

VISUAL NEUROSCIENCE

An International Journal
for Experimental and
Theoretical Research


EDITOR:
James T. McIlwain, *Brown University*

Visual Neuroscience is devoted to the publication of research and theoretical articles in basic visual neuroscience, with primary emphasis on retinal and brain mechanisms that underlie visually-guided behaviors and visual perception. The major goal of *Visual Neuroscience* is to bring together in one journal a broad range of studies which reflect the exciting diversity and originality of contemporary research in this field. Methodologies are drawn from neuroanatomy, neurophysiology, neurochemistry, neuroimmunology, and behavioral science, as well as computational models and computer-assisted formulations. Molecular, cellular, local-circuit, and systems-level analyses in both vertebrate and invertebrate species are presented.

Visual Neuroscience (ISSN 0952-5238) is published bi-monthly. Subscription to Volume 10, 1993: \$320.00 for institutions, \$125.00 for individuals, single parts \$55.00.

Send orders to:
CAMBRIDGE UNIVERSITY PRESS
40 West 20th Street
New York, NY 10011-4211

Outside the USA, Canada, and Mexico:
CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building
Shaftesbury Road
Cambridge CB2 2RU, England



Does the nervous system use equilibrium-point control to guide single and multiple joint movements?

**E. Bizzi, N. Hogan,* F. A. Mussa-Ivaldi
and S. Giszter**

*Department of Brain and Cognitive Sciences and *Department of
Mechanical Engineering, Massachusetts Institute of Technology,
Cambridge, MA 02139*
Electronic mail: emilio@wheaties.ai.mit.edu

Abstract: The hypothesis that the central nervous system (CNS) generates movement as a shift of the limb's equilibrium posture has been corroborated experimentally in studies involving single- and multi-joint motions. Posture may be controlled through the choice of muscle length-tension curve that set agonist-antagonist torque-angle curves determining an equilibrium position for the limb and the stiffness about the joints. Arm trajectories seem to be generated through a control signal defining a series of equilibrium postures. The equilibrium-point hypothesis drastically simplifies the requisite computations for multi-joint movements and mechanical interactions with complex dynamic objects in the environment. Because the neuromuscular system is springlike, the instantaneous difference between the arm's actual position and the equilibrium position specified by the neural activity can generate the requisite torques, avoiding the complex "inverse dynamic" problem of computing the torques at the joints. The hypothesis provides a simple, unified description of posture and movement as well as contact control task performance, in which the limb must exert force stably and do work on objects in the environment. The latter is a surprisingly difficult problem, as robotic experience has shown. The prior evidence for the hypothesis came mainly from psychophysical and behavioral experiments. Our recent work has shown that microstimulation of the frog spinal cord's premotoneuronal network produces leg movements to various positions in the frog's motor space. The hypothesis can now be investigated in the neurophysiological machinery of the spinal cord.

Keywords: contact tasks; equilibrium point; force field; inverse dynamics; microstimulation; motor control; multi-joint coordination; robotics; spinal cord

1. Introduction

The purpose of this target article is to present a critical evaluation of the equilibrium-point hypothesis. Before discussing its strengths and weaknesses, we would like to make clear why this hypothesis was proposed. To this end, we summarize briefly the transformations that are thought to occur when a sensory stimulus (such as an object to be reached) appears in the environment. The first step in carrying out a reaching task involves a transformation performed by cortical parietal cells. These cells receive visual, orbital, and neck afferent information. The integration of the information from these different sources generates a neural code representing the location of an object with respect to the body and the head (Andersen et al. 1985b). The second step involves the planning of the direction of hand motion and presumably its velocity and amplitude. Psychophysical observations by Morasso (1981) have suggested that this planning stage is carried out in extrinsic coordinates that represent the motion of the hand in space. In the same vein, recordings

from single cells in cortical and subcortical areas have shown a correlation between their firing pattern and the direction of hand motion (Georgopoulos et al. 1982, 1983). Whether such a correlation reflects an encoding of spatial coordinates or of muscle synergies is still an object of debate (Caminiti et al. 1990; Georgopoulos 1991; Mussa-Ivaldi 1988); it appears evident, however, that some high center of the brain such as the motor cortex must represent motor behavior in terms of extrinsic spatial coordinates. Subsequent representation in other coordinates (e.g., joint angles or muscle lengths) may also occur as part of the process of implementing the motor plan. This observation was first made in 1935 by Bernstein, who noted that our ability to control movements is independent of movement scale or location (Bernstein 1967).

If the spatial features of a hand movement are planned and represented by some structure of the CNS then there must be another set of neural processes devoted to transforming the desired hand trajectory into muscle activations. A third step in carrying out a reaching task there-

fore consists in the conversion by the CNS of the desired correction, amplitude, and velocity of movement into signals that control the mechanical action of the muscles. The equilibrium-point hypothesis is related to this third step and the communication between the processes of movement planning and movement execution.

Investigators of motor control have become increasingly aware of the computational complexities in the production of muscle forces. Some have proposed that the CNS derives a motion of the joints from the desired path of the end point (inverse kinematics) and that it then derives the forces to be delivered to the muscles (inverse dynamics; Hollerbach & Atkeson 1987). The idea that the CNS performs these inverse computations implies that it can somehow estimate precisely limb inertias, center of mass, and the moment arm of muscles. Small errors in the estimation of these parameters can result in inappropriate movements. Robotic experience with similar approaches has shown that inertial parameter errors as small as 5% can result in instability (Slotine 1985). Most motor control investigations regard this feedforward computation as rather unrealistic. As an alternative, we and others have proposed a different solution to the inverse dynamics problem: the equilibrium-point hypothesis.

2. Definition of the equilibrium-point hypothesis

The hypothesis was first proposed by Feldman (1966b), who viewed joint posture as an equilibrium resulting from the length-dependent forces generated by agonist-antagonist muscles. A key feature of the equilibrium-point hypothesis is that muscles have springlike behavior. Experimental evidence has indicated that muscles behave like tunable springs in the sense that the force they generate is a function of their length and neural activation level (Matthews 1972; Rack & Westbury 1974). The force-length relationship in individual muscle fibers was studied by Gordon et al. (1966), who related the development of tension at different muscle lengths to the degree of overlap between actin and myosin filaments. This overlap limits the formation of cross-bridges. The increase in muscular stiffness observed when the motoneuronal drive increases is considered a direct consequence of the generation of new cross-bridges.

In 1966, Feldman put forward the idea that the CNS may execute a movement by generating CNS signals that change the relative activation of agonist and antagonist muscles. This change in activation generates joint torques; the resulting joint motion will depend upon the muscle torques and the external loads.

