Cortical Populations and Behaviour: Hebb's Thread

BAGRAT AMIRIKIAN and APOSTOLOS P. GEORGOPOULOS,
Veterans Affairs Medical Center and University of Minnesota Medical School

Abstract This paper discusses work on the function of the motor cortex as revealed by single cell recordings in monkeys and artificial neural network modelling. Our key conceptual approach both in behavioural neuroscience and neural network modeling of motor cortical function relies on reconstructing, visualizing, and modelling the activity in *neuronal populations*, indeed a key concept advanced by Hebb (1949). The behaviour investigated ranges from exertion of isometric force to pointing movements to complex cognitive processing. The functional properties of single cells with respect to the direction of movement in space are described as well as a population code which provides a unique measure for this direction. Finally, the results of modeling studies are discussed in which directional population activity is used as an input to an artificial neural network to drive a simulated arm.

The motor cortex occupies a crucial node in brain control of motor function. It provides major outputs to the spinal cord and brainstem, and is heavily interconnected with other cortical areas as well as the cerebellum and the basal ganglia. The question of how the motor cortex performs its function has been intensely investigated during the past century but it is only during the last 25 years that major advances have been made due to the development of a technique (Ricci et al., 1958) by which the ongoing activity of cells in the motor cortex could be recorded in awake, behaving animals (Evarts, 1968). Several studies showed that changes in motor cortical cell activity precedes the development of the motor output and relates quantitatively to its intensity and spatial characteristics (for reviews, see Evarts, 1981; Georgopoulos, 1990). Specifically, when reaching in space, cell activity during the reaction time relates primarily to the direction of the movement and less to its extent (Fu et al., 1995; Georgopoulos, 1990). Under isometric conditions, when a static force is exerted, cell activity relates to the magnitude (see Evarts, 1981) and direction (Kalaska et al., 1989; Taira et al., 1995) of force, and when a force pulse is developed, cell activity relates to the incremental, dynamic force exerted but not to the total force output (Georgopoulos et al., 1992).

Directional Tuning of Single Cells

A major aspect of the spatial characteristics of the motor output, be it movement or isometric force, is its direction in space. Cells in the motor cortex (Caminiti et al., 1990; Georgopoulos et al., 1982; Georgopoulos et al., 1986; Kalaska et al., 1989; Schwartz, 1992; Schwartz et al., 1988), as well as in other structures (Caminiti et al., 1991; Fortier et al., 1989; Kalaska et al., 1983), are directionally selective and are broadly tuned with respect to the direction of movement. This means that the cell activity is highest for a movement in a particular direction (the cell's preferred direction) and decreases progressively with movements farther and farther away from this direction. The changes in cell activity relate to the direction and not the endpoint of the reaching movement (Georgopoulos et al., 1985). Quantitatively, the crucial variable on which cell activity depends is the angle formed between the direction of the movement and the cell's preferred direction: the intensity of cell activity is a linear function of the cosine of this angle (Georgopoulos et al., 1982; Schwartz et al., 1988). An example is shown in Figure 1. The preferred directions of single cells range throughout the 3-D directional continuum (Schwartz et al., 1988).

Directional Coding by Neuronal Populations

The broad directional tuning indicates that a given cell participates in movements of various directions, and that a movement in a particular direction will involve the activation of a whole population of cells. Given that single cells are directionally tuned, we proposed a vectorial neural code for the direction of reaching by the neuronal ensemble (Georgopoulos et al., 1983; Georgopoulos et al., 1986; Georgopoulos et al., 1988), as follows: (a) A particular vector represents the contribution of a directionally tuned cell and points in the cell's preferred direction; (b) cell vectors are weighted by the change in cell activity during a particular movement; and (c) the sum of these vectors (i.e.,

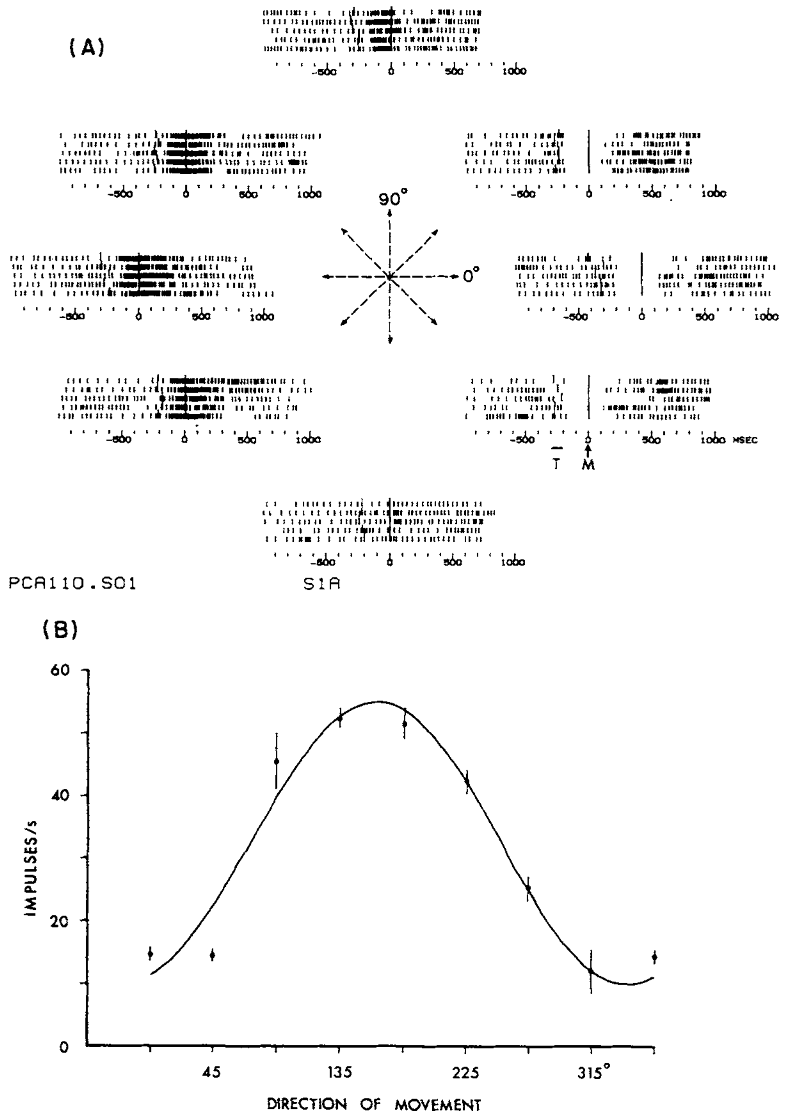

Figure 1. Directional tuning of a cell recorded in the arm area of the motor cortex during 2-D reaching. A, impulse activity during five trials of reaching in the directions indicated in the drawing at the centre. Short vertical bars indicate the occurrence of an action potential. Rasters are aligned to the onset of movement (M). Longer vertical bars preceding the onset of movement indicate the onset of the target (T); those following the movement indicate, successively, the entrance to the target zone and the delivery of reward. B, average frequency of discharge (± SEM) from the onset of the stimulus until the entry to the target zone are plotted against the direction of movement. Continuous curve is a cosine function fitted to the data using multiple regression analysis. (From Georgopoulos et al., 1982; reproduced with permission. Copyright by Society for Neuroscience.)

the population vector) provides the unique outcome of the ensemble coding operation. We found that the population vector points in the direction of the movement (Georgopoulos et al., 1983; Georgopoulos et al., 1986; Georgopoulos et al., 1988, Figure 2). The population vector approach has proved useful not only in studies of motor cortex (Caminiti et al., 1990; Georgopoulos et al., 1983; Georgopoulos et al., 1986; Georgopoulos et al., 1988; Kalaska et al., 1989) but also in studies of other motor areas, including the cerebellum (Fortier et al., 1989), the premotor cortex (Caminiti et al., 1991), area 5 (Kalaska et al., 1983), and area 7 (Steinmetz et al., 1987). Finally, the population vector code need not be confined to studies pertaining to physical space but can be generalized to arbitrary spaces and dimensions. Such an application was made successfully to the coding of faces in the inferotemporal cortex of the monkey (Young & Yamane, 1992).

