Fetz, E. E. & Palmer, S. S. (1985) Patterns of facilitation and suppression of antagonist forelimb muscles from motor cortex sites in the awake monkey. *Journal of Neurophysiology* 53:805-820.
- Cheney, P. D., Mewes, K. & Fetz, E. E. (1988) Encoding of motor parameters by corticomotoneuronal (CM) and rubromotoneuronal (RM) cells identified by spike-triggered averaging in the awake monkey. *Behavioral Brain Research* 28:181-191.
- Coulter, J.D. & Jones, E.G. (1977) Differential distribution of corticospinal projections from individual cytoarchitectonic fields in the monkey. *Brain Research* 129: 335-340.
- Crutcher, M. D. & Alexander, G. E. (1990) Movement-related neuronal activity selectively coding either direction or muscle patterns in three motor areas of the monkey. *Journal of Neurophysiology* 64:164-178.
- Evarts, E. V. (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology* 31:14-27.
- Evarts, E. V. (1981) Role of motor cortex in voluntary movement In: *Handbook of physiology*, Vol. II, pp. 1083-1121. American Physiological Society, Waverly Press.
- Fetz, E. E. (1981) Neuronal activity associated with conditioned limb movements. In: *Handbook of behavioral neurobiology*, Vol. II: *Motor coordination*, ed. A. L. Towé & E. S. Luschei. Plenum Press, pp. 493-526.
- Fetz, E. E. & Cheney, P. D. (1980) Postspike facilitation of forelimb muscle activity by primate corticomotoneuronal cells. *Journal of Neurophysiology* 44:751-772.
- Fetz, E. E., Cheney, P. D., Mewes, K. & Palmer, S. (1989) Control of forelimb muscle activity by populations of corticomotoneuronal and rubromotoneuronal cells. In: *Peripheral control of posture and locomotion*, ed. J. A. H. Allum and M. Hulliger. Elsevier (*Progress in Brain Research* 80:437-449).
- Fetz, E. E. & Finocchio, D. V. (1975) Correlations between activity of motor cortex cells and arm muscles during operantly conditioned response patterns. *Experimental Brain Research* 23:217-240.
- Fetz, E. E., Finocchio, D. V., Baker, M. A. & Soso, M. J. (1980) Sensory and motor responses of precentral cortex cells during comparable passive and active joint movements. *Journal of Neurophysiology* 43:1070-1089.
- Fetz, E. E. Shupe, L. E. & Murthy, V. (1990) Neural networks controlling wrist movements. *Proceedings of 1990 International Joint Conference on Neural Networks II:675-679.*
- Fetz, E. E. & Shupe, L. E. (1990) Neural network models of the primate motor system. In: *Advanced neural computers*, ed. R. Eckmiller. Elsevier/ North Holland, pp. 43-50.
- Flament, D., Fortier, P. A. & Fetz, E. E. (1992) Response patterns and post-spike effects of peripheral afferents in dorsal root ganglia of behaving monkeys. *Journal of Neurophysiology*, 67: 875-889.
- Fortier, P. A. Kalaska, J. F. & Smith, A. M. (1989) Cerebellar neuronal activity related to whole-arm reaching movements in the monkey. *Journal of Neurophysiology* 62:198-211.
- Fuster, J. (1985) The prefrontal cortex and temporal integration. In: *The cerebral cortex*, vol. 4, ed. A. Peters & E. G. Jones, Plenum Press, N.Y. pp. 151-177.
- Georgopoulos, A. P., Kalaska, J. F., Crutcher, M. D., Caminiti, R. & Massey, J. T. (1984) The representation of movement direction in the motor cortex: single cell and population studies. In: *Dynamic aspects of neocortical function*, ed. G. E. Edelman, W. E. Gall & W. M. Cowan. John Wiley, pp. 501-524.
- Gibson, A. R., Houk, J. C. & Kohlerman, N. J. (1985) Relation between red nucleus discharge and movement parameters in trained macaque monkeys. *Journal of Physiology (London)* 358:551-570.
- Godschalk M., Lemon, R. N., Kuypers, H. G. J. M. & Van der Steen, J. (1985) The involvement of monkey premotor cortex neurons in preparation of visually cued arm movements. *Behavioral Brain Research* 18:143-157.
- Hinton, G. E. & Anderson, J. A. (1981) *Parallel models of associative memory*. Lawrence Erlbaum Associates, Hillsdale, N.J.