There are at least two variants of the equilibrium-point hypothesis. Feldman (1986) called them the "alpha" and "lambda" models. We would like to stress that the alpha model, which has been attributed to our group, reflects our views only in part. The following discussion describes our interpretation of the two models.

2.1. The alpha model. A central postulate of the alpha model is that the CNS generates a temporal sequence of signals that specify, at all times, an equilibrium position of a limb and the stiffness of the muscles acting on the limb. Although the terminology of the equilibrium-point hypothesis is by now firmly rooted in the literature, the term

equilibrium position is a source of some confusion. We use the term in the following sense: It is the location at which the limb would rest if the centrally generated commands were "frozen" at any given value and the limb were free to move in the absence of external loads or forces. In the presence of static external loads or forces, the actual equilibrium position of the limb will in general differ from this position. Hence we introduced the term *virtual position* to distinguish the two. A time sequence of central commands gives rise to a time sequence of virtual positions, which is called a *virtual trajectory*.

The experimental evidence supporting this view derives from three sets of experiments performed in monkeys (Bizzi et al. 1984). The movements used in these experiments were visually evoked single-joint flexion and extension of the elbow, which lasted approximately 700 msec for a 60-degree amplitude.

The first set of experiments was performed both with intact monkeys and monkeys deprived of sensory feedback. The monkey's forearm was briefly held in its initial position after a target light that indicated final position had been presented. It was found that movements to the target after the forearm was released were consistently faster than control movements in the absence of a holding action.

Figure 1 shows a plot of the initial accelerative transients against the durations of the holding period in the same animal before and after interruption of the nerves conveying sensory information. The time course of the increase in the amplitude of the initial accelerative transient was virtually identical in the two conditions.

The initial acceleration after the release of the forearm increased gradually with the duration of the holding period, reaching a steady-state value no sooner than 400 msec after muscles' activation. These results demonstrated that the CNS has programmed a slow, gradual shift of the equilibrium position instead of a sudden, discontinuous transition to the final position.

The same conclusions were supported by a second set




Figure 1. The forearm of intact and deafferented animals was held in its initial position while the animal attempted to move toward a target light. Then the forearm was released at various times. This figure is a plot of acceleration (immediately following release) versus holding time. The abscissa shows time in milliseconds; the ordinate shows radians per second squared. Solid circles: intact animal; open circles: deafferented animal. (From Bizzi et al. 1984.)




Figure 2. Forearm movements of deafferented monkeys with a holding action in the final position. While the target light remained off, the servo moved the arm to the target position. Then the target light was activated and the servo was turned off. The arm returned to a position intermediate between the initial and target positions before moving back to the target position. Similar results were obtained in many trials in two monkeys. The upper bar indicates duration of servo action. The lower bar indicates onset of the target light. The broad trace shows arm position; the dashed trace shows torque. B: flexor (biceps); T: extensor (triceps). (From Bizzi et al. 1984.)

of experiments (Bizzi et al. 1984) based on forcing the forearm to a target position through an assisting torque pulse applied at the beginning of a visually triggered forearm movement. The goal of the experiments was to move the limb ahead of the postulated equilibrium position with an externally imposed displacement in the direction of the target. It was found that the forearm, after being forced by the assisting pulse to the target position, returned to a point between the initial and final position before moving to the end point. This outcome results from a restoring force that is generated by the springlike muscle properties. If muscles merely generated force, or if the springlike properties were negligible, this return motion of the limb would not have been observed. Again, these experiments were performed in both intact and deafferented monkeys.

In the third set of experiments, performed in monkeys deprived of proprioceptive feedback, the forearm was first driven passively from the initial position to a new position in the absence of a target light and held there for a

variable amount of time (1 to 3 sec), after which a target light at the new position was activated (Figure 2).

A cover prevented the animal from seeing its arm. After the reaction time to the presentation of the target light, the monkey activated its muscles (flexors in the case of Figure 2) to reach the target position even though the elbow was already there. At this point, usually shortly after the onset of muscle activity, the servo that held the arm was deactivated. Once released, the arm moved toward the original starting position to a point intermediate between the initial position and the target position before moving back to the target position. Note that during the return movement requiring extension, evident flexor activity was present. The amplitude of the return movement depended on the duration of the holding action. If enough time elapsed between the activation of the target light and the deactivation of the servo, the arm remained in the target position upon release. These observations provide further support for the view that the CNS specifies a series of equilibrium positions throughout the movement.

The idea of a moving equilibrium point (termed a virtual trajectory) is a direct consequence of two known facts: (1) that a limb is at static equilibrium in the absence of external loads when all the torques generated by opposing muscles cancel out, and (2) that the neural input to each muscle has the effect of selecting a length-tension curve. If the net stiffness due to muscle action on the limb is positive (see below), it follows that at all times the neural activities directed to all the muscles acting on a limb can be "translated" into a corresponding equilibrium angle, which is given by the balance of the springlike torques that keep the limb at rest (Hogan 1988a).

During the execution of a movement, these equilibria or virtual positions act as centers of attraction. The difference between actual and virtual position generates a springlike force directed toward a virtual position. The course of the movement is then determined by the interaction of the springlike force with limb inertia and viscosity and the velocity-based modulation of the muscle length-tension relationship. On this basis, a computer simulation developed by Hogan (1984) successfully reproduced all of the above experimental observations.

It should be stressed that a comparison of forearm movements in intact and deafferented monkeys revealed some quantitative, but no qualitative, differences. Insofar as the deafferented animal could execute movement, sensory feedback was not essential for movement. The major role of feedback in the experiments of Bizzi et al. (1984) may have been to augment the inherent properties of muscles such as stiffness, a role for which a considerable body of experimental evidence (Hoffer & Andreassen 1981) has been presented. It follows that in these highly trained and deafferented animals, the simple task of moving only one joint was executed primarily by a movement program of central origin. It should be noted that Bizzi et al. (1984) studied large arm movements with undemanding requirements on final position accuracy. In addition, these movements were performed at moderate speed. It is conceivable that under these circumstances feedback plays a minor role. In contrast, during the execution of motor tasks demanding greater accuracy, feedback may well play a much more important role, as shown by Sanes and Evarts (1983a) and by Day and

Marsden (1982). It is known that muscle and cutaneous receptors are most sensitive to signals of small amplitude (Matthews & Stein 1969).