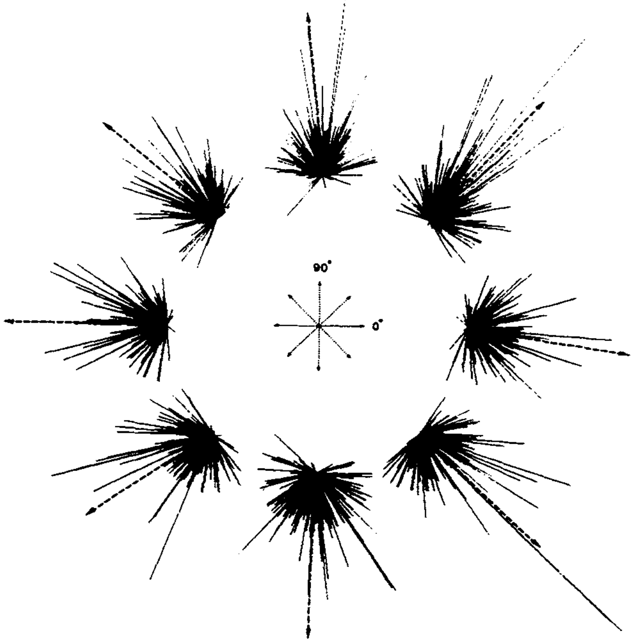

Figure 2. Population vector analysis applied to eight movement directions in 2-D space. Each cluster represents the same population; the movement directions are shown in the diagram at the centre. The population vector (interrupted arrow) points in or near the direction of the movement. (From Georgopoulos et al., 1983; reproduced with permission.)

SOME GENERAL PROPERTIES OF THE NEURONAL POPULATION VECTOR

The neuronal population vector predicts the direction of reaching for movements of different origin. In these experiments monkeys made movements that started from different points. Under these conditions, the population vector predicted well the direction of the reaching movement (Caminiti et al. 1991; Kettner et al., 1988), even if the preferred directions of individual cells shifted somewhat with different movement origins (Caminiti et al., 1990).

The direction of reaching is predicted well by neuronal population vectors in different cortical layers. The average absolute angle between the population vector calculated from cells in the upper layers (II and III) and the direction of movement was 4.31° ± 2.98° (mean ± SD, $n = 8$ movement directions), compared to 2.32° ± 2.06° for the lower layers (V and VI) (Georgopoulos, 1990). This finding suggests that the ensemble operation of the population vector can be realized separately in the upper and lower layers. This is important because that information can then be distributed to different structures according to the differential projections from the upper and lower layers (Jones & Wise, 1977).

The neuronal population coding of the direction of reaching is resistant to loss of cells. The population coding described above is a distributed code and as such does not depend exclusively on any particular cell. This robustness was evaluated by calculating the population vector from progressively smaller samples of cells randomly selected from the original population (Georgopoulos et al., 1988). Indeed, the direction of the population vector can be reliably estimated from as few as 100-150 cells, and even much fewer if optimal algorithms are used (Salinas & Abbott, 1994).

The neuronal population vector predicts the direction of dynamic isometric force (Georgopoulos et al., 1992). The dynamic relations of cell activity in the motor cortex to the direction of 2-D isometric force were investigated recently (Georgopoulos et al., 1992). The following experimental arrangement allowed the dissociation between dynamic and static components of the force exerted. Monkeys produced pure force pulses on an isometric handle in the presence of a constant force bias so that the net force (i.e., the vector sum of the monkey's force and the bias force) was in a visually specified direction. The net force developed over time had to stay in the specified direction and to increase in magnitude in order to exceed a required intensity threshold. The most interesting case is the one in which the directions of the net and bias forces differ, by being, for example, orthogonal. In order for the task to be performed successfully under these conditions, the animal's force has to change continuously in direction and magnitude, so that, at any moment during force development, the vector sum of this force and the bias force is in the visually specified direction. Thus this experimental arrangement effectively dissociated the animal's force vector, the direction of which changed continuously in a trial, from the net force vector, the direction of which remained invariant. Eight net force directions and eight bias force directions were employed. Recordings of neuronal activity in the motor cortex revealed that the activity of single cells was directionally tuned in the absence of bias force, and that this tuning remained invariant when the same net forces were produced in the presence of different directions of bias force. These results demonstrated that cell activity does not relate to the direction of the animal's force. Since the net force is equivalent to the dynamic component of the force exerted by the animal, after a static component vector (equal and opposite to the force bias) is subtracted, these findings suggest that the motor cortex provides the dynamic force signal during force development; other, possibly subcortical, structures could provide the static compensatory signal. This latter signal could be furnished by antigravity neural systems, given that most static loads encountered are gravitational in nature. According to this general view, the force exerted by the subject consists of dynamic and static components, each of which is controlled by different neural systems; these signals would converge in the spinal cord and

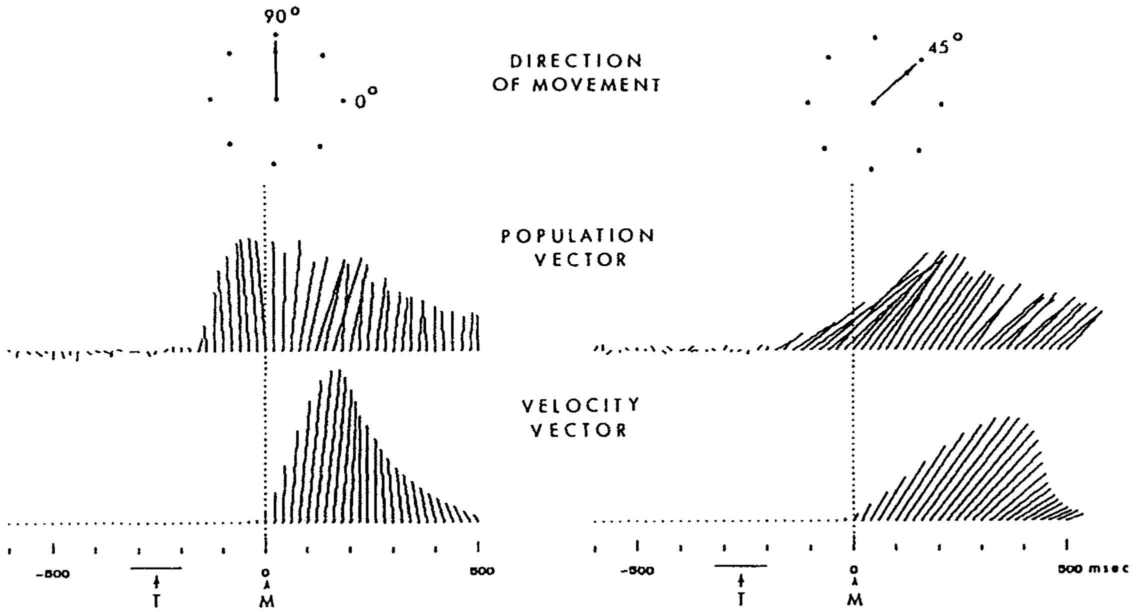

Figure 3. The population vector points in the direction of movement well before the movement begins. The results for two movement directions in 2-D space are illustrated (top); the population vector was calculated every 20 ms (middle); the average instantaneous (20 ms bin) velocity of the movement is also shown (bottom). Before the target onset (T), the population vector is very small in length and its direction varies from moment to moment. Well before the onset of movement (M), it increases in length and its direction points to the direction of the upcoming movement. This finding suggests that even the earliest inputs to the motor cortex are relevant to the direction of the upcoming movement. (From Georgopoulos et al., 1984; reproduced with permission.)

provide an ongoing integrated signal to the motoneuronal pools.

