In:

Self-organization, Computational Maps, and Motor Control

pp. 383-421

Pietro G. Morasso
and
Vittorio Sanguineti
(eds.)

1997

Elsevier -- North Holland

Psychology Series

NEURAL MODELS FOR FLEXIBLE CONTROL OF REDUNDANT SYSTEMS

Frank H. Guenther and Daniele Micci Barreca[1]
Department of Cognitive and Neural Systems, Boston University
677 Beacon St., Boston, MA 02215 (USA)

Abstract

This chapter discusses the explanation of a class of human motor equivalence competencies put forth by the DIVA and DIRECT models of motor skill acquisition and performance. It is suggested that experimental data indicating approximate postural invariance for reaches do not imply that the motor system is utilizing postural targets. Instead, an inverse kinematics transformation utilizing a directional mapping with a "postural relaxation" component is shown to be consistent with these data while also providing motor equivalent capabilities not possessed by models that use postural targets. This transformation is related to robotics techniques utilizing a Jacobian pseudoinverse and to the motor control models of Cruse and colleagues. A self-organizing neural network architecture that learns such a directional mapping is presented, including simulations verifying its ability to explain the approximate postural invariance seen in the experimental data. Side effects of the model's learning process suggest two sources that may contribute to the gentle curvature seen in human reaches: a bias toward movements along the long axis of the manipulability ellipsoid, and a tendency toward more comfortable postures.

1 Introduction: Motor equivalence and redundancy

Motor equivalence is the ability to carry out a task using different motor means. For example, people are capable of producing written letters with very similar shapes using their wrist and fingers or shoulder and elbow

1. Frank Guenther is supported in part by the Alfred P. Sloan Foundation, the National Institutes of Health (1 R29 DC02852-01), and the Office of Naval Research (ONR N00014-95-1-040). Daniele Micci Barreca is supported by the Italian Council for Research.

(Merton, 1972), their dominant or non-dominant arms (Raibert, 1977; Wright, 1990), and even using pens attached to their feet or held in their teeth (Raibert, 1977). Motor equivalence is seen in a wide variety of human behaviors, including handwriting, reaching (e.g., Cruse, Brüwer, and Dean, 1993), and speaking (e.g., Abbs and Gracco, 1984; Lindblom, Lubker, and Gay, 1979; Savariaux, Perrier, and Orliaguet, 1995), and in a wide variety of species, including turtles (Stein, Mortin, and Robertson, 1986) and frogs (Berkinblit, Gelfand, and Feldman, 1986). The ubiquity of motor equivalence is no doubt the evolutionary result of its utility: animals capable of using different motor means to carry out a task under different environmental conditions have a tremendous advantage over those that cannot.

This chapter describes self-organizing neural network models that address a subset of motor equivalent behavior: the ability to use redundant degrees of freedom to compensate for temporary constraints on the effectors while producing movement trajectories to targets. For example, people normally use jaw movements during speech, but they can also produce recognizable speech with a pipe clenched in their teeth by increasing lip and tongue movements to compensate for the fixed jaw. The models described here stress *automatic* compensation; i.e.:

- they successfully compensate for constraints on the effectors even if the constraints have never before been experienced,
- they do not require any new learning under the constraining conditions, and
- they do not invoke special control strategies to deal with constraints.

This kind of automatic compensation can greatly reduce the computational requirments of movement planning, potentially freeing up cognitive resources for more important or more difficult tasks. Finally, because these models are self-organizing neural networks whose parameters are tuned during an action-perception cycle, they also require no explicit knowledge about the physical geometry of the effector system being controlled.

In order to highlight the main hypotheses underlying these models, it is useful to consider a simplified view of the movement control process wherein movement trajectories are planned within some reference frame, and these trajectories are mapped into a second reference frame that relates closely to the effector or articulator system that carries out the movements[2]. For example, one can consider speech production as the process of formulating a trajectory within a planning reference frame to pass through a sequence of targets, each corresponding to a different phoneme in the string being produced. The dimensions of this planning frame might correspond to acoustic quantities or locations and degrees of key constrictions in the vocal tract. The planned trajectory can then be mapped into a set of articulator movements that realize the trajectory. The articulator movements are defined within an effector reference frame that relates closely to the musculature or primary movement degrees of freedom of the speech articulators. The process of mapping from the planning frame to the effector frame need not wait until the entire trajectory has been planned, but instead may be carried out in concurrence with trajectory planning.

This paper addresses several important issues concerning the motor equivalent control of redundant effector systems that arise within this view of the movement control process. The nature of the planning reference frame is addressed in Section 2, where it is posited that maximal automatic compensation is possible if trajectory planning is carried out in a reference frame that relates closely to the task space for the movement (e.g., 3D space for reaching or an acoustic-like space for speaking), rather than a frame that relates more closely to the effector or articulator system. The nature of the mapping from the planning frame to the effector frame is addressed in Section 3. Here it is shown that the flexibility made possible by planning movements in a task-based reference frame can be realized by mapping from directions in this frame to directions in the effector frame, rather than from positions in the planning frame to positions in the effector frame. This approach is similar to robotic control techniques that utilize a generalized inverse of the Jacobian matrix (e.g., Hollerbach and Suh, 1985; Klein and Huang, 1983; Liegeois, 1977; Mussa-Ivaldi and Hogan, 1991; Whitney, 1969). Controllers of this kind do not include explicit postural targets for achieving task space targets. Section 4 addresses the issue of whether such models can be reconciled with experimental data indicating

2. The models described in this chapter focus on kinematic problems, ignoring the effects of inertia and external loads on planned movements. The issue of invariant realization of kinematic commands under varying load conditions is a very important one but is beyond the scope of this chapter; see Bullock and Contreras-Vidal (1993) for a proposed solution that is compatible with the models described here.

that humans use a limited range of the possible postures for reaching a given target. It is shown that a certain class of these models that incorporate a form of "postural relaxation" can indeed capture the main aspects of the reaching data. This class of models is compared to a similar proposal by Cruse and colleagues (e.g., Cruse, Brüwer, and Dean, 1993). In Section 5, a neural network model that utilizes postural relaxation is introduced, and simulations are presented to highlight some of the important properties possessed by this type of model. The model suggests two sources that might contribute to curvature in human reaches: (i) a learning bias toward the long axis of the manipulability ellipsoid, and (ii) a tendency to move toward more comfortable postures.

The discussion in this chapter is closely related to the DIRECT model of targeted reaching (Bullock, Grossberg, and Guenther, 1993; Guenther, 1992) and the DIVA model of speech production (Guenther, 1994, 1995a,b). Detailed descriptions of these models, including hypothesized roles of task space feedback, tactile/proprioceptive feedback, and efference copies of outflow commands during both learning and performance, can be found in the cited publications. This chapter will focus on the inverse kinematics transformation performed by these models. The main properties of this transformation are captured by the simplified block diagram shown in Figure 1. Task space targets (e.g., the location of a target in 3-D coordinates for reaching) are compared to the current position of the end effector in task space to produce a desired movement direction Δx. This is then transformed into a desired movement direction in joint space through a learned directional mapping. This neural network mapping can be tuned during a babbling cycle, and, after learning, it approximates a generalized inverse of the Jacobian matrix relating joint space velocities to task space velocities. Externally imposed constraints or perturbations may interfere with commanded joint rotations as indicated in the block diagram. A second neural network mapping transforms the current joint space position back into task space coordinates. Learned mappings of this form have been called "forward models" (Jordan, 1990; Jordan and Rumelhart, 1992). The primary role of the forward model in the DIRECT and DIVA models is to allow planning of movement trajectories without requiring task space feedback. This differs from the role played by the forward model in Jordan (1990) and Jordan and Rumelhart (1992), where it is used to transform task space error into action space error in order to train an inverse model that maps desired task space positions into effector positions.

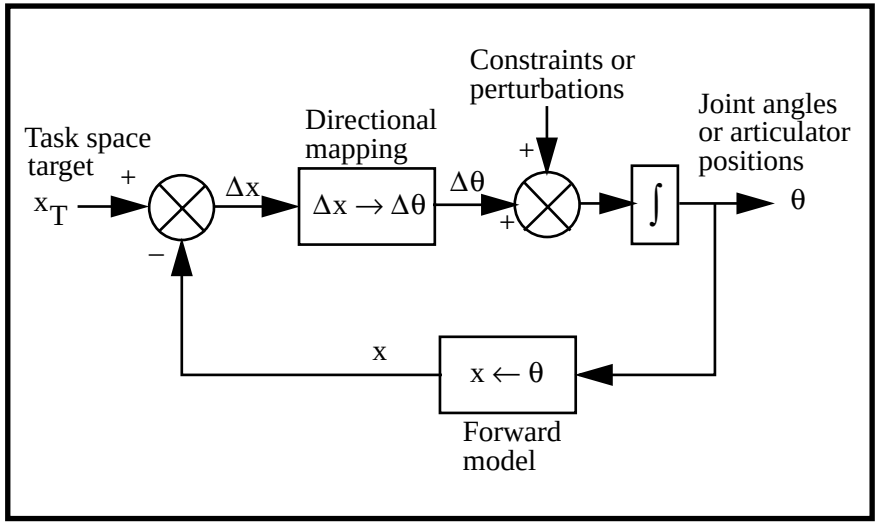

Figure 1: Simplified block diagram of the inverse kinematics tranformation performed by the DIVA and DIRECT models. The directional mapping and forward model are neural network mappings that can be tuned using babbled movements.