In summary, the alpha equilibrium-point model rests upon studies in which fairly large single-joint movements were performed by deafferented animals at moderate speeds. We hasten to add that whereas our animals performed adequately in our restricted condition, there is no question that the performance of deafferented monkeys and humans is far from normal during the execution of multijoint movements. Clearly, sensory input must play a role. The lambda model and its predecessor, Merton's (1953) reflex servo control model, are directed at including sensory input into a motor-control scheme.

2.2. The lambda model. Alpha and lambda models have many common features. For example, both posit a unified description of posture and movement, and both attribute springlike properties to muscles and view movements as being generated by a shift in the equilibrium position of the limb. The main difference between the alpha and lambda models lies in the importance accorded the distinction between feedforward and feedback mechanisms generating an equilibrium position. In the alpha model the torque, T , produced by the muscles about a joint, is expressed as a function of the joint equilibrium angle, θ , and the centrally generated commands. In the deafferented animal, those commands are reflected in the muscle active states, presumably determined by the set of α -motoneuron activities $\{\alpha\}$:

$$T = \psi(\theta, \{\alpha\}). \quad (1)$$

In the absence of external loads or forces, the joint equilibrium angle θ_0 may accordingly be expressed as a function¹ of $\{\alpha\}$, $\theta_0 = g(\{\alpha\})$, provided the stiffness is nonzero. In subsequent work, the same formalism was applied with success (Flash 1987) in a study simulating multijoint arm trajectories. In our view, it is the apparent mechanical behavior of the neuromuscular system, whether due to reflex action or intrinsic muscle behavior, that determines the stability and dynamic behavior of the limbs and how they interact with objects in the environment. Thus, in our formulation we deliberately make no attempt to distinguish between central and reflex effects on the α -activity. In contrast, Feldman (1986) has proposed that the net effect of the central commands impinging upon the α -motoneurons is to modulate the threshold of the stretch reflex. More precisely, in Feldman's lambda model the activity of the α -motoneurons can be expressed as a function

$$\alpha = \alpha(x - \lambda) \quad (2)$$

of a muscle's length, x , and a centrally determined control parameter, λ . This parameter has the same dimension of muscle length and $\alpha(x - \lambda)$ is a threshold function ($\alpha = 0$ if $x \leq \lambda$ and $\alpha > 0$ otherwise). Thus, each muscle force, F , is expressed as a function of the difference between muscle length, x , and a control parameter, λ , that is $F = F(x - \lambda)$. Note that for $x > \lambda$, the function follows an invariant characteristic, and for $x \leq \lambda$ the active force is zero. Feldman and coworkers have suggested using the lambda model to account for the formation of known electromyographic (EMG) patterns (Berkinblit et al. 1986;

Feldman 1986). They suggested that this important goal could be achieved by relating the parameter λ on one hand to the equilibrium position, θ_0 , and on the other to the net motoneuron activation, α .

Our main critique of the lambda model is related to its striking resemblance to Merton's servo reflex control hypothesis. Like the lambda model, Merton's (1953) hypothesis postulated that voluntary movements are initiated and controlled by the CNS as a central modulation of feedback. More precisely, the mismatch between extra- and intrafusal muscle fiber lengths generates an input signal to the alpha motoneurons via the monosynaptic pathways. Merton proposed considering this input as a length-error signal that the CNS specifies via the gamma system. Therefore, according to this hypothesis, movements are initiated and controlled by a specification of gamma fusimotor activity.

It is well known that the experimental evidence does not support Merton's hypothesis for the following reasons: (1) There is no gamma lead with respect to alpha activity, (2) the gain of the stretch reflex during movement is too small, and (3) deafferented animals can perform goal-directed movements, albeit in a clumsy way. The lambda model avoids the first criticism by postulating an (unspecified) relation between the centrally determined control parameter, λ , and the α - and γ -motoneuron activities. In this respect the λ model resembles the modification of Merton's hypothesis known as the servo-assistance hypothesis (Granit 1970; Stein 1974). The lambda model avoids the second criticism on the grounds that a significant stretch-reflex gain is not necessary for postural stability. However, in that case the relative contribution of reflexes to the expression of movement would presumably also be insignificant. The third criticism still applies to the lambda model.

Like Feldman, we believe that reflex activity may modify supraspinal commands. We part company, however, on the question of the relative contribution of reflexes to the expression of movement. The experimental work of Bizzi et al. (1978) investigated the contribution of reflex mechanisms in generating the forces produced by the neck muscles when loads were unexpectedly applied during centrally programmed head movements in monkeys. The results indicated that the compensating torque of reflex origin was less than 10% to 30% of that required for perfect compensation. Similar results were obtained by Valbo (1973a), Grillner (1972), and Allum (1975). The conclusion from these experiments is that the reflex apparatus contributes in a modest way to force generation. These observations may be consistent with the fact that feedback-loop delays may cause instability if reflex gains are too high (Hogan et al. 1987).

Regarding the ability of deafferented animals to perform goal-directed movements, Feldman has suggested that in such a pathological condition the CNS replaces the lambda control with a coactivation strategy (Feldman 1986). This could be the case. We would like to point out, however, that even such a cocontraction strategy would require the remarkable ability of the CNS to control and coordinate directly a complex pattern of alpha activations. The available evidence suggests that these activations are correctly mapped by the deafferented animal into a desired equilibrium location, as suggested by the alpha model.

2.3. The alpha/lambda debate. Feldman (1986) has suggested that the alpha and lambda models are mutually exclusive.² In contrast, we believe that these models are mutually compatible and also that one, the lambda model, is a subset (or constrained version) of the other, the alpha model.

To illustrate this point let us consider a single-joint system. The torque-angle function of this system is given by Equation 1 in the alpha model. In the lambda model, at a single joint, each muscle's length is unique demonstrated by the joint angle and each muscle makes a specific contribution to the joint angle and a set of command parameters $\{\lambda\}$:

$$T = \phi(\theta, \{\lambda\}) \quad (3)$$

As was the case with the alpha model, provided the joint stiffness is nonzero, the equilibrium condition, $\phi(\theta_0, \{\lambda\}) = 0$, defines a map from $\{\lambda\}$ to the equilibrium position, θ_0 , $\theta_0 = f(\{\lambda\})$.

However, the specific form assumed for the dependence on λ , a threshold function, means that for a single muscle there exists a range of joint angles (corresponding to the muscle lengths $x \leq \lambda$) for which that muscle stiffness is nearly zero. The required nonzero joint stiffness may be achieved by cocontraction of opposing muscles across the joint, just as in the alpha model.