These findings established that the coding of directional information applies to the motor output in general, even in the absence of joint motion. Moreover, it showed that the direction specified by the motor cortex is not that of the total force exerted by the animal but that of the dynamic component of the force, that is the component of the force remaining after a constant, static force is subtracted.

TIME-VARYING PROPERTIES OF THE NEURONAL POPULATION VECTOR

The neuronal population vector predicts the movement trajectory in continuous, tracing movements (Schwartz, 1993, 1994). In this experiment monkeys were trained to trace smoothly with their index finger sinusoids displayed on a screen, from one end to the other. The direction of the population vectors, calculated successively in time along the trajectory, changed throughout the sinusoidal movement, closely matching the smoothly changing direction of the finger path. Moreover, a neural "image" of the sinusoidal trajectory of the movement was obtained by connecting successive population vectors tip-to-tail. This finding suggests that the length of the population vector carries information concerning the instantaneous velocity of the movement.

The neuronal population vector predicts the direction of reaching during the reaction time. This is the simplest case of predicting the direction of an upcoming movement. Given that the changes in cell activity in the motor cortex precede the onset of movement by approximately 160-180 ms, on average (Georgopoulos et al., 1982), it is an important finding that the population vector predicts the direction of the upcoming movement during that period during which the movement is being planned (Georgopoulos et al., 1984; Georgopoulos et al., 1988). An example is shown in Figure 3.

The neuronal population vector predicts the direction of reaching during an instructed delay period. In these experiments the monkeys were trained to withhold the movement for a period of time after the onset of a visual cue signal and to move later in response to a "go" signal. During this instructed delay period, the population vector in the motor cortex computed every 20 ms gave a reliable signal concerning the direction of the movement that was triggered later for execution (Georgopoulos et al., 1989b).

Neural Mechanisms of Cognitive Processing: Neuronal Populations as Keys For Understanding

The results summarized above underscore the operational usefulness of the neuronal population vector for monitoring

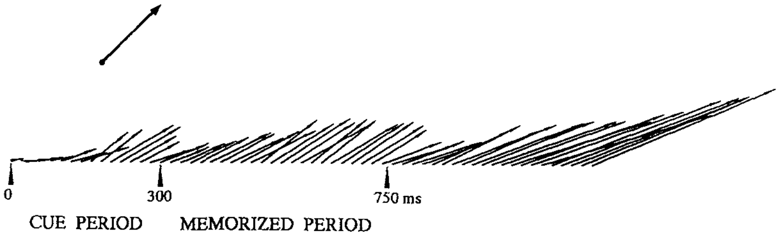

Figure 4. Population vectors in the memorized delay task for the direction indicated are plotted every 20 ms. The arrow on top indicates the direction of the *cue* signal present during the first 300 ms of the delay period. (From Smyrnis et al., 1992; reproduced with permission.)

in time the directional tendency of the neuronal ensemble. We took advantage of this property and used the population vector as a probe to decipher the neural processing of directional information during various cognitive operations.

MEMORY HOLDING

In these experiments (Smyrnis et al., 1992), two rhesus monkeys were trained to move a handle on a 2-D working surface in directions specified by a light on the plane. They first captured with the handle a light on the centre of the plane and then moved the handle in the direction indicated by a peripheral light (*cue* signal). The signal to move (*go* signal) was given by turning off the centre light. The following tasks were used. In the *non-delay* task the peripheral light was turned on at the same time as the centre light went off. In the *memorized* delay task the peripheral light stayed on for 300 ms (cue period) and the centre light was turned off 450-750 ms later (delay period). Finally, in the *non-memorized delay* task the peripheral light stayed on continuously whereas the centre light went off 750-1050 ms after the peripheral light came on. Recordings in the arm area of the motor cortex ($N = 171$ cells) showed changes in single cell activity in all tasks. The population vector was calculated every 20 ms, following the onset of the peripheral light. We were interested in two aspects of the information carried by the population vector. One concerns its direction, which can be interpreted as the directional information carried by the directional signal; the other aspect concerns the length of the population vector, which can be regarded as the strength of the directional signal carried. The direction of the population vector during the memorized delay period was close to the direction of the target (Figure 4). The length of the population vector reflects the strength of the directional signal. It is interesting that the population vector length was similar in the cue period but it was longer during the memorized versus the non-memorized part of the delay. This is shown in Figure 5 which illustrates the time course of the length of the population vector in the two delay tasks. Three phases can be distinguished in this time course. First, there is an initial increase of the vector length during the 300 ms of the delay period; this increase is similar for both tasks. Second, this increase subsides during the rest of the non-memorized delay period but continues at a somewhat higher level during the memorized delay period; the latter difference is indicated in Figure 5 by stippling. Finally, there is a steep increase in the population vector length following the *go* signal, at the end of the delay period. Thus the memorized task is distinguished from the non-memorized one by the higher population signal during that part of the delay period during which the instructed direction had to be kept in memory.

The directional information carried by the population vector in the memorized task identifies the memorized information in a direct fashion. Moreover, this analysis provided an insight concerning the time course of encoding and holding directional information. For that purpose we used the length of the population vector which can be regarded as reflecting the strength of the directional signal in the neuronal ensemble. The population vector length showed an initial increase which started approximately 100 ms following the cue onset and peaked at 250 ms. This increase was very similar in both the memorized and the non-memorized delay tasks (Figure 5). We interpret this initial peak as reflecting an *encoding* process. A second phase followed which differed in the memorized and the non-memorized tasks in that a higher, sustained signal was present during the memorized delay period but not during the non-memorized delay (stippled area in Figure 5). We interpret this as reflecting a *holding-in-memory* process. Following the onset of the "go" signal, the population vector length increased similarly in all tasks used. These findings are interesting because the increase in the signal during the memorized delay period was observed *in the absence* of the target; however, one would have expected that the signal would be stronger in the presence rather than in the absence of the visual stimulus. This finding strengthens our interpretation of this increase as a memory signal, in contrast to a sensory one and raises the more general possibility that the motor cortex may be particularly involved when only part of the visual information about an upcoming movement is provided.

MEMORIZED COMPLEX TRAJECTORY

The studies summarized above dealt with motor cortical

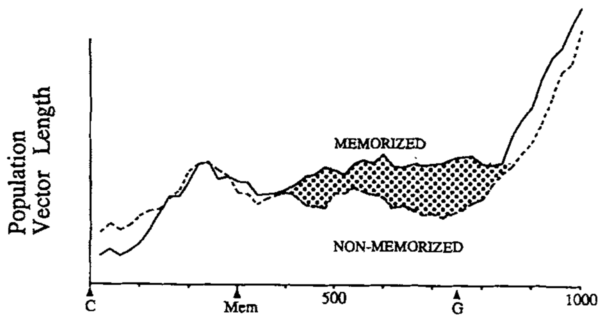

Figure 5. Length of mean resultant of the population vector is plotted against time for the two delay tasks. G, minimum time of onset of the *go* signal. (From Smyrnis et al., 1992; reproduced with permission.)

activity during a task that required memorization of the direction of a straight movement. A different question concerns the neural mechanisms subserving memorized, complex movement trajectories. This problem was investigated in a recent study (Ashe et al., 1993) in which two monkeys were trained to perform from memory an arm movement with an orthogonal bend, up and to the left, following a waiting period. They held a 2-D manipulandum over a spot of light at the centre of a planar working surface. When this light went off, the animals were required to hold the manipulandum there for 600-700 ms and then move the handle up-and-to-the-left to receive a liquid reward. There were no external signals concerning the go time or the trajectory of the movement. Following 20 trials of the memorized movement trajectory, 40 trials of visually triggered movements in radially arranged directions were performed. The activity of 137 single cells in the motor cortex was recorded extracellularly during performance of the task. A good percentage (62.8%) of cells changed activity during the waiting period. Other cells did not change activity until after the 600 ms minimum waiting time was over, and, occasionally, cell activity changed almost exactly at 600 ms after the centre light was turned off. However, the most interesting observation was that a few cells changed activity *exclusively during the execution of the memorized movement* (see Figure 5 in Ashe et al., 1993); these cells were completely inactive during performance of similar movements in the visually guided control task. These findings suggest that *performance of a movement trajectory from memory may involve a specific set of cells,* in addition to the cells activated during both visually guided and memorized movements. This idea is in accord with the results of recent modeling studies (Lukashin et al., 1994).