2 The planning space

A central issue in biological motor control concerns the nature of the coordinate frame for movement planning. In other words, what variables are explicitly controlled by the central nervous system during reaching movements? The most commonly encountered view in the motor control literature posits that the nervous system directly controls the spatial characteristics of movement, rather than the joint angle characteristics. Morasso's key study of the kinematic properties of planar arm movements (Morasso, 1981) provided some of the strongest experimental support for spatial planning. In this experiment, Morasso recorded nearly straight trajectories of the hand and smooth, bell-shaped spatial velocity profiles.

These characteristics of the hand trajectory appeared to be invariant across different movements in different regions of the workspace. In contrast, the temporal patterns of joint angles did not follow straight lines in joint space, often exhibited double-peaked velocity profiles, and occasionally even exhibited joint reversals.

Although the results of Morasso (1981) are usually taken as strong evidence for spatial planning, several investigators have pointed out that endpoint trajectories are not completely straight but are instead gently curved in many parts of the workspace, particularly in the sagittal plane (e.g., Hollerbach, Moore, and Atkeson, 1986). This might appear to be evidence for joint space trajectory planning, but the curvature seen in these movements is insufficient to support simple joint space interpolation, which would lead to much larger curvature. To account for this, Hollerbach, Atkeson, and Moore (1986) proposed a modification to joint space interpolation, which they termed "staggered joint interpolation", in which different joints begin moving at different times in order to produce straighter trajectories. However, this model cannot account for the joint reversals seen in the study of Morasso (1981). Uno, Kawato, and Suzuki (1989) proposed a model based on the minimization of the torque change along the trajectory. According to this minimum torque model, curvatures is an inherent side effect of a control strategy which controls joints instead of hand trajectories. Although this model provides one possible explanation for much of the curvature seen in reaches, a study by Wolpert, Ghahramani, and Jordan (1995) showed that increasing perceived curvature of movements through altered visual feedback caused subjects to change their movements to produce visually straighter hand paths, at odds with non-spatial planning models such as the minimum torque change model.

Another group of researchers has tried to explain the curvature of human reaches under the assumption of spatial trajectory planning. For example, Wolpert, Ghahramani, and Jordan (1994) demonstrated that some of the curvature of hand trajectories can be attributed to perceptual distortion. I.e., in some parts of the workspace, a curved reach appears straighter than it actually is. However, perceptual distortion alone did not appear to be sufficient to fully account for the curvature of reaches in this study. Flash (1989) suggested that curvature arises as the result of interactions between the viscoelastic properties of muscles and the inertial properties of the arm while following a straight-line "equilibrium trajectory". This chapter suggests two factors related to the inverse kinematics transformation that might also contribute to curvature in reaches planned as straight lines in task space: a learning bias toward movements along the long axis of the manipulability ellipsoid (described in Section 3.2), and a tendency toward joint rotations that lead to more comfortable postures (described in Sections 4 and 5).

A different rationale for spatial trajectory planning arises from the viewpoint that maximally flexible performance can be achieved if movements are planned in a reference frame that relates as closely as possible to the task space for the movement (e.g., 3-D space for reaching or acoustic space for speaking), rather than a frame that relates closely to the effectors or articulators (e.g., Bullock, Grossberg, and Guenther, 1993; Guenther, 1992, 1994, 1995a,b; see also Saltzman and Kelso, 1987; Saltzman and Munhall, 1989). This rationale largely motivates the models discussed in this chapter. For example, it is well-known that speakers are typically capable of reaching acoustic targets for vowels in the presence of constraints that prevent certain movements of the lips or jaw during speaking (e.g., Abbs and Gracco, 1984; Lindblom, Lubker, and Gay, 1979; Savariaux, Perrier, and Orliaguet, 1995). This kind of compensation requires completely different final positions for the unconstrained articulators, and, therefore, completely different articulator or effector space trajectories. Therefore, systems that explicitly plan trajectories in a coordinate frame relating closely to the effectors or articulators, such as a muscle length or joint angle coordinate frame, must take the constraint into consideration during the planning process. In contrast, if the motor control system plans movements to speech targets as acoustic trajectories (Bailly, Laboissière, and Schwartz, 1991; Guenther, 1995b; Perkell et al., 1993) and maps these planned trajectories into articulator/effector space movements in a manner that provides automatic compensation for externally imposed constraints (as described in the next section), then the complexity of the movement planning process is greatly reduced since the constraints can be largely ignored[3].

3 Directional mappings

Trajectories planned in task space must still be carried out by articulator or effector movements. One possibility is to use a position-to-position mapping from task space to effector space; e.g., each point in 3-D space can be mapped to a joint configuration that is satisfactory for this point. Another possibility is to use a directional mapping from desired movement directions in task space into movement directions in effector space (e.g., joint rotations). The DIRECT and DIVA models use the latter form of mapping because it provides the automatic compensation for externally imposed constraints on effector motion motivated in the previous section and described below.

The use of a directional mapping for movement control is closely related to robotic controllers that utilize a generalized inverse of the Jacobian matrix[4] (e.g., Baillieul, Hollerbach, and Brockett, 1984; Hollerbach and Suh, 1985; Klein and Huang, 1983; Liégeois, 1977; Mussa-Ivaldi and Hogan, 1991; Whitney, 1969). The relationship between spatial velocity of the end effector and the joint velocities of a manipulator such as an arm is given by the following equation:

$$\dot{x} = J(\theta)\dot{\theta} \tag{1}$$

where $\dot{x}$ is the spatial velocity vector of the hand, $\dot{\theta}$ is the joint velocity vector, and $J(\theta)$ is the manipulator's Jacobian matrix, whose elements depend only on the joint configuration θ. To obtain a joint rotation vector that moves the hand at a desired spatial velocity, we can rearrange this equation:

$$\dot{\theta} = J^{-1}(\theta)\dot{x} \tag{2}$$

where $J^{-1}(\theta)$ is an inverse of the Jacobian matrix. For a redundant manipulator, a unique inverse for J does not exist. In this case, J^{-1} is a generalized inverse, or pseudoinverse, of the Jacobian matrix. The most commonly used generalized inverse is the Moore-Penrose (MP) pseudoinverse, which has the desirable property of returning the minimum norm joint rotation vector that can produce the desired spatial velocity.

Directional mappings as discussed in this chapter are generally related to, but often slightly different from, a pseudoinverse of the Jacobian matrix. In particular, learned approximations to a pseudoinverse that do not strictly satisfy Equation 2 will be discussed. With such a directional mapping, the spatial trajectories produced by the inverse kinematics transformation schematized in Figure 1 are not straight lines in task space, but instead are gently curved.

3.1 Motor equivalence

The ability to reach targets in pseudoinverse-style controllers such as the DIVA and DIRECT models is very robust to error in the directional mapping. This can been seen in the following example. Imagine an intended straight-line movement of the hand to a target in 3D space, as schematized in Figure 2. Assume that a 30° error in the directional mapping causes the actual trajectory to veer upward from the desired straight-line trajectory. The desired task space movement direction (indicated by dashed arrows in the figure) always points from the current position of the hand to the target. As the actual trajectory moves further away from the desired trajectory, the task space direction vector points more and more downward to counteract this error in movement direction. The system thus "steers in" toward the target. As long as the directional mapping is off by less than 90° and some form of feedback regarding end effector position is available, the target will successfully be reached, although for large directional errors the trajectory will deviate significantly from a straight line in planning space.

This property has several important implications for biological movement control. First, it suggests how a person can easily overcome constraints on the effectors (such as a cast limiting arm movement during reaching or a bite block limiting jaw movement during speaking) that effectively introduce error in the directional mapping, and thus provides an

3. This is not to say that path planning always takes place in a purely spatial coordinate frame without regard to arm geometry and constraints. For example, studies by Dean and Brüwer (1994) and Sabes, Wolpert, and Jordan (in preparation) show that the geometry of the arm is taken into consideration when planning trajectories around obstacles.

4. An interesting treatment of a class of models that utilize the transpose of the Jacobian matrix, rather than a generalized inverse, is provided by Mussa Ivaldi, Morasso, and Zaccaria (1988).