The lambda model is related to the alpha model by a set of equations (Equation 2) that establish a dependence of the alpha signal upon the muscle length, x , and the centrally generated command, λ . Given the relation between muscle length and joint angle, we may rewrite these equations as

$$\{\alpha\} = \{\alpha(\theta - \lambda)\} \quad (4)$$

The equilibrium condition ($T = 0$) applied to the alpha model provides a map g from the set $\{\alpha\}$ to the equilibrium angle θ_0 . The same equilibrium condition applied to the lambda model provides a different map f from the set $\{\lambda\}$ to the equilibrium angle. In set-theoretic notation, the first map is a set of ordered pairs, $g = \{(\theta_0, \{\alpha\})\}$, and the second is a set of pairs, $f = \{(\theta_0, \{\lambda\})\}$. It is easy to show³ that under Feldman's conditions, f is a subset of g . The converse is not true. For example, unlike the lambda model, the alpha model does not require an invariant characteristic for the form of the torque-angle relation. In this case, there are elements of g that do not belong to f .

In summary, we believe that the alpha and lambda models are by no means mutually exclusive. Each model tends to direct attention toward a different aspect of motor control. The lambda model has been primarily applied to the explanation of EMG patterns. In contrast, we have been more concerned with the mapping of different motor behaviors – such as posture, movement, and contact – into the corresponding sets of equilibrium positions. In this regard, the fact that reflexes may induce a restriction of the alpha model may be of significant value.

We would like to reiterate that we feel uncomfortable being cast in the role of advocating a pure alpha model which we have never proposed. The experimental evidence reported in Bizzi et al. (1976; 1982; 1984), Polit and Bizzi (1978), and in Hogan (1982; 1984) was directed at establishing whether the CNS adopts a final-position control-strategy in order to generate arm trajectories in animals performing highly practiced forearm move-

ments. The results indicate that the transition from the initial to the final position was implemented by a gradual shift in the control signal establishing both a trajectory and a final equilibrium condition. To make the observation we used deafferented animals. However, our experiments were not intended to ascertain the relative contribution in the intact animal of feedforward commands versus feedback signals.

3. Multijoint posture and movement

The study of multijoint arm movement presents problems radically different from those of the single joint. In a multijoint situation, if a displacement is externally imposed on the hand, the amount of stretch experienced by the muscle depends not only upon the amplitude of the perturbation, but also upon its direction. Then, a single number is no longer sufficient to describe the force-displacement relation. This relation is expressed by a matrix whose elements characterize the ratio of each component of the restoring force vector to each component of the applied displacement vector.

To deal with the more complex situation of multijoint arm movement, a specific experimental approach to the study of posture was developed (see Mussa-Ivaldi et al. 1985). This approach was based on measuring the net springlike behavior of the multijoint arm by displacing the hand in several directions from rest and measuring the restoring forces, $F = (F_x, F_y)$. Because the displacements of the hand, $\delta r = (\delta x, \delta y)$, were small enough to justify neglecting higher order effects, a linear relation of the following form was assumed:

$$F_x = K_{xx} \delta x + K_{xy} \delta y$$

$$F_y = K_{yx} \delta x + K_{yy} \delta y. \quad (5)$$

Then, by measuring forces and displacements in different directions, it was possible to estimate the K coefficients from a linear regression applied independently to both of the above expressions. These coefficients could be represented by a single entity – a table, or matrix, expressing the multidimensional stiffness of the hand:

$$K = \begin{bmatrix} K_{xx} & K_{xy} \\ K_{yx} & K_{yy} \end{bmatrix} \quad (6)$$

With this notation, Equation 5 assumes a more compact form, $F = K \delta r$.

We first determined whether the behavior of the multijoint arm was in fact springlike. The curl of the force field $F = F(\delta r)$ must be zero for a springlike system (Hogan 1984). In terms of the stiffness matrix, the off-diagonal coefficients K_{yx} and K_{xy} (which were measured independently) would have to be identical. Our measurements showed that this was the case.

Because the curl of the force field was zero, the hand stiffness matrix K was symmetric and could be represented as an ellipse characterized by three parameters: magnitude (the total area derived from the determinant of the stiffness matrix); orientation (the direction of maximum stiffness); and shape (the ratio between maximum and minimum stiffness). The hand stiffness was actually estimated by Mussa-Ivaldi et al. (1985). Subjects were asked to maintain the hand at a set of workspace locations. At each location, K was derived from measured force and

displacement vectors as outlined in Equation 1. The corresponding stiffness ellipses captured the main geometrical features of the springlike-force field associated with a given hand posture and provided an understanding of how the arm interacts with the environment.

To sum up, the experimental evidence indicates that an equilibrium position of the hand is established by the coordinated interaction of spring like forces generated by the arm muscles (Mussa-Ivaldi et al. 1985). According to the virtual-trajectory hypothesis, which was tested first in the context of single-joint movements (Bizzi et al. 1984), the multijoint arm trajectories are achieved as the CNS gradually shifts the centrally determined virtual position between the initial and final positions. In this control scheme the hand tracks its virtual equilibrium point and torque need not be computed explicitly.

Evidence supporting this equilibrium-point hypothesis in the context of multijoint hand movements was obtained by combining observations of hand movements with computer simulation studies. A model developed by Flash (1987) has successfully captured the kinematic features of measured planar-arm trajectories. In the simulation, Flash made the assumption that the hand's virtual trajectories (but not necessarily the actual trajectories) are invariantly straight. In addition, she assumed that each virtual trajectory has a unimodal velocity profile, regardless of the target locations in the workspace.

The arm dynamics were simulated by obtaining torques derived from the difference between actual and virtual positions multiplied by the stiffness (Flash 1987). It must be stressed that the stiffness parameters used in the simulation of movements were derived from experimentally measured postural stiffness values. The results of the simulation showed that with straight virtual trajectories, the actual movements were slightly curved. Moreover, the direction of curvature, in different workspace locations and with different movement directions, was in good agreement with the experimentally observed movements. This result suggests that during movement planning, the CNS ignores the inertial and viscous properties of the arm and directly translates the desired trajectory into a sequence of equilibrium positions. Therefore, when the movement is executed, the inertial and viscous forces act as perturbations, causing deviations of the actual path with respect to the planned path.

The success of the simulation in capturing the kinematic details of measured arm movements is important as a step toward providing us with a framework for understanding the CNS's trajectory formation in the multijoint context. This work indicates a planning strategy whereby the motor controller may avoid complex computations such as the solution of the inverse dynamics problem. Recent findings by Flanagan et al. (1990) agree partially with Flash's simulations. Flanagan's results support the notion that multijoint movements are planned in end-point coordinates. In contrast with Flash, their results suggest that the equilibrium position is shifted at a constant velocity.