MENTAL ROTATION

This task required a transformation of an intended movement direction. In these studies we first carried out psychological experiments in human subjects. Then we trained monkeys to perform the same task and recorded the activity of single cells in the brain of these animals during performance of the task. Finally, we tried to connect the neural results with those of the human studies and interpret the latter on the basis of the former; the objective is to get as close as possible in relating neurophysiology and cognitive psychology. We describe below these steps as they were applied to a particular problem of a mental transformation of movement direction.

The task required subjects to move a handle at an angle from a reference direction defined by a visual stimulus on a plane. Since the reference direction changed from trial to trial, the task required that, in a given trial, the direction of movement be specified according to this reference direction. In the psychological studies (Georgopoulos & Massey, 1987), human subjects performed blocks of 20 trials in which the angle above and its departure (counterclockwise or clockwise) were fixed, although the reference direction varied. Seven angles (5-14°) were used. The basic finding was that the time to initiate a movement (reaction time) increased in a linear fashion with the angle. The most parsimonious hypothesis to explain these results is that subjects arrive at the correct direction of movement by shifting their motor intention from the reference direction to the movement direction, traveling through the intermediate angular space. This idea is very similar to the mental rotation hypothesis of Shepard and Cooper (1982) to explain the monotonic increase of the reaction time with orientation angle when a judgement has to be made whether a visual image is normal or mirror-image: In both cases a mental rotation is postulated. In fact, the mean rates of rotation and their range among subjects were very similar in the perceptual (Shepard & Cooper, 1982) and motor (Georgopoulos & Massey, 1987) studies. Moreover, when the same subjects performed both perceptual and motor rotation tasks, their processing rates were positively correlated (Pellizzer & Georgopoulos, 1993) which indicates similar processing constraints for both tasks.

In the neurophysiological studies (Georgopoulos et al. 1989a; Lurito et al., 1991), two rhesus monkeys were trained to move the handle 90° and counterclockwise from the reference direction; these trials were intermixed with others in which the animals moved in the direction of the target. When the time-varying neuronal population vector was calculated during the reaction time, it was found that it rotated from the stimulus (reference) to the movement direction through the counterclockwise angle, when the animal had to move away from the stimulus, or in the direction of the stimulus, when the animal had to move towards it. This is illustrated in Figure 6. It is remarkable that the population vector rotated at all, and also that it rotated through the smaller, 90° counterclockwise angle. These results showed clearly that the cognitive process in this task truly involved a rotation of an analogue signal. The occurrence of a true rotation was further documented by showing that there was a transient increase during the

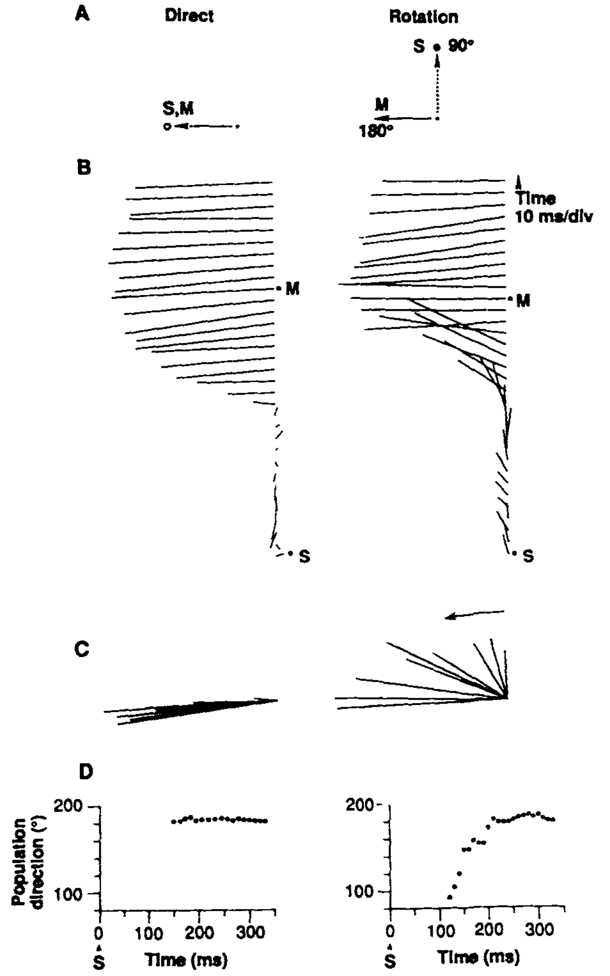

Figure 6. Results from a "direct" (left) and "rotation" (right) movement. (A) Task. Unfilled and filled circles indicate dim and bright light, respectively. Interrupted and continuous lines with arrows indicate stimulus (S) and movement (M) direction, respectively. (B) Neuronal population vectors calculated every 10 ms from the onset of the stimulus, S, at positions shown in (A) until after the onset of the movement (M). When the population vector lengthens, for the direct case (left) it points in the direction of the movement, whereas for the "rotation" case it points initially in the direction of the stimulus and then rotates counterclockwise (from 12 o'clock to 9 o'clock) and points in the direction of the movement. (C) Ten successive population vectors from (B) are shown in a spatial plot, starting from the first population vector that increased significantly in length. Notice the counterclockwise rotation of the population vector (right panel). (D) Scatter plots of the direction of the population vector as a function of time, starting from the first population vector that increased significantly in length following stimulus onset (S). For the direct case (left panel) the direction of the population vector is in the direction of the movement (~180°); for the rotation case (right panel) the direction of the population vector rotates counterclockwise from the direction of the stimulus (~90°) to the direction of the movement (~180°).

middle of the reaction time in the recruitment of cells with preferred directions in-between the stimulus and movement directions; this indicated that the rotation of the population vector was not the result of varying activation of just two cell groups, one with preferred directions centred on the stimulus and another in the movement direction. Therefore, this rotation process, "sweeping" through the directionally tuned ensemble, provided for the first time a direct visualization of a dynamic cognitive process. In this respect, it is noteworthy that the population vector is, in essence, a measure that can take continuous values in direction.

In summary, the results of these studies provide the neural correlates of a dynamic cognitive representation (Freyd, 1987). The essential contribution of this work is in the neural identification and visualization of the time-varying, dynamic representation of direction in the motor cortex when a transformation of this direction is required and achieved. Interestingly, the mean rotation rate and the range of rates observed for different reference directions were very similar to those obtained in the human studies.

CONTEXT-RECALL MEMORY SCANNING

In the studies above we identified the neural correlates of a mental rotation process as the orderly rotation of the neuronal population vector from a stimulus to a movement direction, through successive directions within a specified angle. This rotation exemplified the spatial rule operating in the mental rotation task, which required the production of a movement at an angle from a stimulus direction. In this study (Pellizzer et al., 1995) we sought, instead, to determine the neural correlates of a cognitive process the rule of which was based not on a spatial constraint but on the serial position of stimuli in a sequence. Given an arbitrary sequence of stimuli on a circle, one of which was identified as the test stimulus, the motor response had to be towards the stimulus that followed the test stimulus in the sequence. This task is a visuomotor version (Georgopoulos & Lurito, 1991; Pellizzer & Georgopoulos, 1993) of the context-recall memory scanning task (Sternberg, 1969). Previous psychophysical studies (Pellizzer & Georgopoulos, 1993) suggested that the processing mechanisms differ between the mental rotation and context-recall tasks. In order to determine the neural mechanisms in the latter task, we recorded single cell activity in the motor cortex of a monkey trained to perform a context-recall and a control, instructed delay task. Moreover, we reanalyzed the neural data from the mental rotation study (Lurito et al., 1991) to compare them with those obtained in the present study.