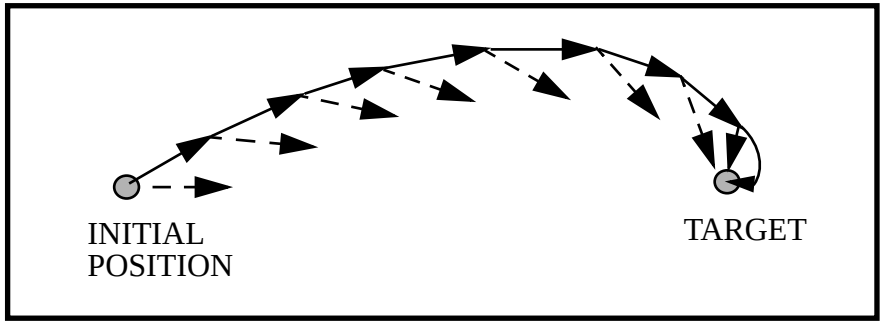

Figure 2: Robustness to error in the directional mapping for targeted movements. Here a 30° error in the mapping causes the actual trajectory to veer from the desired straight-line trajectory. The desired task space movement direction at each point along the trajectory is indicated by the dashed arrows. As the actual trajectory moves further away from the desired trajectory, the task space direction vector points more and more downward to counteract this error in movement direction, allowing the system to "steer in" toward the target. As long as the directional mapping is off by less than 90°, the target will be successfully reached.

explanation for one form of motor equivalence. Simulations verifying the abilities of the DIRECT and DIVA models to overcome errors in the directional mapping due to shifting of the visual field, joint blockage, and blockage of one or more speech articulators are provided elsewhere (Bullock, Grossberg, and Guenther, 1993; Guenther, 1992, 1994, 1995a,b). Second, it implies that even a coarsely learned directional mapping, such as that possessed by an infant in the early months of life, can be used to reach objects or produce speech sounds, although with imperfect movement trajectories. Finally, it shows how error correction capabilities can automatically arise from the same mechanism used to control normal movements, unlike a controller that aims for postural targets and must somehow choose a new postural target if the normal target is inaccurate or unreachable due to external constraints.

3.2 Learning issues

Several interesting issues arise when one considers how a neural network can learn a directional mapping between task space coordinates and effector coordinates. One such issue concerns what appears to be a rather broad class of models in which the amount of learning that occurs during a given movement is scaled by the size of the spatial movement of the end effector. Models that possess this property include versions of DIRECT that use gradient descent in a linear neural network (Fiala, 1995) or in a radial basis function network (Cameron, 1995, and the version described later in this chapter).

To understand the learning properties of this class of models, it is useful to consider the *manipulability ellipsoid* (Yoshikawa, 1985) that relates joint rotations to spatial velocities. Consider the set of joint velocity vectors of unit length when at a given joint configuration for a three-joint arm constrained to planar movements. These vectors lie on the unit sphere in joint velocity space. If one plots the spatial velocity vectors that would result from these joint velocity vectors (one example is shown Figure 3), the result is a set of vectors that fall inside an ellipse in the plane of movement. Correspondingly, moving in a spatial direction aligned with the long axis of this manipulability ellipsoid requires much less joint rotation than moving in a direction along the short axis. The eccentricity and direction of the ellipsoid vary as functions of joint configuration.

Simulations using gradient descent learning in the networks of Fiala (1995) and Cameron (1995) show that residual error in the directional mapping after training (e.g., due to having too few cells or too little training to fully learn the mapping) tends to warp the actual movement direction away from the desired direction toward the long axis of the manipulability ellipsoid. Inspection of the networks after training suggests that this warping can be roughly characterized by the following equation:

$$\phi_{actual} = \phi_{desired} - \rho \frac{l_{long} - l_{short}}{l_{long}} \sin[2(\phi_{desired} - \phi_{ellipsoid})] \qquad (3)$$

where ρ is a scalar (in degrees or radians) that depends on the amount of training and the number of cells in the network, l_{long} and l_{short} are the lengths of the long and short axes of the manipulability ellipsoid, $\phi_{desired}$

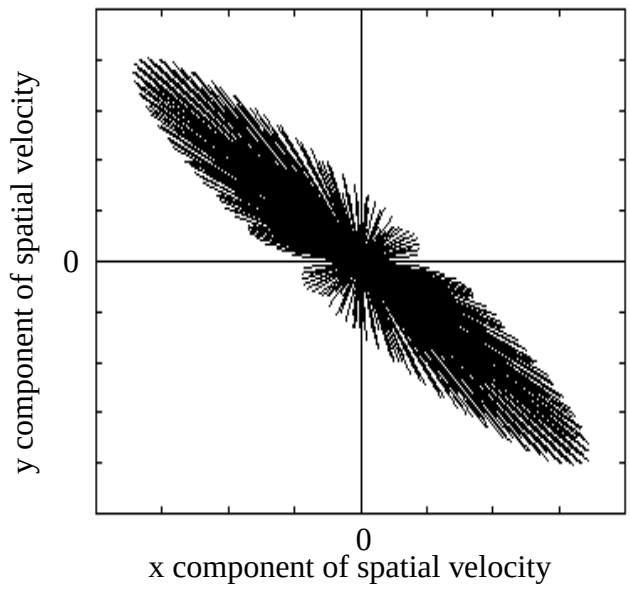

Figure 3: The spatial velocity vectors produced by unit vectors in joint velocity space take the form of a *manipulability ellipsoid* (Yoshikawa, 1985). Spatial movements along the long axis of the ellipsoid require less joint rotation than movements along the short axis. The shape and orientation of the ellipsoid vary with joint configuration.

is the desired spatial movement direction, ϕ_{actual} is the spatial movement direction after warping by the directional mapping, and $\phi_{ellipsoid}$ is the spatial direction of the long axis of the manipulability ellipsoid. A polar plot of the residual error as a function of movement direction for a radial basis function network is plotted in Figure 4 along with the manipulability ellipsoid at the corresponding joint configuration. As suggested by Equation 3, the error is roughly zero for desired movements along the long and short axes of the ellipsoid, and it reaches maxima for desired movement directions falling halfway between the long and short axes. Although the network's error function at different joint configurations usually took a form similar to that shown in Figure 4, there were significant deviations from this form in some regions of workspace, emphasizing that Equation 3 is only an approximation used here to provide some insight into the general form of the residual error rather than a precise characterization.

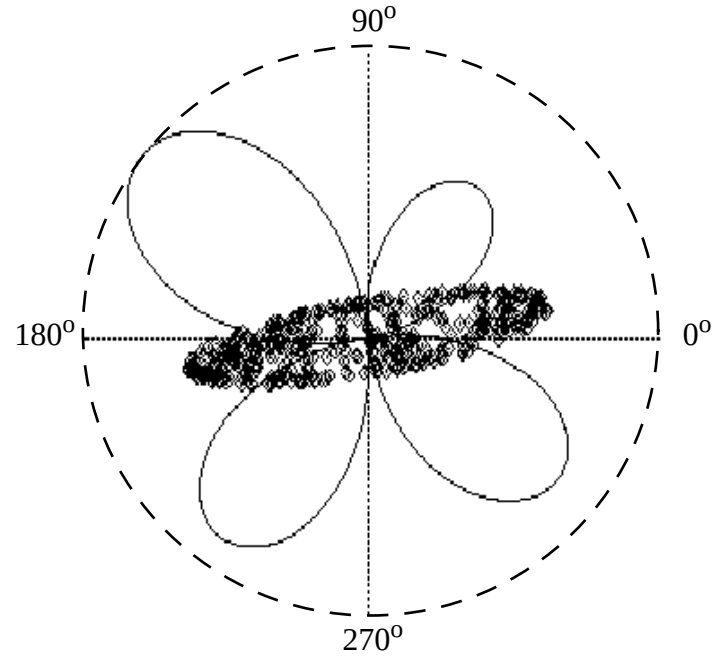

Figure 4: Polar plot of residual error at one joint configuration in a neural network after learning a directional mapping. The elliptical cloud of points forms the manipulability ellipsoid at this joint configuration and is included for comparison purposes. The petal-shaped curves indicate the residual error as a function of movement direction; the distance from the center of the plot corresponds to the magnitude of the directionl error (the dotted circle corresponds to 4° directional error). Error is approximately zero along the long and short axes of the ellipsoid, and it reaches maxima at points roughly halfway between the long and short axes. This residual error leads to movements that are warped toward the long axis of the ellipsoid.

An intuitive feel for why this error pattern arises can be gained by considering that, for the same amount of joint rotation, movements along the

long axis of the manipulability ellipsoid produce more spatial displacement of the end effector than movements in other directions. More learning on average will thus occur for movements along the long axis of the ellipsoid in systems where the amount of learning is larger during larger spatial displacements of the end effector. If the error does not converge to zero due to too few cells or too little training in a neural network learning the directional mapping, then one would expect the learning to be biased toward the joint rotation patterns that caused the most learning; i.e., those producing movements along the long axis of the ellipsoid. This intuition suggests that a similar form of warping in the directional mapping will arise in a variety of neural network models in which the amount of learning is scaled by the size of the spatial movement. Given that the learning of directional mappings has been hypothesized to occur in brain regions such as the cerebellum (e.g., Pellionisz, and Llinás, 1985), this in turn suggests that this kind of warping may contribute to the curvature seen in human reaches. Figure 5 compares the movement paths produced by Subject 1 in the study of Morasso (1981) with the paths produced by the simple computational model of Equation 3 when the manipulability ellipsoid was calculated using the arm segment lengths of the same subject. The direction and amount of curvature in the model's paths are largely, but not entirely, consistent with the experimental results. It is expected, however, that many different factors contribute to the curvature seen in human reaches, potentially including perceptual distortion (Wolpert, Ghahramani, and Jordan, 1994) and a tendency toward more comfortable postures as described in Section 5.