It should be noted that the simulations described by Flash were for relatively slow movements. The question whether fast movements could be achieved in the same way is dealt with in section 7.2. It is also important to note that the good agreement between simulated and experimental trajectories was contingent upon using stiffness

fields whose shape and orientation were identical with those recorded under static conditions by Flash; any change in these two parameters led to substantially different simulated trajectories. Hence, a question left unanswered by Flash's experiments was whether the neural signals to the muscles involved in the execution of natural movements could significantly alter the shape and orientation of the stiffness field, thus undermining the significance of the simulation.

An answer to the latter question was provided by Mussa-Ivaldi et al. (1987; see also Bizzi & Mussa-Ivaldi 1990). Mussa-Ivaldi et al. (1987) found that the shape and orientation of the stiffness field did not change when externally imposed disturbances acting in different directions were applied to the hand. These disturbances generated large shifts in the EMG activation of arm muscles but failed to modify these two parameters. The results of Mussa-Ivaldi et al. provided the evidence for assuming that these parameters may not change when the hand moves through the locations at which the field was measured.

Flash's simulation showed that her model can generate multijoint arm trajectories. The experiment of McKeon et al. (1984), which complement Flash's results (1987), provided qualitative evidence supporting the equilibrium-point hypothesis in the context of human multijoint movements. McKeon et al. (1984) asked subjects to perform pointing movements between two targets while gripping the handle of a two-link manipulandum. A clutch mounted on the inner joint of the manipulandum was used to brake the inner link under computer control. Because the clutch was activated randomly at the onset of a movement, the hand trajectory was restricted to a circular path with a radius equal to the length of the outer link of the manipulandum. While the clutch was engaged, the force exerted on the handle was always strongly oriented so as to restore the hand to the unconstrained path and not to the end point of the path. This result is in accordance with the equilibrium-point trajectory: The muscle's springlike properties and the proprioceptive reflexes generate forces attracting the hand toward the original path.

4. Control of contact tasks

The equilibrium-point hypothesis also provides a simple but highly effective way to solve the much more complex problem of contact control tasks in which the limb must exert force stably and do work on objects in the environment. In general, the manipulated object can have its own dynamic behavior, which may be arbitrarily complex. Because of the mechanical interaction, that dynamic behavior is added to the already complex dynamic behavior of the limb.

The ability to control contact with objects is clearly a fundamental prerequisite for the use of tools, one of the distinctive features of human behavior. [See also Parker & Gibson: "A Developmental Model for the Evolution of Language and Intelligence in Early Hominids" *BBS* 1 (3) 1979; Chevalier-Skolnikoff: "Spontaneous Tool Use and Sensorimotor Intelligence in *Cebus* Compared With Other Monkeys and Apes" *BBS* 12 (3) 1989; MacNeilage et al.: "Primate Handedness Reconsidered" *BBS* 10 (2) 1987; and Greenfield: "Language, Tools and Brain" *BBS*

14 (4) 1991.] The subtlety and difficulty of this problem is disguised by the ease with which humans manipulate objects. Experience with robots has repeatedly shown that even the apparently trivial problem of controlling the force exerted on a surface has proven surprisingly difficult. In robotics, an "obvious" approach is to measure the force of contact and send that information to the controlling computer so that it can adjust or regulate the force exerted. Unfortunately, this approach has been plagued by a phenomenon called *contact instability*. Robotics researchers in numerous laboratories have reported that a robot capable of executing unrestrained motions stably and accurately will break into a pathologically uncontrollable chattering instability upon contact with a rigid surface, bouncing off the surface and bumping it repeatedly. This problem has been identified as one of the prominent challenges of robotics (Paul 1987). Yet biological systems clearly have little difficulty contacting and manipulating objects.

The necessary and sufficient condition for a manipulator to remain stable when coupled to an arbitrarily complex passive object has recently been derived mathematically and verified experimentally. Details are provided in Colgate (1988; Colgate & Hogan 1988; Hogan 1988b). The essence of the result is that an arbitrary collection of passive objects such as springs and masses can temporarily store energy, but cannot supply it indefinitely. Consequently, if there is nonzero dissipation associated with the motion of the system it will converge on a stable state of minimum energy, a result first proved by Lord Kelvin.

In contrast, a typical actuator (e.g., a robot motor or a muscle) can supply energy indefinitely (or at least over time scales that are long compared to the characteristic dynamic behavior of the system they act upon; i.e., muscle may continuously supply energy for far longer than the duration of a typical voluntary movement). If that energy supply is improperly controlled (e.g., so that the energy supplied exceeds that dissipated), unstable behavior may result.

However, if the actuator control system is designed so that the apparent behavior of the actuator is that of an object that can only temporarily store energy (e.g., a spring), then Lord Kelvin's result is recovered: That actuator can be connected to an arbitrary collection of springs and masses and the combined system will be stable. This is precisely what is achieved by the springlike behavior of single muscles or agonist-antagonist muscle groups about a single joint.

In the multijoint case, new factors arise because of the possibility of complex interactions between joints. If the off-diagonal terms in the stiffness matrix were not identical, $K_{xy} \neq K_{yx}$, it would imply that energy could be supplied indefinitely by perturbing the hand so as to make small circular motions about the equilibrium point. It is therefore highly significant that our psychophysical experiments on human subjects have established that the entire multijoint upper limb mimics the behavior of a passive, multijoint spring (Mussa-Ivaldi et al. 1985) even though that requires finely balanced interjoint feedback (Hogan 1985a).

The relation between the properties of the hand stiffness and the equilibrium-point hypothesis may be summarized as follows: In the single-joint case, in order to

define a map relating central commands to a virtual position the stiffness must be nonzero. In the multijoint case, the corresponding requirement is that the stiffness matrix, K , must be nonsingular (determinant $(K) \neq 0$). For static stability about that equilibrium point, the stiffness matrix must be positive definite, a stronger requirement. Passive springlike behavior adds a further requirement: The stiffness matrix must be symmetric.

To establish fully stability or passivity of the arm, its dynamic behavior must be considered in addition to its stiffness; measurements of the arm's dynamic response to perturbation (its mechanical impedance) are required. In the absence of that information we conclude that our observation of multijoint springlike arm behavior is consistent with the theoretical requirements to preserve stability on contact with passive objects, although a stronger conclusion would require a more thorough analysis of dynamic behavior.