In the control task, a yellow stimulus was presented in one of eight directions and stayed on for 400 ms, after which it turned blue. This provided the go signal for the monkey to exert a force pulse such that a force feedback cursor exceeded a certain threshold. In the context-recall task, three yellow stimuli (list stimuli) were presented successively (every 400 ms) at different positions on the circle, and stayed on the screen. After an additional 400 ms, one of these stimuli (except the last) turned blue. This identified the test stimulus, and provided the go signal: now, the monkey had to move the cursor in the direction of the stimulus that followed the test stimulus in the sequence. We found that during the response time, the patterns of neural activity in the motor cortex initially resembled those associated with the direction of the second stimulus. When the test stimulus was the first in the sequence, cell activity continued to reflect the direction of the second stimulus which, in this case, was the appropriate motor response. However, when the test stimulus was the second in the sequence, neural activity changed abruptly to reflect the pattern associated with the direction of the third stimulus which was now the appropriate motor response. This switch was abrupt, occurred ~100-150 ms after the go signal, and was evident both in the activity of single cells and in the time-varying neuronal population vector which changed direction within ~50-60 ms.

These findings identify the neural correlates of a switching process that is different from a mental rotation described previously (Lurito et al., 1991). Additional evidence for the basically different nature of the two neural processes was provided by an analysis of the directional selectivity of cells that changed activity during the response time. The idea is that in a rotation process, the set of cells that change activity during the response time should include cells the preferred directions of which would be in the intermediate direction between the stimulus and response directions; indeed, this was observed (see Figure 13 in Lurito et al., 1991). In contrast, a switching operation such as postulated for the context-recall task, should not involve the activation of cells in directions intermediate between the test stimulus (S2) and motor response (S3). Indeed, this was observed (see Figure 4 in Pellizzer et al., 1995). It seems then that the time taken to derive the motor direction in the mental rotation task reflects a transformation, whereas the time taken in the context-recall task reflects a selection process. Finally, it should be noted that these studies provided an insight into the neural mechanisms of these processes in a particular brain area, namely the motor cortex, but it is obvious that other brain areas are likely to be involved. Additional experiments are needed to delineate the identification of such areas and elucidate their relative contributions to task performance.

Neural Network Modeling Studies

GENERAL

The motor cortex can be essentially regarded as a neural network. The elements of the network are directionally tuned and one of the possible operations of the network is the computation of the neuronal population vector. The broad directional tuning of single cells seems to be a general effect of the population operation, because this property appeared in the hidden layer of a three-layer network trained to calculate the population vector (Lukashin, 1990). Another question concerns the interactions among cells. It is known that there are extensive local interconnections (Huntley & Jones, 1991) as well as functional neuronal interactions (Asanuma & Brooks, 1965; Stefanis & Jasper, 1964). An objective then is to (a) identify these interactions among the directionally tuned cells in the motor cortex and discover the rules that govern their presence, (b) study an artificial neural network made of directionally tuned elements with massive interconnections, and (c) compare the results obtained in the real (i.e., motor cortical) and artificial networks. We found the following (Georgopoulos et al., 1993). First, interactions between cells in the motor cortex relate to the directional tuning of the cells in a pair: interactions are more than twice as frequent when the cells are tuned than when they are not, and, for directionally tuned cells, the interaction ranges from strongly positive (i.e., excitatory) to strongly negative (i.e., inhibitory) as the angle between the preferred directions of the cells in a pair varies from 0° (i.e., same preferred direction) to 180° (i.e., opposite preferred directions). Second, the same trend was observed between the directionally tuned elements in a massively interconnected, dynamical artificial network during a stable period of computation of the neuronal population vector. Third, it was found that for the network to be stable, the strength of the synaptic interactions has to be low: in the best (i.e., most stable) case the mean synaptic strength tends to $2/N$, where N is the number of elements in the network. This is in keeping with the fact that cortical cells in an area are extensively but weakly interconnected (Martin, 1988). These findings validate the correspondence between the motor cortical and the artificial neural network and open the possibility of using this network in the temporal domain to explore the mechanisms of the cognitive operations described in the preceding section.

The time-varying directional operations discussed in the preceding sections have been modelled recently using a massively interconnected artificial neural network that consists of directionally tuned neurons and produces as an outcome the neuronal population vector (Lukashin & Georgopoulos, 1994a, b). This network has reproduced well many of the experimental findings and has led to a novel hypothesis on how memorized trajectories of complex movements could be stored in the synaptic connections of overlapping neural networks (Lukashin et al., 1994). The idea is that there is a general purpose network that is involved in all kinds of movements, memorized or not, but which carries no information about memorized trajectories of specific shapes (e.g., circles, ellipses, scribbles) and which, if activated alone, would produce straight line trajectories. It is now hypothesized that there are also networks highly specific for a particular trajectory (e.g., clockwise circle) which are interconnected with the general purpose network; when a specific trajectory needs to be performed, the