One desirable consequence of this learning bias concerns singularities in the workspace, where geometric limitations of the manipulator make some spatial movement directions impossible. Producing a desired spatial velocity when moving toward a singularity generally requires higher and higher joint velocities, a well-known problem of pseudoinverse techniques (e.g., Baillieul, Hollerbach, and Brockett, 1984). When approaching a singularity, the manipulability ellipsoid "flattens" as some movement directions become impossible; for example, a 2D ellipsoid as shown in Figure 3 collapses into a line segment as the short axis shrinks toward the origin. In systems where learning is biased toward the long axis of the ellipsoid, the problem of extremely high joint velocities is somewhat alleviated because a desired spatial movement direction that does not align with the long axis of the ellipsoid is warped toward the long axis by the directional mapping, with the amount of warping roughly scaling with the ratio of the long axis to the short axis. The system thus learns movements that are largely aligned with the long axis of the ellipsoid and therefore require relatively small joint velocities. Fiala (1995) provides a demonstration of the well-behaved performance of this kind of network at workspace singularities.

A second potential benefit of biasing movements toward the long axis of the manipulability ellipsoid is a reduction in the amount of total joint rotation required to reach a target. A simulation was run to compare the amount of joint rotation required to follow the gently curved hand paths arising from Equation 3 to straight paths. In both cases, the MP pseudoinverse was used to transform the movement path into joint rotations, and the square root of the sum of the squared joint increments was calculated at each time step for movements to 80 randomly chosen targets. The gently curved paths produced by Equation 3 required 13.9% less total joint rotation than the straight paths. The amount of rotation saved by the neural networks described above, however, varied significantly with the exact network architecture and amount of training.

The final two learning issues to be addressed here are related to potential shortcomings of "direct inverse" learning techniques pointed out by Jordan and Rumelhart (1992). A direct inverse learning approach is one in which movement commands are generated in effector space (typically randomly during training), and the system learns a mapping from the task space consequences of these movements to the movement commands that caused them. This inverse mapping can later be used to command effector space movements to achieve task space goals. The DIRECT and DIVA models currently utilize a direct inverse learning scheme.

One shortcoming pointed out by Jordan and Rumelhart (1992) is that learning in direct inverse models is not "goal-directed"; i.e., it is not sensitive to errors in sensation space, and there is therefore no direct way to find an action that corresponds to a particular desired sensation. The validity of this claim, however, depends on the generalization properties of the direct inverse learning system. For example, in the direct inverse model described later in this chapter, learning generalizes to all spatial directions at each

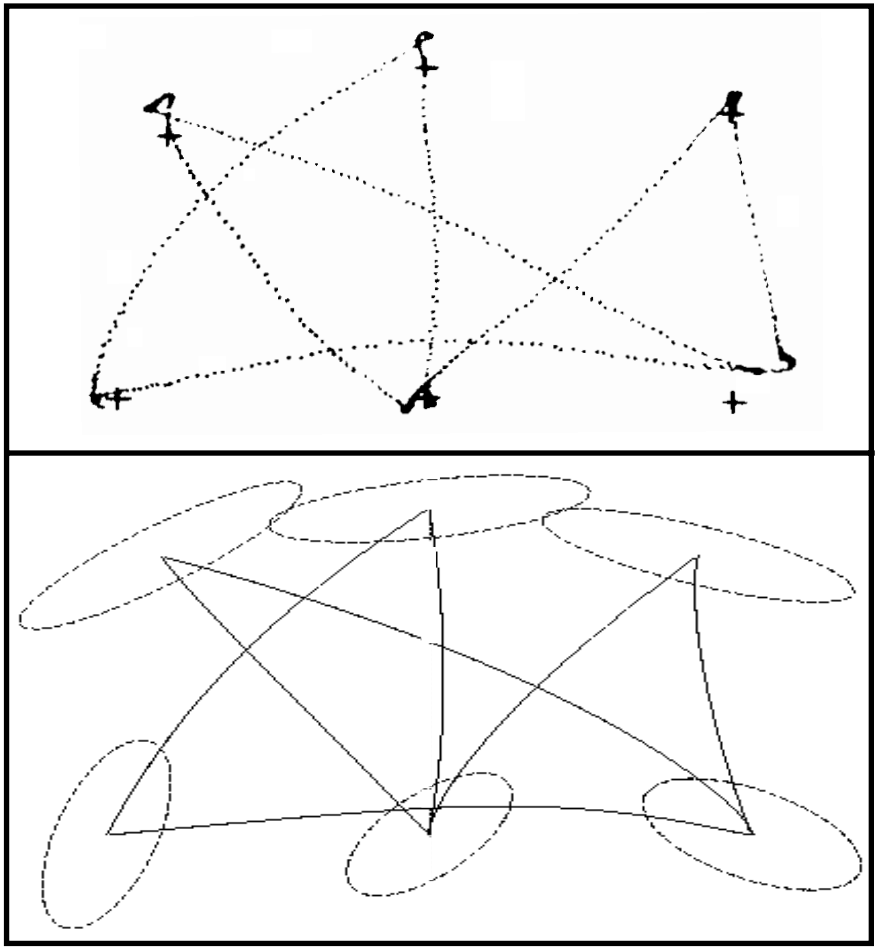

Figure 5: (Top) Movement paths reported for Subject 1 in the study of Morasso (1981) [adapted from Morasso, 1981]. (Bottom) Movement paths produced by the simple computational model of Equation 3 using the arm segment lengths for Subject 1 to calculate the manipulability ellipsoid (shown for each target location).

sampled joint configuration; this is because the model learns a directional mapping that is an approximation to the Jacobian pseudoinverse at each joint configuration, and the approximate Jacobian pseudoinverse learned for one movement direction can be used for all other movement directions. Consider the task of reaching from one point to a second point, where the task is to be learned by repeatedly attempting the reach. On the first attempt, the model will most likely move in the wrong direction, but this movement will drive learning in the directional mapping that applies to *all* movement directions. Subsequent reaches will become more and more accurate as the approximate Jacobian pseudoinverse improves. In other words, the model will learn the task simply by repeatedly attempting to perform it; this is goal-directed learning. Although it is convenient to utilize random movements during training in order to insure coverage of the workspace, direct inverse models such as the one described in Section 5 of this chapter can thus also use a goal-directed learning process.

The second potential shortcoming concerns convexity and learning in redundant systems. Jordan and Rumelhart (1992) point out that most direct inverse learning techniques learn an average of possible effector space solutions for a given task space goal. Such a system can thus learn an invalid solution if the solution space is non-convex, since an average of solutions is only guaranteed to be a solution for convex solution spaces. The solution space of joint configurations that achieve a desired spatial position is non-convex, and the convexity problem can indeed prevent the successful reaching of targets in direct inverse models that learn positional mappings. However, this is not a serious problem for many direct inverse models that learn directional mappings, such as the DIVA and DIRECT models. This is verified by simulation results of the DIRECT model successfully performing reaches using a redundant arm (Bullock, Grossberg, and Guenther, 1993; Guenther, 1992) and the DIVA model successfully reaching acoustic targets using a highly redundant articulator set (Guenther, 1994, 1995a,b). Consideration of two properties of directional mappings help clarify why this is the case. First, directional mappings are locally linear, even for redundant systems. This means that if one only considers a small region of joint space, the set of joint velocity vectors that produce a desired spatial velocity is convex. Therefore, systems that effectively learn different directional mappings in different regions of joint space can largely, if not completely, avoid this problem. For example, the radial basis network described in Section 5 utilizes different parameters in different regions of workspace (corresponding to different radial basis

functions) and smoothly interpolates between these parameter sets. Second, systems that use directional mappings can successfully reach targets even if the directional mapping contains a large amount of error (discussed in Section 3.1). Therefore, any residual error that might exist, e.g. from assuming linearity over too large a region of joint space, will not prevent the system from reaching targets, but will instead only lead to curvature in the movement trajectories. If this residual error is related to the manipulability ellipsoid as described above, the resulting slight curvature might have useful side effects such as better performance near singularities and reduced total joint rotation. It should be noted, however, that even though direct inverse models are apparently sufficient for learning directional mappings in redundant systems, the forward modeling learning scheme described in Jordan and Rumelhart (1992) could also be used to learn directional mappings and might still offer advantages over direct inverse techniques.

A final note concerns the use of the word "convex" to describe different aspects of movement control models. Guenther (1995a) describes how a *convex region theory* for the targets of speech provides a simple, unifying explanation for many long-studied speech production phenomena, including contextual variability, carryover coarticulation, anticipatory coarticulation, and a collection of speaking rate effects. It is important to note, however, that the speech sound targets learned by the model are convex in *task* space, not effector space. These convex region targets are generalizations of the traditional point targets assumed by the vast majority of models of reaching and speaking, and the model is *not* limited to convex solution regions in effector space. For example, Guenther (1995b) describes how the model successfully handles a case where the convex task space target region maps into disconnected regions in effector space.