The way equilibrium-point control may be used in contact tasks is illustrated conceptually in Figure 3. Although the real-life situation may be more complex, the basic mechanics are sufficiently similar to illustrate the concept. Figure 3A depicts a hand being moved downward to contact a surface and push on it. The hand is assumed to be controlled so that a simple relation between its force and its position is maintained. For simplicity, the force, F , is assumed to be proportional to the separation of the virtual position, X_v , and the actual position, X , of the hand:

$$F = K(X_v - X) \quad (7)$$

where K is a constant, the stiffness at the hand.




Figure 3. Diagram illustrating how virtual position may be used to control contact tasks. A: The hand is controlled by a relation between force exerted, F , and the difference between a virtual position, X_v , and the actual hand position, X . B: A trajectory of the virtual position, X_v , that could be used to move to a surface and exert force on it. C: While the hand is not in contact, its actual position, X , tends to follow the virtual position. On contact, the position of the hand is constrained by the surface. D: While the hand is not in contact, the force exerted by the hand, F , is constrained to zero. On contact, the force is proportional to the distance the virtual position moves "into" the surface.

Next assume that the virtual position, X_v , is moved slowly and steadily downward, as shown in Figure 3B. As long as the hand is clear of the surface, the force exerted is zero and the actual position of the hand will approximately track the equilibrium position, as depicted in Figure 3C. As soon as the hand contacts the surface, the actual position can no longer follow the virtual one because of the physical constraint imposed by the surface. But once the hand is in contact with the surface, the problem is to control the force exerted. This is easily accomplished by controlling the depth the equilibrium position "penetrates" into the surface. The fact that it may penetrate into a surface is the main reason the equilibrium position defined by the action of neuromuscular system is called *virtual*. Figure 3D shows that as the virtual position moves deeper into the surface, the force rises in the proportion required by the relation between force and position.

This simple example illustrates that *the same process* that can simplify the computational burden of controlling unrestrained motions can also be an effective way to control the force exerted on an object. The idea is a generalization of the unified description of posture and movement outlined above. If the response of the limb to perturbations is appropriate, the equilibrium position, which may simplify motion control, is also an effective way to control interaction. During interactive tasks, the behavior of the limb is dominated by the dynamic response of the skeletal and neuromuscular system to perturbation. Functional interactive behavior such as tool use can therefore be controlled by modulating that dynamic response, the mechanical impedance of the neuromuscular system (Hogan 1985b).

5. Neurophysiological basis of the equilibrium-point hypothesis

Until recently, the equilibrium-point hypothesis rested mostly on data derived from psychophysical and behavioral experiments. Recently, Mussa-Ivaldi et al. (1990), Giszter et al. (1991c; 1992b) and Bizzi et al. (1991) addressed directly the problem of providing a neurophysiological underpinning for the equilibrium-point hypothesis. To this end, they investigated the characteristics of the spinal circuitry involved in producing coordinated multijoint leg movements and postures. For these studies they used the spinal frog.

The spinal frog is a simplified preparation that retains significant multijoint motor abilities (Berkinblit et al. 1986; Fukson et al. 1980; Giszter et al. 1989). It is well known, for example, that the spinal frog is capable of generating a coordinated sequence of multijoint hindlimb movements to remove a noxious stimulus from the skin. This "wiping reflex" requires complex information processing. Thus, the spinal cord must contain enough circuitry to coordinate the motion of multiple limb segments.

One possible approach to understanding the motor behavior of a spinalized frog consists of postulating that a noxious stimulus on the skin triggers some form of an "inverse-dynamics" computation within the spinal cord. This computation must ultimately generate a coordinated pattern of joint torques in the hindlimb. In contrast,

according to the equilibrium-point hypothesis, the motion of the hindlimb is generated by the development of neural patterns that specify a sequence of equilibrium points with the limb's workspace. In support of the latter view, evidence for extensive cocontraction during flexion and wiping has indeed been found (Schotland et al. 1989).

We have addressed these different hypotheses in experiments in which we microstimulated the gray matter of spinalized frogs. According to the view that favors inverse dynamics, the activation of a region in the spinal gray matter is expected to generate a timed pattern of joint torques. These torques need not define an equilibrium point within the workspace.

Alternatively, the equilibrium-point hypothesis predicts that we should be able to induce a stable equilibrium of the leg within its range of action by activating the spinal gray matter. The equilibrium-point hypothesis also implies that the development of neural patterns corresponds to a movement of the equilibrium point.

Mussa-Ivaldi et al. (1990), Giszter et al. (1992a; 1992b), and Bizzi et al. (1991) elicited motor responses by microstimulating the spinal gray matter in a region located from the base of the dorsal horn to the upper ventral horn. An important methodological feature of these experiments involved measuring the x and y force components at the ankle with a 6-axis force transducer. The x - y plane corresponded approximately to the horizontal plane, where most of the wiping movements tended to occur.

In these studies, the electrical stimulation of the spinal gray matter at threshold levels for movement always coactivated groups of muscles. In order to record the forces generated by the leg, a two-part procedure in a single recording session was followed: First, the frog's ankle was placed at one location in the leg's horizontal workspace (i.e., that region of the horizontal plane that can be reached by the ankle). Second, the direction and amplitude of the force at the ankle elicited by stimulating a site in the spinal cord were recorded. While stimulating the same site, the investigators in these studies repeated this procedure with the ankle placed at each of 14 locations covering the whole range of the workspace. The force vectors that were recorded varied in direction and amplitude as the experimenters placed the leg at different workspace locations. The measured force vectors were used to estimate the force field in a large region of the ankle's workspace. Remarkably, in most instances, the spatial distribution of these vectors resulted in a field characterized by a single equilibrium point (i.e., a point at which the amplitude of the F_x , F_y force components was zero). The fields, with their associated equilibria, were found to be distributed in several locations throughout the leg's workspace.

Note that the type of field shown in Figure 4A does not result merely from the mechanical properties of the musculoskeletal system. A radically different pattern of forces was found when the stimulating electrode was placed in the ventral roots or within gray matter regions populated by the motoneurons (Figure 4B). In this case, the structure of the field was often characterized by forces that were parallel or divergent, with no detectable equilibrium point. The striking differences obtained after stimulation in regions consisting predominantly of interneurons compared to those with predominantly motoneurons indicated that the stimulation and activation of




Figure 4. Force fields. A: The force field obtained from the quiescent frog before the stimulation of the cord. The equilibrium of the force field is indicated by the filled circle. B: The alteration of the force field as a result of stimulating a site in the spinal cord grey matter in the lateral interneuron region. C: A force field resulting from stimulating a different spinal cord site (1 mm rostral to the site that generated the field shown in B) but still in the lateral interneuron region. Clearly, the equilibria lie in different locations. D: A force field resulting from stimulating motor fibers (8th ventral root) with the same currents used in the stimulation of the lateral interneuron region. Clearly, no equilibrium is present. We obtained similar results when the electrode was placed among the motoneurons.

the spinal cord's premotoneuronal network caused a balanced recruitment of motoneurons that imposed a structure on the forces generated by the limb muscles.