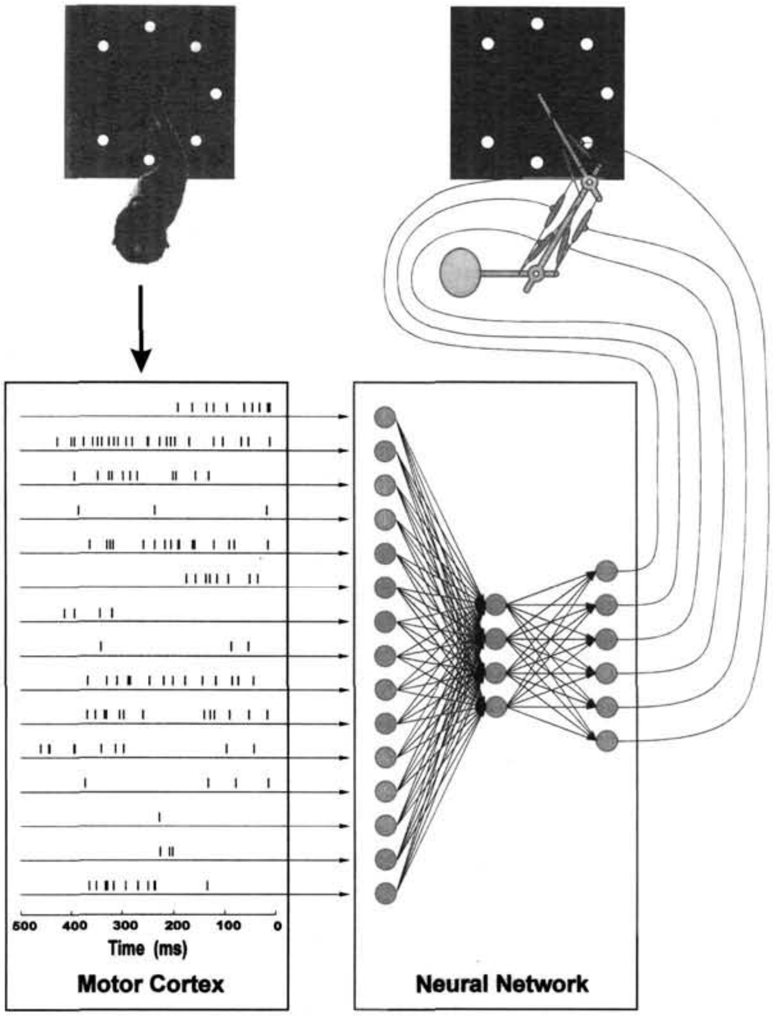

Figure 7. The decoding computation scheme used to transform neuronal commands encoded in a series of action potentials into a force exerted by the simulated actuator. The top left part of the figure illustrates a monkey exerting a force against the immovable handle in one (180°) of eight instructed directions. An example of the motor cortical activity recorded while the animal performed this task is represented in the bottom left panel. The spike trains were recorded for different trials but for the same instructed direction of force. These neuronal signals drive the simulated actuator sketched in the top right part of the figure. The actuator muscles are modeled as nonlinear springs with an exponential length-tension relationship: $f(l) = k\{\exp[\alpha(l - l_0) - 1\}$ where $f(l)$ is a contraction force developed by a muscle at length l, l_0 is the muscle rest length, and k and α are constants. A three-layered feedforward neural network (the directed connections are shown by thin arrows) transforms cortical signals into coordinated activation of actuator muscles. Activities of units at the input layer are the spike trains taken as they are from an experimentally obtained data file. The cortical activity converges in four model units located at the intermediate layer. The intermediate units provide integrated (over inputs and over time) signals to six output units. Activities of intermediate and output units were calculated using sigmoid activation function, $y(x) = 1/2[1 + \tanh(x - \beta)]$, where β is a threshold and x is a synaptic input entering into the unit. The output activity of the network changes the muscles rest lengths (each output unit innervates one muscle; a linear relation between the rest length l_0 and output unit activity was used: $l_0(y) = a + by$, where a and b are constants). If the actuator is free to move, this will cause the contraction of muscles and, as a result, the actuator's endpoint will move into a new equilibrium position. If, however, the actuator's endpoint is blocked by an immovable object (the possibility that was explored here), then the tensions developed by muscles due to the change of the muscle rest lengths will produce the endpoint force against that object.

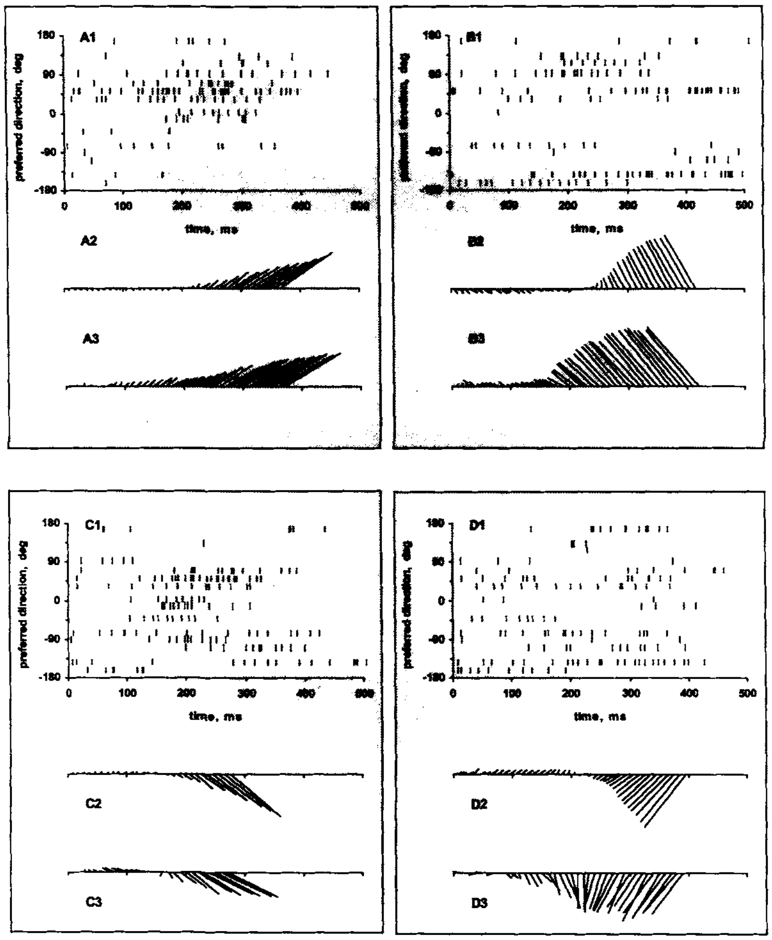

Figure 8. The motor cortical activity and the forces exerted by monkey and by simulated actuator. Four rasters *A1, B1, C1,* and *D1* show spiking activity recorded while monkey developed isometric forces in four different instructed directions (45°, 135°, -45°, and -135°, respectively). The spike trains are aligned to the time of visual stimulus (zero time), instructing the monkey to begin the motor action. Each raster includes spike trains for the same 15 cells ordered along the vertical axis in accordance with their preferred directions. The time-varying forces developed by the monkey are depicted in *A2, B2, C2,* and *D2*. The force vectors measured every 10 ms are displayed as line segments; the tails of force vectors are aligned along the time axis (the horizontal line). The time scale is the same as shown in the rasters above. Finally, *A3, B3, C3,* and *D3* show forces exerted by the actuator in response to the neuronal activity presented in the rasters. The forces were calculated using the decoding algorithm described (see text and Figure 3) with a 10-ms time step and were arranged along the time axis in the same way as experimentally measured forces. The measured and calculated forces are normalized to the same magnitude assigned arbitrarily.

appropriate specific network fuses with the general purpose network and, now acting as one network, produces the desired trajectory. It is remarkable that the size of the specific network need only be less than 5% of the size of the general purpose network for the desired trajectory to be effectively stored and reproduced (Lukashin et al., 1994). It is noteworthy that such very specific cells have been observed at low proportions in neurophysiological recordings during performance of memorized trajectories (Ashe et al., 1993), as discussed above.

The idea of the existence of very specialized networks raises the question of the degree of specialization and of how such networks are created in the first place. We can only speculate on these issues. With respect to the general-purpose network, it is reasonable to assume that it is present at birth, since it is assumed to subserve all movements. There are several possibilities concerning the specialized networks. One possibility is that a number of small-size networks, specific for basic shapes (e.g., straight lines, curves, and some combinations thereof ["motor shape primitives"]), are present at birth. Then, motor learning for other, complex motor acts would consist of adjusting the connections strengths between the general-purpose and specific networks. This idea implies that all of the specialized primitives are used routinely, although not as frequently as the general-purpose network. Another possibility is that innate specific networks code for more complicated shapes and are large in number. The mechanism of motor learning would then be similar as that described above, but in this case only a small number of the specialized networks would be used. This means that a number of the complex, specialized networks may never be used. This situation would be similar to that encountered in the immune system in which there is a large potential for making a large number of antibodies, of which, however, only some may be actually made, depending on the exposure of the organism to specific antigens. In both cases there is a selection: selection of a specialized trajectory or selection of an antibody, both from a large ensemble available. Finally, an intermediate hypothesis would be that we start with motor shape primitives, but that the more complex trajectories resulting from the combination of these primitives with the general network become themselves very specialized and behave as such in the formation of other trajectories in novel associations.

TRANSFORMATION OF SPIKING ACTIVITY INTO ISOMETRIC FORCE EXERTED BY AN ACTUATOR

The population vector code, by combining activities of broadly tuned cells, provides an unambiguous and reliable estimation of the upcoming motor output (movement or force). However, the population-coding hypothesis concerns only the cortical representation of motor commands. It does not answer the question of how the motor commands encoded in the cell activities are translated into coordinated contraction of limb muscles to generate a desired motor action. This problem is closely related to the design of adaptive systems that transform neuronal signals chronically recorded from the motor cortex into the physiologically appropriate motor output of artificial actuators such as multijoint prosthetic limbs.

Despite apparent similarities between these two problems, the suitable approaches to attack each one pursue different goals. The methods used to solve the first problem are heavily based on anatomical, neurophysiological, and biomechanical properties of the actual biological structures involved. The ultimate goal of this data-driven approach is to understand how the real biological systems implement the motor control. In contrast, the methods used to solve the second problem are usually based on theoretical analysis, the ultimate goal of which is to develop a computational algorithm that utilizes the raw biological signals to drive an artificial actuator. Here, the computational scheme does not try to be biologically realistic. Any control algorithm that successfully solves the problem is acceptable. Despite the difference in goals, the importance of interplay between the data- and theory-driven approaches should not be underestimated. The ideas and concepts developed in one field drive the other, and vice versa.