3.3 Motor cortical cells

Another point of interest regarding directional mappings is their compatibility with a number of neurophysiological studies investigating functional properties of single cells in motor cortex. The most salient aspect of these studies is that a typical motor cortical cell involved in reaching has a preferred direction of movement which will cause the cell to fire maximally, and the cell's activity will fall off roughly as the cosine of the difference between the movement direction and the cell's preferred direction (e.g., Georgopoulos et al., 1982; Kalaska et al., 1989). The mathematical analysis of Mussa-Ivaldi (1988) showed that cells relating to movement direction in muscle length space (i.e., cells coding muscle shortening velocities) will necessarily have cosine-shaped tuning curves when analyzed with respect to the spatial direction of movement. This property is evident in recent neural network models (e.g., Burnod et al, 1992; Bullock, Grossberg, and Guenther, 1993) and is illustrated for cells in the DIRECT model in Figure 6. Cosine-shaped tuning curves are not built into the model cells but instead arise as a result of learning the mapping between spatial directions and joint rotations. Additional similarities between the effector direction vector cells and motor cortical cells include the observations that a cell's preferred direction will typically change for different starting positions of the hand in 3-D space (Caminiti et al., 1990) and for different joint configurations when the hand is in the same 3-D spatial position (Scott and Kalaska, 1995). (The latter property is not present in the model of Burnod et al., 1992, because the tuning curves of cells in this model depend on spatial position of the hand rather than joint configuration.)

It is important to note, however, that the gross aspects of motor cortical cell firing preferences, such as directional tuning curves, appear to be compatible with a large range of movement control hypotheses (Mussa Ivaldi, 1988; Sanger, 1994). A much more detailed breakdown of cell properties, including aspects such as force dependency, tonic activity as a function of position, and temporal pattern of activity during a reach (e.g., Kalaska, Cohen, Hyde, and Prud'homme, 1989), must be taken into consideration in order to gain a more insightful picture of the role of motor cortex in movement control. See Bullock, Cisek, and Grossberg (1995) for a modeling study that proposes distinct roles for many of the subclasses of cortical cells in a neural network model related to DIRECT.

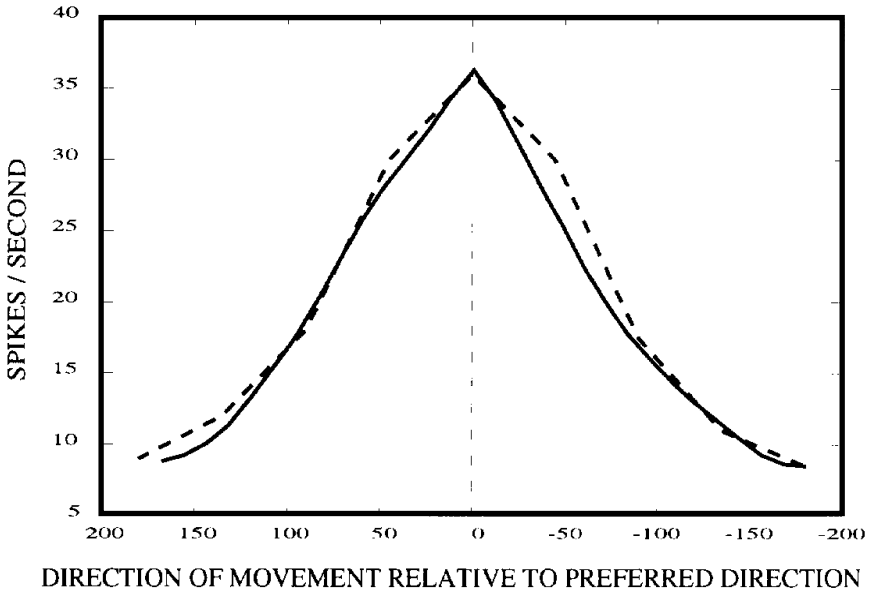

Figure 6: Comparison of average directional tuning curves obtained from single cell studies in motor cortex by Kalaska et al. (1989) (broken line) to tuning curves of cells in the motor direction vector of the DIRECT model (solid line) [adapted from Bullock, Grossberg, and Guenther (1993)]. Model firing rates are in arbitrary units and have been scaled along the y axis to cover the same range as the Kalaska et al. data. After learning the directional mapping, the model's cells display many of the properties seen in motor cortex cells.

4 Is postural target information necessary or desirable?

Another important issue in biological motor control concerns whether the motor control system utilizes postural targets in planning movements. For example, Rosenbaum and colleagues have posited that the motor control system performs a reach to a target by first choosing an appropriate posture for that target using a set of postures stored in memory (Rosenbaum, Engelbrecht, Bushe, and Loukopoulos, 1993). The Rosenbaum et al. model then forms a movement trajectory using linear interpolation in joint space from the initial configuration to the chosen posture. This model is contradicted by psychophysical studies showing smoothly interpolated trajectories in task space but not in joint space (e.g., Morasso, 1981), but the more general question of whether the motor system uses postural targets in some manner remains open.

Several psychophysical studies have investigated this issue, with mixed results. Studies of pointing movements with the elbow fully extended indicate that the final posture of the arm is relatively invariant for a given target position (Hore, Watts, and Vilis, 1992; Miller, Theeuwen, and Gielen, 1992). For pointing movements on a planar surface, Cruse, Brüwer, and Dean (1993) reported that the final postures "were virtually independent of the configuration at the start of the pointing movement" (p. 131), and for reaches to grasp an oriented object, Desmurget et al (1995) similarly report that "the final limb angles were highly predictable" (p. 905).

However, all of these studies imposed constraints on the arm that removed some of the redundancy available during free reaches. Subjects in the Desmurget et al. were required to orient the hand so as to grasp an object, and all reaches started from the same initial arm configuration. The Hore et al. (1992) and Miller et al. (1992) studies effectively removed the elbow degree of freedom by requiring full extension of the arm, and the experiments reported in Cruse et al. (1993) were constrained to a plane. In a less constrained three-dimensional reaching task, Soechting, Buneo, Herrmann, and Flanders (1995) reported that the final postures of reaches to a given target depended on the starting configuration of the arm. Furthermore, Cruse (1986) reported that although the final postures for planar reaches did not vary nearly as much as was possible given the geometry of the arm, there was a small but significant effect of the initial posture on the final posture, and Cruse and Brüwer (1987) reported that choosing starting postures with extreme joint angles increased the variability of final postures.

It therefore appears that although the motor system uses a far smaller range of final postures than is possible given the redundancy of the arm, some variability in final posture is seen, particularly when starting from very different initial postures. Furthermore, reaches can be successfully

completed even when constraints on the joints cause unusual final postures. For example, Cruse, Brüwer, and Dean (1993) reported that applying a force that opposes elbow flexion causes subjects to reach different final postures. The ability to perform kinematically normal targeted reaches using pointers that change the effective length of the forearm (e.g., Lacquaniti, Soechting, and Terzuolo, 1982) also speaks against the use of fixed postural targets since a completely different final posture is required to reach the same spatial position for each different forearm length.

In contrast to models like that of Rosenbaum et al., pseudoinverse-style controllers do not generally associate a particular posture with each target in planning coordinates. In fact, Klein and Huang (1983) discuss how this leads to a well-known potential drawback of pseudoinverse techniques. In a typical pseudoinverse-style controller, repeated tracing of a closed shape such as a square by the end effector can lead the effector system to "curl up" into unusual configurations, potentially reaching physical limits of the effector system. Klein and Huang (1983) show that the MP pseudoinverse has this problem, which is related to the non-integrability of a differential equation related to the pseudoinverse.

Mussa-Ivaldi and Hogan (1991) describe a different generalized inverse of the Jacobian matrix which is integrable and therefore avoids this problem. If no obstacles or joint constraints act on the effector system, this approach has the effect of assigning a joint configuration to each spatial position of the end effector (given the initial configuration of the system), even though no target configuration needs to be explicitly represented in the planning process. This raises an important point concerning the need for postural target information: a lack of postural targets in the movement planning process does not necessarily imply that the motor control system will not produce final postures that appear stereotypical for a given target. For example, if no joint constraints or obstacles are present, a controller using the generalized inverse of Mussa-Ivaldi and Hogan (1991) will always reach the same joint configuration for a given spatial target position.