- Humphrey, D. R., Schmidt, E. M. & Thompson, W. D. (1970) Predicting measures of motor performance from multiple spike trains. *Science* 170: 758-762.
- Kalaska, J. F., Caminiti, R. & Georgopoulos, A. P. (1983) Cortical mechanisms related to the direction of two-dimensional arm movements: relations in parietal area 5 and comparison with motor cortex. *Brain Research* 51: 247-260.
- Kalaska, J. F., Cohen, D. A., Hyde, M. L. & Prud'homme, M. (1989) A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *Journal of Neuroscience* 9:2080-2102.
- Kohonen, T. (1982) Self-organized formation of topologically correct feature maps. *Biological Cybernetics* 43:59-69.
- Kuperstein, M. (1988) Neural model of adaptive eye-hand coordination for single postures. *Science* 239:1308-1311.
- Lemon, R. N., Mantel, G. W. H. & Muir, R. B. (1986) Corticospinal facilitation of hand muscles during voluntary movement in the conscious monkey. *Journal of Physiology (London)* 381:497-527.

- Martin, J. H. & Ghez, C. (1985) Task-related coding of stimulus and response in cat motor cortex. *Experimental Brain Research* 57:427-442.
- Mewes, K. (1988) Characteristics of rubromotoneuronal cells and their role in the control of the hand in primates. Ph.D. Thesis, University of Kansas.
- Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H. & Acuna, C. (1975) Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *Journal of Neurophysiology* 38:871-908.
- Muir, R. B. & Lemon, R. N. (1983) Corticospinal neurons with a special role in precision grip. *Brain Research* 261:312-316.
- Mussa-Ivaldi, F. A. (1988) Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neuroscience Letters* 91:106-111.
- Niki, H. & Watanabe, M. (1976) Prefrontal unit activity and delayed response: relation to cue location versus direction of response. *Brain Research* 105:79-88.
- Palmer, S. S. & Fetz, E. E. (1985) Discharge properties of primate forearm motor units during isometric muscle activity. *Journal of Neurophysiology* 54:1178-1193.
- Penfield, W. & Boldrey, E. H. (1937) Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain* 60:389-443.
- Pribram, K.H., Nuwer, M. & Baron, R. (1974) The holographic hypothesis of memory structure in brain function and perception. In: *Contemporary developments in mathematical psychology 2*, ed. D. Krantz, R.C. Atkinson, R. Luce & P. Suppes, Freeman Press, San Francisco, CA.
- Rumelhart, D. E., Hinton, G. E. & Williams, R. J. (1986) Learning internal representations by error propagation. In: *Parallel distributed processing*, vol. 1, ed. D. E. Rumelhart & J. L. McClelland. MIT Press, Cambridge, MA, pp. 318-362.
- Schieber M. A. & Thach, W. T. (1985) Trained slow tracking. II. Bidirectional discharge patterns of cerebellar nuclear, motor cortex, and spindle afferent neurons. *Journal of Neurophysiology* 55:1228-1270.
- Schmidt, E. M., Jost, R. G. & Davis, K. K. (1975) Reexamination of the force relationship of cortical cell discharge patterns with conditioned wrist movements. *Brain Research* 83:213-223.
- Soso, M. J. & Fetz, E. E. (1980) Responses of identified cells in post-central cortex of awake monkeys during comparable active and passive joint movements. *Journal of Neurophysiology* 43:1090-1110.
- Tanji, J. & Evarts, E. V. (1976) Anticipatory activity in motor cortex neurons in relation to direction of intended movement. *Journal of Neurophysiology* 39:1062-1068.
- Thach, W. T. (1978) Correlation of neural discharge with pattern and force of muscular activity, joint position and direction of intended next movement in motor cortex and cerebellum. *Journal of Neurophysiology* 41:650-676.
- Turner, R. S. (1991) Movement and instruction related activity in the globus pallidus of the monkey. Ph.D. Thesis, University of Washington.
- Weinrich, M. & Wise, S. P. (1982) The premotor cortex of the monkey. *Journal of Neuroscience* 2:1329-1345.
- Woolsey, C. N. (1958) Organization of somatic sensory and motor areas of the cerebral cortex. In: *Biological and biochemical bases of behavior*, ed. H. F. Harlow & C. N. Woolsey. University of Wisconsin Press, pp. 63-81.