This was illustrated by Lukashin et al. (1996a) who addressed the question of how neuronal signals might be used to drive an artificial actuator so that its motor output would correspond to the performance of a real limb. The suggested computational scheme, which transformed impulse activity recorded from the monkey motor cortex during performance of the isometric force task to the force produced by a simulated actuator, is not biologically realistic. However, this decoding scheme was inspired by, and based on, experimental findings obtained in anatomical, neurophysiological and psychophysical studies conducted on three different organisms: (a) experiments on microstimulation of the frog's spinal cord (Bizzi et al., 1991; Mussa-Ivaldi et al., 1994); (b) studies of human arm stiffness characteristics (Mussa-Ivaldi et al., 1985; Tsuji et al., 1995); and (c) single cell recordings in the motor cortex of monkey (Georgopoulos et al., 1992). The motor cortical activity used in Lukashin et al. (1996b) as command signals came from single cell recording experiments (Georgopoulos et al., 1992) discussed above, when there was no force bias applied ($\mathrm{B} = 0$). In this case, the force S exerted by the monkey had only the dynamic component I and was the only force acting on the isometric handle. Therefore the force S developed over time had to be in the instructed direction and had to increase in magnitude in order to exceed a required threshold. The bottom left part of Figure 7 shows an example of impulse activity of $\mathrm{N} = 15$ different cells. These spike trains were recorded in different trials but for the same instructed direction of force (180° in the 2D workspace, Figure 7, top left). The key idea is that these cortical signals must now drive a simulated actuator (Figure 7, top right) in such a way that it would exert an isometric force in the same direction as the monkey did. The actuator is a planar two-joint, six-muscle model of the arm. The transformation of cortical signals into motor output of the actuator is performed by an artificial neural network (Figure 7, bottom right) connected to the actuator. The network receives experimentally measured impulse activity as a time varying input to the input layer and transforms it into a time varying pattern of activity at the output layer. This results in contraction of actuator "muscles" by means of changing the muscle rest lengths. Finally, a set of the muscle rest lengths unambiguously defines the direction and magnitude of the endpoint force exerted by the actuator against an immovable object.

In the framework of this computational scheme, the performance of the model (i.e., the relation between the input cortical signals and the force exerted) depends mainly on the network connectivity, which must provide a synergistic activation of all muscles to generate a required motor action. To ensure physiologically normal motor output of the actuator, the network was trained (Lukashin et al., 1996a) on experimental data obtained from studies of human arm stiffness (Mussa-Ivaldi et al., 1985). As a result, the stiffness properties of the model arm were similar to those measured for the human arm.

After training, the underlying set of synaptic weights was fixed and the performance of the model was tested against the whole neurophysiological data set (Georgopoulos et al., 1992) that included both experimentally measured motor cortical commands and resulting motor actions. An important issue of the performance is the robustness of the decoding scheme with respect to the size of the population of cells generating neuronal signals, with respect to the variations in the composition of cells included in the population of a given size, and, finally, with respect to changes in the cell activity from trial to trial.

The uniformity of the distribution of the cells' preferred directions throughout space is of particular importance in reconstructing the directional signal encoded in the population activity. Therefore, all other factors being equal, the best performance for a fixed size population is expected for such a composition of cells whose preferred directions form a nearly uniform distribution. It was found (Lukashin et al., 1996b) that when the uniformity requirement is fulfilled the performance and robustness of the model improves gradually as the size, N, of the population increases, revealing a tendency for saturation as N approaches to 15-20 cells.

Typical results for $N = 15$ demonstrating the time-evolving performance of the model are displayed in Figure 8. The forces exerted by the monkey and by the actuator are shown for four instructed directions together with the corresponding motor cortical activity. It can be seen (Figure 8, templates *A1, B1, C1, D1*) that directional tuning of cortical impulse activity is barely perceptible due to the large variability in cell discharge. However, the computation scheme successfully decodes these signals, and the time-varying forces developed by the actuator (Figure 8, templates *A3, B3, C3, D3*) are very similar to those developed by the monkey arm (Figure 8, templates *A2, B2, C2, D2*). Following an initial period of time, which lasts 100-200 ms, the direction of force exerted by the actuator stabilizes and the magnitude of force increases. The stabilized direction of force is close to the instructed direction, for which the cortical activity was recorded. This was also observed for the remaining four instructed directions for this particular ensemble of cells and for other ensembles of cells ($N \geq 15$) and trials, thus suggesting a high degree of robustness of the decoding algorithm.

Concluding Remarks

One of the major challenges of brain theory is to elucidate the neural basis of behaviour. This was also the prime concern of Hebb (1949), and, in a sense, our work on neuronal populations and their modeling can be thought of as following the thread that he started. Through that thread, it became possible to interpret the functional properties of single cells with respect to relatively simple yet behaviourally meaningful motor actions. The finding that neurons are broadly tuned to behavioural variables led to the key idea of distributive coding. The population vector algorithm allows to read out this code by transforming aggregates of purely temporal spike trains into a spatio-temporal vector. This discovery of cortical representation of motor commands, combined with the elucidation of neural mechanisms by which these commands generate a particular behavioural pattern unfolding in time, should further advance our understanding of neural basis of motor behaviour. The research in this direction also provides an impetus to the field of designing adaptive systems that transform neuronal signals recorded from the brain into physiologically accurate motor output of multijoint prosthetic limbs.

This work was supported by United States Public Service grants NS17413 and PSMH48185, and the US Department of Veterans Affairs, and the American Legion Brain Sciences Chair.

Send correspondence to second author at Brain Sciences Center (11B), Veterans Affairs Medical Center, One Veterans Drive, Minneapolis, MN 55417.

References

Asanuma, H., & Brooks, V. B. (1965). Recurrent cortical effects following stimulation of internal capsule. *Archives Italiennes de Biologie, 103,* 220-246.

Ashe, J., Taira, M., Smyrnis, N., Pellizzer, G., Georgakopoulos, T., Lurito, J. T., & Georgopoulos, A.P. (1993). Motor cortical activity preceding a memorized movement trajectory with an orthogonal bend. *Experimental Brain Research, 95,* 118-130.

Bizzi, E., Mussa-Ivaldi, F. A., & Giszter, S. F. (1991). Computations underlying the execution of movement: A biological perspective. *Science, 253,* 287-291.

Caminiti, R., Johnson, P. B., Galli, C., Ferraina, S., Burnod, Y., & Urbano, A. (1991). Making arm movements within different parts of space: The premotor and motor cortical representation of a coordinate system for reaching at visual targets. *Journal of Neuroscience, 11,* 1182-1197.

Caminiti, R., Johnson, P.B., & Urbano, A. (1990). Making arm movements within different parts of space: Dynamic aspects in the primate motor cortex. *Journal of Neuroscience, 10,* 2039-2058.

Evarts, E. V. (1968). Relation of pyramidal tract to force exerted during voluntary movement. *Journal of Neurophysiology, 31,* 14-27.

Evarts, E. V. (1981). Role of motor cortex in voluntary movements in primates. In V. B. Brooks (Ed.), *Handbook of physiology - The nervous system, vol II - Motor control* (pp. 1083-1120). Bethesda: American Physiological Society.

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.

Freyd, J. J. (1987). Dynamic mental representation. *Psychological Review, 94,* 427-438.

Fu, Q.-G., Flament, D., Coltz, J. D., & Ebner, T. J. (1995). Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *Journal of Neurophysiology, 73,* 836-854.

Georgopoulos, A. P. (1990). Neurophysiology of reaching. In M. Jeannerod (Ed.), *Attention & Performance XIII* (pp. 227-263). Hillsdale, NJ: Erlbaum.

Georgopoulos, A. P., Ashe, J., Smyrnis, N., & Taira, M. (1992). Motor cortex and the coding of force. *Science, 256,* 1692-1695.