Unfortunately, the same property that allows the Mussa-Ivaldi and Hogan pseudoinverse to overcome the problem of curling up into unusual configurations during repeated tracing movements leads to a different problem in a world where joint constraints or obstacles often make unusual joint configurations necessary. Whereas under unconstrained conditions the Mussa-Ivaldi and Hogan pseudoinverse will lead to the same final posture for a given target regardless of starting posture, blocking a joint will cause the system to use a different final posture, often requiring an "uncomfortable" angle for some other joint or joints. This is schematized in Figure 7. The ability to use the uncomfortable but necessary final posture shown on the right side of Figure 7 is indeed desirable, and it highlights an important advantage of this type of controller over controllers that explicitly specify a final posture as the target of a movement: although a controller using the Mussa-Ivaldi and Hogan pseudoinverse will normally produce the same posture for a given target position regardless of starting position, it still possesses the property of automatic compensation for constraints that prevent the arm from reaching the normal configuration (see Section 3). The problem, however, arises when the constraint is removed. Because they do not correct for uncomfortable postures, systems using the MP pseudoinverse or the pseudoinverse proposed by Mussa-Ivaldi and Hogan will maintain the uncomfortable wrist angle for all future reaches to this target and to much of the rest of the workspace. For the Mussa-Ivaldi and Hogan (1991) pseudoinverse, the implicit mapping from target positions to joint configurations has been changed for all points in the workspace.

This highlights a dilemma faced by biological motor systems: on the one hand, they should not rely on explicit postural targets as this removes the ability to automatically compensate for constraints on the effector system, but on the other hand they need to avoid the problem of extreme or uncomfortable joint configurations that can occur with most pseudoinverse-style control techniques.

A solution to this dilemma lies in a different pseudoinverse-style approach. Given a redundant arm in a particular posture, for each desired movement direction there exists a set of joint rotation vectors that will all move the arm in the desired direction. If a controller utilizes a joint rotation vector from this set that also moves the arm in the direction of a more comfortable posture (e.g., toward the center of the joint ranges), the system will have a tendency to end up in comfortable joint configurations. Since spatial movement directions are mapped into joint rotations, the system will maintain the motor equivalence capabilities of pseudoinverse-style controllers.

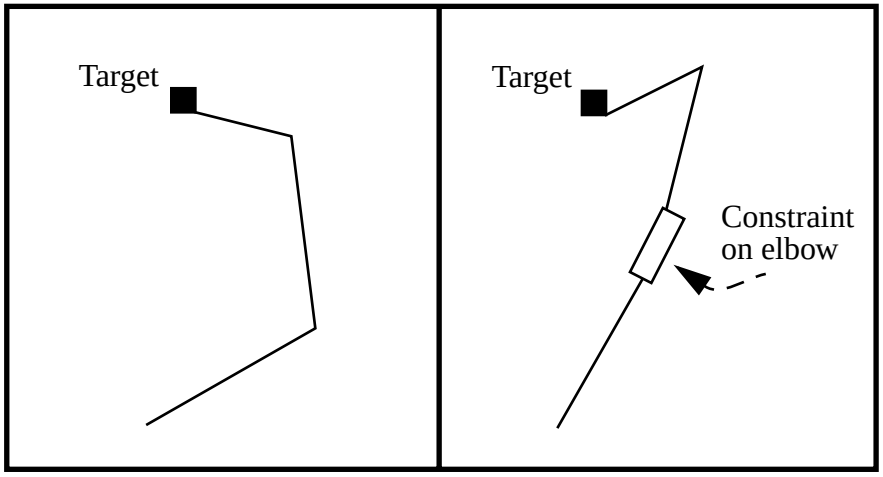

Figure 7: Potential problem for pseudoinverse-style controllers. The left side schematizes the normal configuration used by a three-joint arm to reach a target. The right side shows what happens if a constraint is placed on the elbow. Here, the pseudoinverse controller will successfully reach the target, using an "uncomfortable" wrist configuration forced by the constraint. The problem arises when the constraint is removed: because they do not correct for uncomfortable postures, systems using only the Moore-Penrose (MP) pseudoinverse or the pseudoinverse proposed by Mussa-Ivaldi and Hogan (1991) will maintain the uncomfortable wrist angle for all future reaches to this target and to much of the rest of the workspace.

Furthermore, such a system avoids the problem of indefinitely maintaining an uncomfortable posture that may have been forced by temporary joint constraints or obstacles since it will try to move toward a more comfortable posture with every movement. Such a solution will not generally utilize an integrable pseudoinverse and can therefore end up in different postures depending on starting position, but this property is a necessary one to avoid the problem of maintaining uncomfortable configurations forced by temporary joint constraints or obstacles. The potential ill effects of non-integrability pointed out by Klein and Huang (1983) and Mussa-Ivaldi and Hogan (1991) are also avoided due to the tendency toward comfortable postures: when repeatedly tracing a closed path, the postures will converge toward comfortable postures for each point along the path. Finally, the tendency toward comfortable postures will also greatly limit the range of postures the controller will use, in keeping with the experimental results described above.

Pseudoinverse-style control schemes with the desired property, which we will refer to as *postural relaxation* in this chapter, have been proposed in the robotics literature (e.g., Klein and Huang, 1983; Liégeois, 1977). A typical approach is to calculate joint rotations according to the following equation:

$$\dot{\theta} = J^{-1}(\theta)\dot{x} + \dot{\theta}_0 \tag{4}$$

where $\dot{\theta}_0$ is a joint rotation vector that causes no net displacement of the end effector. A typical choice for $J^{-1}(\theta)$ in this equation is the MP pseudoinverse. If we choose $\dot{\theta}_0$ so that it tends to move the arm toward a more comfortable configuration, the system will have the desired properties. This can be done by choosing

$$\dot{\theta}_0 = \alpha(I - J^{-1}J)\nabla_\theta H \tag{5}$$

where α is a constant scaling parameter, $\nabla_\theta H$ is the gradient of some function H that we seek to minimize, and I is the identity matrix. The matrix $(I - J^{-1}J)$ has the effect of projecting the vector $\nabla_\theta H$ onto the nullspace of J, thus assuring that $\dot{\theta}_0$ causes no net displacement of the end effector. If we choose H so that it decreases for increasing "comfort", then the controller will have the desired property of choosing movement directions that tend toward more comfortable joint configurations. One obvious choice is

$$H = \sum_{i=1}^{N} \beta_i \left(\frac{\theta_i - \theta_i^c}{\theta_i^r} \right)^2 \tag{6}$$

where θ_i^c is the center of the ith joint's range (or some other position of maximal comfort), θ_i^r is the extent of the ith joint's range, and β_i is a scaling parameter. This function can also be thought of as describing the amount of energy needed to maintain a posture θ. It has the effect of

penalizing postures for which one or more joints are near their mechanical limits, therefore promoting a more "natural" configuration of the arm. Alternatively, cost functions can be introduced to quantify other aspects of postural fitness. For example, one can consider the manipulability (Yoshikawa, 1985) associated with a given joint configuration. A desirable posture, in this case, might be one where the manipulability ellipsoid is nearly circular, indicating that the amount of joint rotation required to move a fixed distance is roughly equal for all movement directions. The stiffness ellipsoid (e.g., Mussa-Ivaldi, Hogan, and Bizzi, 1985; Flash and Mussa-Ivaldi, 1990) might also be used in this manner.

A related approach to movement planning has been discussed in the motor control literature by Cruse and colleagues (e.g., Cruse, 1986; Cruse and Brüwer, 1987; Cruse, Brüwer, and Dean, 1993). Cruse (1986) initially proposed that the final postures of reaching movements might be chosen so as to minimize a cost function similar to the cost function of Equation 6. In this form, Cruse's model is closely related to the model of Rosenbaum and colleagues described above. The model was modified in Cruse and Brüwer (1987) and Cruse, Brüwer, and Dean (1993). In the later formulations, this "minimum cost principle" was first combined with joint rotations calculated using the MP pseudoinverse. This component of the model produces straight-line trajectories; to explain curved movement paths, the model also included a competing "end-point control" process which by itself would produce a trajectory by linearly interpolating between initial configuration and a final posture of the arm in joint space. The net effect of combining the purely local pseudoinverse and minimum cost strategies with the global end-point control process is a moderately curved trajectory whose curvature represents a compromise between a straight-line trajectory and a highly curved trajectory formed through joint space interpolation.

The next section describes an implementation of postural relaxation in the DIRECT model that is closely related to both the robotics approach described in Equations 4-6 and the Cruse et al. (1993) model. The current model differs from these approaches in that model parameters are learned during an action-perception cycle instead of relying on explicit knowledge of the manipulator geometry in the form of the Jacobian matrix. This is an important property not only for biological controllers that presumably do not have explicit knowledge of the Jacobian, but also for artificial controllers when the mathematical relationship between task space and effector space is not known in the form of a Jacobian matrix (e.g., the relationship between formant frequencies and articulator positions in the speech production models of Bailly, Laboissière, and Schwartz, 1991, and Guenther, 1995b). The model also differs from the approach described in Equations 4-6 in that it tends to produce slightly curved trajectories. Finally, the model differs from the Cruse et al. model in that curvature arises from purely local computations due to biases in the learning process rather than through competition between a local straight-line planning process and a global planning process that attempts to move the arm to a postural target.

5 An adaptive pseudoinverse-style controller with postural relaxation

A version of the DIRECT model utilizing the network described in the following paragraphs for the directional mapping in Figure 1 was trained to move throughout the workspace, then tested to demonstrate some of the desirable properties of the postural relaxation control scheme. A three-joint planar arm with a shoulder angle range of -160° to 90° (with respect to the horizontal axis), an elbow angle range of 0° to 180°, and a wrist angle range of -160° to 160° was used for the simulations; elbow and wrist angles are measured with respect to the alignment with the previous link.