6. Temporal evolution of the force fields

After the delivery of a stimulus to the spinal gray matter, each measured force component changed with time. Consequently, the force field as a whole changed with time. The dependence of the force field on time is captured graphically by a sequence of "frames" (Figure 5). Each frame shows the force field measured at a given latency from the onset of the stimulus.

The first frame (latency = 0) shows the resting force field, that is, the field as it was before the stimulus had produced any mechanical effect. The subsequent frames are separated by intervals of 86 msec. They show the effect of the stimulus as a smooth change in the overall pattern of forces. In several instances we have observed the following sequence of events (as indicated in Figure 5):

(1) After a brief delay from the onset of the stimulus (about 50 msec), the pattern of forces began to change and the equilibrium position started to "move" in a given direction (Figure 5, frames 1 to 3).




Figure 5. Temporal sequence of fields following the stimulation of a spinal site. The six frames are ordered by increasing latency from the stimulus and are separated by intervals of 86 msec. The filled circle indicates the equilibrium point.

(2) Then (Figure 5, frame 4), the equilibrium position reached a point of maximum displacement within the workspace. This point was maintained for a time interval that depended on the stimulation parameters (current, train duration, etc.). At the same time, the field forces reached a maximum amplitude around the equilibrium position, corresponding to a maximum in end-point stiffness.

(3) Finally (Figure 5, frames 5 & 6), the force vectors started to decrease in amplitude and rotate toward their original directions. At the same time, the equilibrium point moved back to the resting location. This sequence of static equilibria is by definition an "equilibrium trajectory": As the neuromuscular activity changes gradually in time, the equilibrium undergoes a gradual shift. Furthermore, after the EMG activities have returned to their resting value, the equilibrium returns to the resting location.

Summing up, these results show that the stimulation of the upper and middle layer of the spinal cord in conjunction with the positioning of the leg in different workspace locations produces a force field with a single equilibrium point. The equilibrium point represents a locus at which the leg would be at steady state. Further experiments are in progress to relate natural activation patterns to the data obtained from microstimulation of spinal-cord circuits. Basically, the neural signals originating from the spinal cord specify directional forces converging onto a location in the leg's workspace. Not surprisingly, perhaps, these results imply that natural activation of a group of spinal-cord premotor circuits also specifies the limb's final position and configuration.

7. Conclusions: Strengths and weaknesses of the equilibrium-point hypothesis

7.1. Strengths. The equilibrium-point hypothesis is strongly rooted in the biomechanics of muscles and in particular in their tunable springlike behavior (Hogan

1984; Rack & Westbury 1969): The isometric force generated by a muscle depends on the level of neuromuscular activity as well as the length of the muscle. In other words, the state of activation of a muscle does not determine tension alone but a whole length-tension curve.

It is significant that the springlike behavior of muscles is in conflict with the traditional engineering notion of an "ideal actuator." According to this notion, the output of an ideal actuator – for example, the torque produced by a torque motor – should be independent of the operating state (e.g., position and velocity). This requirement is analogous to the specification of an ideal voltage source in an electrical circuit. We believe there is a deep relationship between the characteristics of the actuators and the control and computational problems that have been central to research in robotics.

A clear example of the relation between actuator characteristics and control is the problem of contact instability: Measurements to date indicate that the dynamic behavior of the neuromuscular system has precisely those properties that are required to guarantee that contact with passive objects in the environment cannot induce instability. Thus, the apparent characteristics of the actuators circumvent a problem that has plagued the application of robots to contact tasks.

An example of the relation between actuator characteristics and computation is the inverse-dynamics problem: To make this computation, one finds the torque that must be applied to the joints to produce a desired motion when the inertial parameters of a manipulator are given. The formulation of this problem implicitly assumes the availability of either an ideal force generator or the computational machinery for translating the output of the inverse dynamics computation (a torque) into a motor command. From this perspective, the biological muscle would be a poor actuator.

The above argument can be reversed, however: Because the biological actuators are springlike, the inverse-dynamics problem does not need to be solved. In fact, according to the equilibrium-point hypothesis, the CNS can express the desired trajectory of a limb directly as a sequence of equilibrium positions. Then the muscles' springlike properties transform the difference between the actual and the desired position of the limb into a springlike restoring force. The actual motions that result are inexact but are produced without computing any dynamics. Consequently, there is no need to postulate neural structures to perform these complex computations.

Of course, the equilibrium-point hypothesis does not eliminate all computational problems: a pattern of neural activity may define a virtual trajectory, but there remains the formidable problem of how to select an appropriate pattern of neural activation to produce a desired virtual trajectory. Nevertheless, because it is based only on the static characteristics of muscles and their reflex connections and requires no knowledge of the dynamic parameters of the limbs (e.g., the inertias), this problem is significantly simpler than the direct computation of muscle forces or joint torques.

7.2. Weaknesses. Flash's simulation work indicates that the equilibrium-point hypothesis is indeed adequate for generating movements at moderate speeds (Flash 1987).

To move faster along the same path, either of two approaches may be taken: Increase the stiffness of the muscles or modify the equilibrium-point trajectory. With respect to the first option, note that the stiffness must be scaled with the square of the speed. That is, to move twice as fast, the muscle system must be four times as stiff. The question, then, is whether such levels of stiffness are biologically plausible.

The stiffness of the human arm has been estimated by measuring the response to disturbances applied during the movement. Stiffness values in the range of 4.0 to 36.0 newton-meters/radians have been recorded. These values are much lower than the theoretical value necessary to produce very fast forearm movements. The work reported by Bennett et al. (1989; Bennett 1990), Xu et al. (1989; 1990a; 1990b; 1991) and Lanman et al. (1978) identifies muscle stiffness during movement. The results indicate that the stiffness of the arm drops with the start of the movement and rises just before reaching the final position. Clearly, these findings do not support the idea that faster movements are achieved by increasing muscle stiffness.