Georgopoulos, A. P., Caminiti, R., Kalaska, J. F., & Massey, J. T. (1983). Spatial coding of movement: A hypothesis concerning the coding of movement direction by motor cortical populations. *Experimental Brain Research Supplement, 7,* 327-336.

Georgopoulos, A. P., Crutcher, M. D., & Schwartz, A. B. (1989b). Cognitive spatial-motor processes 3. motor cortical prediction of movement direction during an instructed delay period. *Experimental Brain Research, 75,* 183-194.

Georgopoulos, A. P., Kalaska, J. F., & Caminiti, R. (1985). Relations between two-dimensional arm movements and single cell discharge in motor cortex and area 5: Movement direction versus movement endpoint. *Experimental Brain Research Supplement, 10,* 176-183.

Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience, 2,* 1527-1537.

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 G. M. Edelman, W. E. Gall, & W. M. Cowan (Eds.), *Dynamic aspects of neocortical function* (pp. 501-524). New York: Wiley.

Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience, 8,* 2928-2937.

Georgopoulos, A. P., & Lurito, J. T. (1991). Cognitive spatial-motor processes. 6. Visuomotor memory scanning. *Experimental Brain Research, 83,* 453-458.

Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989a). Mental rotation of the neuronal population vector. *Science, 243,* 234-236.

Georgopoulos, A. P., & Massey, J. T. (1987). Cognitive spatial-motor processes 1. The making of movements at various angles from a stimulus direction. *Experimental Brain Research, 65,* 361-370.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science, 233,* 1416-1419.

Georgopoulos, A. P., Taira, M., & Lukashin, A. (1993). Cognitive neurophysiology of the motor cortex. *Science, 260,* 47-52.

Hebb, D. O. (1949). *The organization of behavior.* New York: Wiley.

Huntley, G. W., & Jones, E. G. (1991). Relationship of intrinsic connections to forelimb movement representations in monkey motor cortex: A correlative anatomic and physiological study. *Journal of Neurophysiology, 66,* 390-413.

Jones, E. G., & Wise, S. P. (1977). Size, laminar and columnar distribution of efferent cells in the sensory-motor cortex of monkeys. *Journal of Comparative Neurology, 175,* 391-438.

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. *Experimental Brain Research, 51,* 247-260.

Kalaska, J. F., Cohen, D. A. D., 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.

Kettner, R. E., Schwartz, A. B., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. III. Positional gradients and population coding of movement direction from various movement origins. *Journal of Neuroscience, 8,* 2938-2947.

Lukashin, A. V. (1990). A learned neural network that simulates properties of the neuronal population vector. *Biological Cybernetics, 63,* 377-382.

Lukashin, A. V., Amirikian, B. R., & Georgopoulos, A. P. (1996a). Neural computations underlying the exertion of force: A model. *Biological Cybernetics, 74,* 469-478.

Lukashin, A. V., Amirikian, B. R., & Georgopoulos, A. P. (1996b). A simulated actuator driven by motor cortical signals. *NeuroReport, 7,* 2597-2601.

Lukashin, A. V., & Georgopoulos, A. P. (1994a). A neural network for coding of trajectories by time series of neuronal population vectors. *Neural Computation, 6,* 19-28.

Lukashin, A. V., & Georgopoulos, A. P. (1994b). Directional operations in the motor cortex modeled by a neural network of spiking neurons. *Biological Cybernetics, 71,* 79-85.

Lukashin, A. V., Wilcox, G. L., & Georgopoulos, A. P. (1994). Overlapping neural networks for multiple motor engrams. *Proceedings of the National Academy of Sciences of the U.S.A., 91,* 8651-8654.

Lurito, J. L., Georgakopoulos, T., & Georgopoulos, A. P. (1991). Cognitive spatial-motor processes. 7. The making of movements at an angle from a stimulus direction: Studies of motor cortical activity at the single cell and population levels. *Experimental Brain Research, 87,* 562-580.

Martin, K. A. C. (1988). From single cells to simple circuits in the cerebral cortex. *Quarterly Journal of Experimental Physiology, 73,* 637-702.

Mussa-Ivaldi, F. A., Giszter, S. F., & Bizzi, E. (1994). Linear combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences of the U.S.A., 91,* 7534-7538.

Mussa-Ivaldi, F. A., Hogan, N., & Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. *Journal of Neuroscience, 5,* 2732-2743.

Pellizzer, G., & Georgopoulos, A. P. (1993). Common processing constraints for visuomotor and visual mental rotation. *Experimental Brain Research, 93,* 165-172.

Pellizzer, G., Sargent, P., & Georgopoulos, A. P. (1995). Motor cortical activity in a context-recall task. *Science, 269,* 702-705.

Ricci, G., Doane, B., & Jasper, H. (1957). Microelectrode studies of conditioning: Technique and preliminary results. In *Premier congrès international des sciences neurologiques* (pp. 401-415). Brussels: Snoeck-Ducaju.

Salinas, E., & Abbott, L. F. (1994). Vector reconstruction from firing rates. *Journal of Computational Neuroscience, 1,* 89-107.

Schwartz, A. B. (1992). Motor cortical activity during drawing movements: Single unit activity during sinusoid tracing. *Journal of Neurophysiology, 68,* 528-541.

Schwartz, A. B. (1993). Motor cortical activity during drawing movements: population representation during sinusoid tracing. *Journal of Neurophysiology, 70,* 28-36.

Schwartz, A. B. (1994). Direct cortical representation of drawing. *Science, 265,* 540-542.

Schwartz, A. B., Kettner, R. E., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *Journal of Neuroscience, 8,* 2913-2927.

Shepard, R. N., & Cooper, L. A. (1982). *Mental images and their transformations.* Cambridge: MA: MIT Press.

Smyrnis, N., Taira, M., Ashe, J., & Georgopoulos, A. P. (1992). Motor cortical activity in a memorized delay task. *Experimental Brain Research, 92,* 139-151.

Stefanis, C., & Jasper, H. (1964). Recurrent collateral inhibition in pyramidal tract neurons. *Journal of Neurophysiology, 27,* 855-877.

Steinmetz, M. A., Motter, B. C., Duffy, C. J., & Mountcastle, V. B. (1987). Functional properties of parietal visual neurons: Radial organization of directionalities within the visual field. *Journal of Neuroscience, 7,* 177-191.

Sternberg, S. (1969). Memory scanning: Mental processes revealed by reaction time experiments. *American Scientist, 54,* 421-457.

Taira, M., Boline, J., Smyrnis, N., Georgopoulos, A. P., & Ashe, J. (1995). On the relations between single cell activity in the motor cortex and the direction and magnitude of three-dimensional static isometric force. *Experimental Brain Research, 109,* 367-376.

Tsuji, T., Morasso, P. G., Goto, K., & Ito, K. (1995). Human hand impedance characteristics during maintained posture. *Biological Cybernetics, 72,* 475-485.

Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science, 256,* 1327-1331.

Sommaire

Cet article porte sur la fonction du cortex moteur tout comme le révèlent les enregistrements électriques de cellules individuelles chez des singes de même que la modélisation d'un réseau neuronal formel. Notre approche conceptuelle, clé tant au plan des neurosciences du comportement qu'à celui de la modélisation d'un réseau neuronal de la fonction du cortex moteur, repose sur la reconstitution, la visualisation, et la modélisation de l'activité des populations neuronales, qui a été en fait un concept clé développé par Hebb (1949). L'enquête sur le comportement s'étend de l'effort de la force isométrique à la direction des mouvements et à celui du traitement cognitif complexe. Les propriétés fonctionnelles de cellules individuelles en regard de l'orientation du mouvement dans l'espace sont décrites de même qu'un codage par population qui fournit une mesure unique pour cette direction. Enfin, les résultats des études de modélisation dans lesquelles l'activité orientée de la population est utilisée comme entrée dans un réseau neuronal formel pour diriger un bras simulé sont discutés.