The mapping from the desired movement direction vector $\dot{x}$ to the joint rotation vector $\dot{\theta}$ is formed in the model according to the following equation:

$$\dot{\theta} = A(\theta)\dot{x} + B(\theta). \tag{7}$$

The elements of the matrix A and the vector B are the outputs of hyperplane radial basis function (RBF) networks with adaptive centers (see Cameron, 1995; Poggio and Girosi, 1989; Stokbro, Umberger, and Hertz, 1990) and are calculated according to the following equations:

$$a_{ij} = \sum_k \frac{g_k^a(\theta)}{\sum_l g_l^a(\theta)} (w_{kij}^a + z_{kij}^a \cdot c_k^a) \tag{8}$$

$$b_i = \sum_k \frac{g_k^b(\theta)}{\sum_l g_l^b(\theta)} (w_{ki}^b + z_{ki}^b \cdot c_k^b) \tag{9}$$

where w_{kij}^a and w_{ki}^b are parameters, z_{kij}^a and z_{ki}^b are parameter vectors, and c_k^a and c_k^b are vectors whose components are:

$$c_{kl}^a = \frac{\theta_l - \mu_{kl}^a}{\sigma_{kl}^a} \qquad c_{kl}^b = \frac{\theta_l - \mu_{kl}^b}{\sigma_{kl}^b} \tag{10}$$

where μ_{kl}^a, μ_{kl}^b, σ_{kl}^a, and σ_{kl}^b are parameters corresponding to the mean and variance of basis function k along dimension l. The basis functions $g_k^a(\theta)$, $g_k^b(\theta)$ are Gaussian:

$$g_k^a = \exp\left(-\frac{1}{2}\sum_l c_{kl}^{a\,2}\right) \qquad g_k^b = \exp\left(-\frac{1}{2}\sum_i c_{kl}^{b\,2}\right) \tag{11}$$

The parameters w_{kij}^a, w_{ki}^b, z_{kij}^a, z_{ki}^b, μ_{kl}^a, μ_{kl}^b, σ_{kl}^a, and σ_{kl}^b are learned using gradient descent during the action-perception cycle used to train the DIRECT model. The action-perception cycle is induced by instating random joint velocity vectors $\dot{\theta}_B$ (where the subscript B denotes a babbled movement). For the parameters used to produce the a_{ij}, the cost function that is minimized by gradient descent is:

$$H = \sum_i (\dot{\theta}_{Bi} - \dot{\theta}_i)^2 \tag{12}$$

where the $\dot{\theta}_i$ form the joint rotation vector calculated from the RBF network outputs according to Equation 7. The values a_{ij} form the approximate Jacobian pseudoinverse learned by the network. The parameters used to produce the b_{ij}, which form the postural relaxation portion of the network, are adjusted using the cost function:

$$H = \sum_i (\dot{\theta}_{Bi} - \dot{\theta}_i)^2 + \sum_{i=1}^{N} \beta_i \left(\frac{\theta_i - \theta_i^c}{\theta_i^r}\right)^2 \tag{13}$$

where θ_i^c, θ_i^r, and β_i are defined in Equation 6. See Jordan (1990) for a discussion of possible roles for constrained cost functionals such as Equation 13 in motor learning.

The first simulation illustrates the model's performance when tracing a closed path. In this case, a sequence of 100 back-and-forth movements between two targets were simulated. The arm was placed in an "awkward" posture with an extreme wrist angle before moving to the first target. The results of this simulation are shown in Figure 8. The top panel of this figure shows the movements produced by a controller utilizing the MP pseudoinverse. This controller maintains the extreme wrist angle throughout the movement sequence. A similar result would be expected for a controller utilizing the pseudoinverse proposed by Mussa Ivaldi and Hogan (1991). The bottom panel shows the postures reached by the DIRECT model utilizing the postural relaxation scheme described above. Whereas the awkward posture is maintained with the MP pseudoinverse, the DIRECT controller "unfolds" with successive movements, leading to more comfortable postures for both targets. A similar tendency of experimental subjects to only gradually move toward more comfortable postures (in accordance with a minimum cost principle such as Equation 6) was noted by Cruse and Bruwer (1987, p. 143), who stated that the arm

> "...seems to approach angle values which correspond to the minimum cost principle only gradually during the ongoing movement... If there is enough time during the movement, the angle values near the end of the movement correspond quite well to those of the minimum cost principle. In other cases ... even the joint angle values at the end of the movement were different depending on the starting arm position and so revealed the history dependence of the control system."

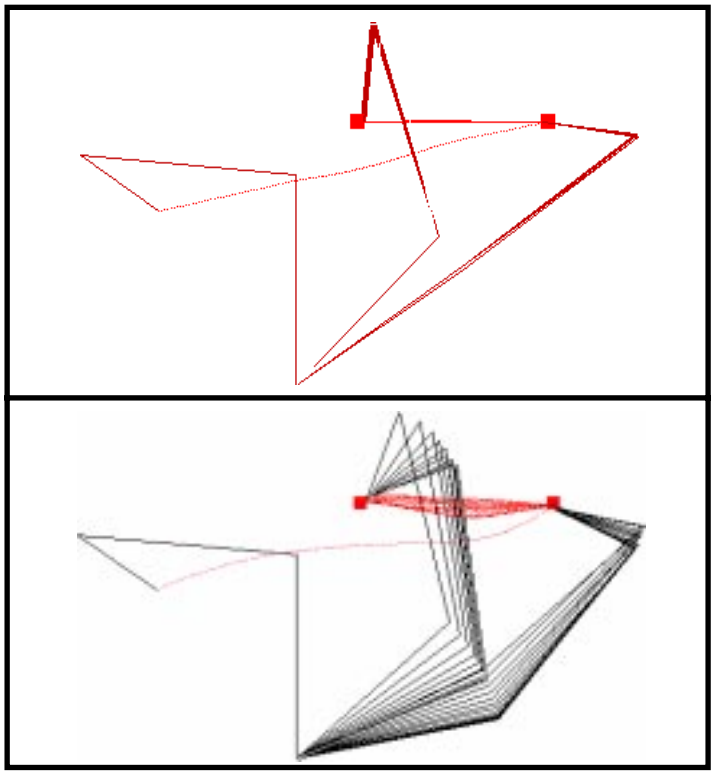

Figure 8: Simulation of back-and-forth movements between two targets starting from an awkward posture with an extreme wrist angle. The top panel shows movements produced by a controller using the MP pseudoinverse. The bottom panel shows movements produced by the DIRECT model utilizing the adaptive postural relaxation scheme. Whereas the awkward posture is maintained with the MP pseudoinverse, the postural relaxation controller "unfolds" with successive movements, converging toward a more comfortable posture for each target.

Another point of interest concerns the slight curvature of the movement trajectories between the two targets for the DIRECT controller. One factor contributing to this curvature is the cost function of Equation 13, which leads to a balance between perfectly straight trajectories (due to the first summation term) and movements toward more "comfortable" postures (due to the second summation term). The second factor is the tendency to produce slightly curved movements in the direction of the long axis of the manipulability ellipsoid discussed in Section 3.

The second simulation, shown in Figure 9, compares the final postures produced by an MP pseudoinverse controller to those of the DIRECT model when reaching to the same target from a number of different initial configurations. As in the experimental results of Cruse (1986), the variability of the final postures produced by the postural relaxation controller is much smaller than the variability produced by the MP pseudoinverse. Again, significant curvature is seen in the trajectories formed by the DIRECT model.