It is possible that the speed of movement could be increased by another method – changing the equilibrium-point trajectory. Simulations of fast movements indicate that the equilibrium-point trajectory must lead initially and then lag behind the actual desired position during the course of the movement (see, e.g., Hogan 1984). The equilibrium position may actually overshoot the final desired position. Although this formulation of the model increases the efficiency of the system, it also increases the computational complexity of the problem. The equilibrium-point trajectory required to produce the movement is no longer simply a copy of the desired movement. The attractive computational simplicity of equilibrium-point control is therefore lost in the latter approach.

In summary, neither an increase in stiffness nor a modification of the equilibrium-point trajectory seems to be a biologically plausible mechanism for generating fast movements. Recently, a third option has been explored by McIntyre (1988; 1990; McIntyre & Bizzi 1992) who has developed a model competent to deal with the generation of faster movements in the context of the equilibrium-point hypothesis. McIntyre's model is basically a simple combination of the alpha and lambda hypotheses. The alpha command to the muscles is obtained from the sum of three elements: (1) a centrally defined signal, representing a desired equilibrium-point trajectory (position and velocity), (2) a position feedback signal, and (3) a velocity feedback signal. Consistent with the alpha model, the desired position signal is used as a feedforward component of the alpha activity. Consistent with the lambda model and the alpha-gamma coactivation, the same desired position cue is used as a reference signal in the position feedback loop. Furthermore, the desired velocity is used as a reference signal for the velocity feedback loop.

In equilibrium-point terms, the combination of feedback and feedforward signals can be regarded as follows: The feedforward signal (1) specifies a sequence of equilibrium points that corresponds to the desired trajectory. In the absence of feedback, this signal would be sufficient to drive the limb toward the final desired location, albeit

in a sluggish way because of the limit on the achievable limb stiffness. Fortunately, feedback signals (2) and (3) come to the rescue by modifying effectively the equilibrium trajectory generated by the alpha activity. The addition of these feedback signals implements a position and a derivative controller that serve to improve the performance of the system. Remarkably, the suppression of feedback would impair but not abolish motor performance, a fact that has been observed in the deafferented preparation and not accounted for by the original lambda model. At the same time, this feedback signal allows the motor system to produce faster movements at a given level of stiffness while retaining the simplicity of an equilibrium trajectory specification at the input. Computer simulations testing such a control scheme have shown that the system is stable and produces fast movements at stiffness levels below those required by the equilibrium-point hypothesis. The movement's speed and velocity profiles are comparable to those achieved by a human subject asked to move "as fast as possible."

One major weakness of the equilibrium-point hypothesis is that it is difficult to test. The central concept is that posture and movement are subserved by the same processes. Static stability is arguably one of the defining requirements of posture; consequently, the equilibrium-point hypothesis makes the assumption that during movement as well as posture the limbs exhibit stability. Note that this is not a requirement for the motion of a mechanical system. Nor is it a fundamental requirement for a biological system, although it is physiologically plausible, given the known springlike behavior of muscles and their reflex connections.

The essence of the equilibrium-point hypothesis is that centrally planned motor intentions are expressed and transmitted to the periphery using the virtual trajectory. Evidence in support of this hypothesis may be derived from observations of simple patterns in the virtual trajectories underlying observed behavior (e.g., Flash's [1987] observation that the same straight virtual trajectory could give rise to a wide range of different reaching movements). The major drawbacks of this approach are the difficulty of arriving at a concise definition of "simplicity" and the difficulty of measuring limb stiffness under relevant conditions.

Much of the difficulty of arriving at a confirmation (or disconfirmation) of the equilibrium-point hypothesis stems from the problem of defining a (perhaps artificial) boundary between central and peripheral processes. Where should one look for a neural expression of the virtual trajectory?

Except in the deafferented case, it is of little value to define the boundary at the level of alpha motoneurons. Defining it at the level of commands descending into the

spinal motoneuron pools seems more reasonable, and our recent studies of spinal frogs support the idea of a virtual trajectory expressed in the collective activity of spinal motoneuron pools. That activity, however, is modified by long loop reflexes.

The theory that motor intentions are expressed and transmitted to the periphery using the virtual trajectory has direct implications for studies of cell discharge in the brain. The important point is that according to the theory, neither the forces generated by the muscles nor the actual motions of the limbs are explicitly computed; they arise from the interplay between the virtual trajectory and the neuromuscular mechanics. Hence, neither the forces nor the motions need be explicitly represented in the brain. If this theory is correct, then cell discharge studies (e.g., Cheney & Fetz 1980; Evars et al. 1983; Georgopoulos et al. 1983; 1982; Kalaska et al. 1983) might be better interpreted in terms of virtual trajectories and neuromuscular stiffness (or, more generally, impedance) than in terms of forces or motions (see especially Humphrey & Reed 1983).

ACKNOWLEDGMENTS

Research supported by Office of Naval Research Grant N00014/88/K/0372, National Institutes of Health Research Grants NS09343, AR26710, and AR40029 and National Science Foundation Grant 8914032-BCS.

NOTES

1. This is a consequence of a fundamental theorem on implicit functions. According to this theorem the equilibrium condition, $\psi(\theta_0, \alpha) = 0$, defines a unique map from the set $\{\alpha\}$ to θ_0 , $\theta_0 = g(\alpha)$, provided that the joint stiffness $\partial\psi/\partial\theta$ is different from zero. If the equilibrium position is also to be stable, it is necessary (though not sufficient) that the stiffness be positive. In this respect the equilibrium-point hypothesis can be regarded as a way to represent a high-dimensional control variable, the set $\{\alpha\}$, by means of another one, θ_0 , which has the same dimension as the variable, θ , which describes the actual movement.

2. Part of Feldman's critique is addressed to a view of the alpha model that we can hardly share. For example, Feldman (1986) attributes to the alpha model the notion (1) that the muscle activation only modulates the muscle's stiffness but not its rest-length, and (2) that in an intact preparation, the level of alpha activity does not depend on feedback variables. Clearly, both statements fly in the face of physiological common sense. Endorsing them would merely transform the alpha model into a straw-man hypothesis whose refutation could not add anything to our knowledge of motor control.

3. The proof is as follows. Let $(\theta_0^*, \{\lambda^*\}) \in f$ be a pair satisfying the lambda-equilibrium condition $\Phi(\theta_0^*, \{\lambda^*\}) = 0$; then the pair $(\theta_0^*, \{\alpha^*\})$ with $\{\alpha\} = \{\alpha\theta_0^* - \lambda^*\}$ satisfies the alpha-equilibrium condition $\psi(\theta_0^*, \{\alpha^*\}) = 0$. That is $(\theta_0^*, \{\alpha^*\}) \in g$.