6 Concluding Remarks

This chapter has addressed several issues related to pseudoinverse-style controllers as models of human motor control. Consideration of the problem of automatic compensation for externally imposed constraints on the effector system suggests that movements should be planned in task space coordinates rather than in coordinate frames more closely related to the effector system. This is consistent with experimental studies suggesting spatial planning of reaching movements and acoustic space planning of speech movements. Using a directional mapping to transform desired task space movement directions into effector space movement directions leads to the ability to automatically compensate for constraints on articulators. This kind of control is closely related to robotics techniques that use a pseudoinverse of the Jacobian matrix. Next, potential problems for pseudoinverse-style controllers as models of biological movement control were addressed. Experimental data suggest that the variability of final arm posture for reaches is generally much less than would be predicted by a pure pseudoinverse controller. Although these data are often cited as evidence for the explicit control of final posture in reaching movements, it is shown that pseudoinverse-style controllers that include a form of "postural relaxation" (related to robotics controllers that combine null-space motions with the pseudoinverse and to the biological motor control models of Cruse and colleagues) can account for these data without using explicit postural tar-

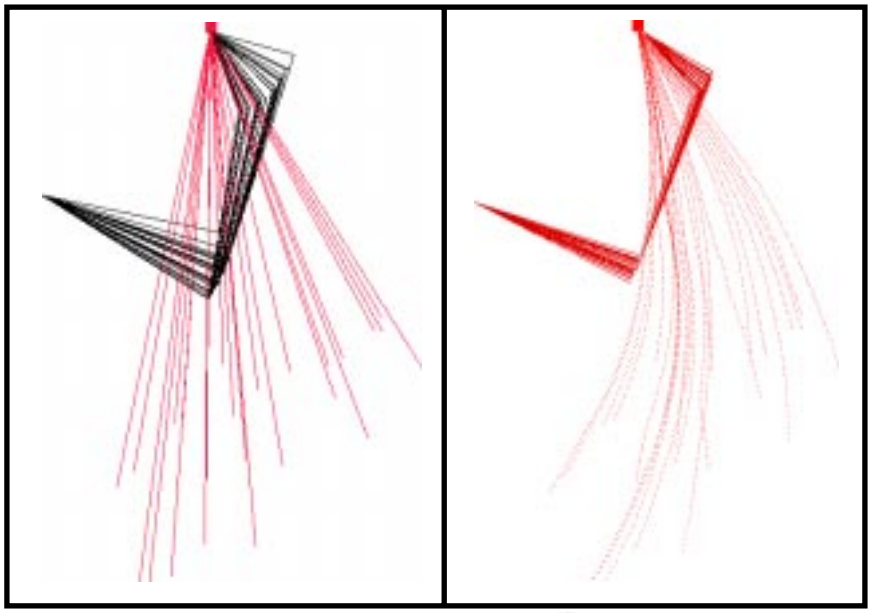

Figure 9: Reaches to the same target position from a variety of initial postures. The left panel shows reaches produced by a controller utilizing the MP pseudoinverse. The right panel shows reaches produced by the DIRECT model utilizing the adaptive postural relaxation scheme. The final postures produced by the DIRECT model cover a much smaller range of joint configurations (centered near a "comfortable" posture for the target) than the pseudoinverse controller, in keeping with the experimental results of Cruse (1986) and others.

gets. Such a controller also maintains the motor equivalence capabilities of pseudoinverse-style controllers that are not inherent to controllers using explicit postural targets. Another potential problem for pseudoinverse-style controllers is that they can produce and indefinitely maintain uncomfortable or extreme joint angles. The postural relaxation scheme was shown to overcome this problem as well. An adaptive neural network controller incorporating postural relaxation was defined and simulated to verify these properties. The model suggests two sources that might contribute to curvature in human reaches: (i) a learning bias toward the long axis of the manipulability ellipsoid, and (ii) a tendency to move toward more comfortable postures.

7 References

Abbs, J.H., & Gracco, V.L. (1984). Control of complex motor gestures: Orofacial muscle responses to load perturbations of lip during speech. *Journal of Neurophysiology*, **51**, pp. 705-723.

Baillieul, J., Hollerbach, J., & Brockett, R.W. (1984). Programming and control of kinematically redundant manipulators. *Proceedings of the 23rd IEEE Conference on Decision and Control*, pp. 768-774. New York: IEEE.

Bailly, G., Laboissière, R., & Schwartz, J.L. (1991). Formant trajectories as audible gestures: an alternative for speech synthesis. *Journal of Phonetics*, **19**, pp. 9-23.

Berkinblit, M.B., Gelfand, I.M., & Feldman, A.G. (1986). A model of the aiming phase of the wiping reflex. In S. Grillner, P.S.G. Stein, D.G. Stuart, H. Forssberg, & R.M. Herman (eds.), *Neurobiology of Vertebrate Locomotion* (pp. 217-227). London: Macmillan.

Bullock, D., Cisek, P.E., & Grossberg, S. (1995). Cortical networks for control of voluntary arm movements under variable force conditions. Boston University Center for Adaptive Systems Technical Report CAS/CNS-95-019. Submitted for publication.

Bullock, D., & Contreras-Vidal, J.L. (1993). How spinal neural networks reduce discrepancies between motor intention and motor realization. In K.M. Newell & D.M. Corcos (eds.), *Variability and Motor Control* (pp. 183-221). Champaign, IL: Human Kinetics Publishers.

Bullock, D., Grossberg, S., & Guenther, F. H. (1993). A self-organizing neural network model for redundant sensory-motor control, motor

equivalence, and tool use. *Journal of Cognitive Neuroscience*, **5**, pp. 408-435.

Burnod, Y., Grandguillaume, P., Otto, I., Ferraina, S., Johnson, P.B., & Caminiti, R. (1992). Visuomotor transformations underlying arm movements toward visual targets: A neural network model of cerebral cortical operations. *Journal of Neuroscience*, **12**, pp. 1435-1453.

Cameron, S. (1995). *Self-organizing neural networks for visual navigation and adaptive control*. Doctoral dissertation, Boston University, Boston.

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**, pp. 2039-2058.

Cruse, H. (1986). Constraints for joint angle control of the human arm. *Biological Cybernetics*, **54**, pp. 125-132.

Cruse, H., & Brüwer, M. (1987). The human arm as a redundant manipulator: The control of path and joint angles. *Biological Cybernetics*, **57**, pp. 137-144.

Cruse, H., Brüwer, M., & Dean, J. (1993). Control of three- and four-joint arm movement: Strategies for a manipulator with redundant degrees of freedom. *Journal of Motor Behavior*, **25**(3), pp. 131-139.

Dean, J., & Brüwer, M. (1994). Control of human arm movements in two dimensions: paths and joint control in avoiding simple linear obstacles. *Experimental Brain Research*, **97**, pp. 497-514.

Desmurget, M., Prablanc, C., Rossetti, Y., Arzi, M., Paulignan, Y., Urquizar, C., & Mignot, J. (1995). Postural and synergic control for three-dimensional movements of reaching and grasping. *Journal of Neurophysiology*, 74, pp. 905-910.

Fiala, J.C. (1995). *Neural network models of motor timing and coordination*. Doctoral dissertation, Boston University, Boston.

Flash, T. (1989). Generation of reaching movements: Plausibility and implications of the equilibrium trajectory hypothesis. *Brain, Behavior, and Evolution*, **33**, pp. 63-68.

Flash, T., & Mussa-Ivaldi, F. (1990). Human arm stiffness characteristics during the maintenance of posture. *Experimental Brain Research*, **82**, pp. 315-326.

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**, pp. 1527-1537.

Guenther, F.H. (1992). *Neural models of adaptive sensory-motor control for flexible reaching and speaking*. Doctoral dissertation, Boston University, Boston.

Guenther, F.H. (1994). A neural network model of speech acquisition and motor equivalent speech production. *Biological Cybernetics*, **72**, pp. 43-53.

Guenther, F.H. (1995a). Speech sound acquisition, coarticulation, and rate effects in a neural network model of speech production. *Psychological Review*, **102**, pp. 594-621.

Guenther, F.H. (1995b). A modeling framework for speech motor development and kinematic articulator control. In K. Elenius & P. Branderud (eds.), *Proceedings of the XIIIth International Congress of Phonetic Sciences*, vol. 2, pp. 92-99. Stockholm: KTH and Stockholm University.

Hollerbach, J.M., Moore, S.P., & Atkeson, C.G. (1986). Workspace effect in arm movement kinematics derived by joint interpolation. In G. Ganchev, B. Dimitrov, & P. Patev (eds.), *Motor Control* (pp. 197-208). New York: Plenum.

Hollerbach, J.M., & Suh, K.C. (1985). Redundancy resolution of manipulators through torque optimization. *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 1016-1021). New York: IEEE.

Hore, J., Watts, S., & Vilis, T. (1992). Constraints on arm position when pointing in three dimensions: Donders' law and the Fick gimbal strategy. *Journal of Neurophysiology*, **68**, pp. 374-383.

Jordan, M.I. (1990). Motor learning and the degrees of freedom problem. In M. Jeannerod (ed.), *Attention and Performance XIII: Motor Representation and Control* (pp. 796-836). Hillsdale, NJ: Erlbaum.

Jordan, M.I., & Rumelhart, D.E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, **16**, pp. 307-354.

Kalaska, J.F., Cohen, D.A.D., Hyde, M.L., & Prud'homme, M. (1989). A comparison of movement direction-related vs. load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *Journal of Neuroscience*, **9**, pp. 2080-2102.

Klein, C.A., & Huang, C. (1983). Review of pseudoinverse control for use with kinematically redundant manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-13**(2), pp. 245-250.

Lacquaniti, F., Soechting, J.F., & Terzuolo, C.A. (1982). Some factors pertinent to the organization and control of arm movements. *Brain Research*, **252**, pp. 394-397.

Liégeois, A. (1977). Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-7**(12), pp. 868-871.

Lindblom, B., Lubker, J., & Gay, T. (1979). Formant frequencies of some fixed-mandible vowels and a model of speech motor programming by predictive simulation. *Journal of Phonetics*, **7**, pp. 147-161.

Merton, P.A. (1972). How we control the contraction of our muscles. *Scientific American*, **226**, pp. 30-37.

Miller, L.E., Theeuwen, M., & Gielen, C.C. (1992). The control of arm pointing movements in three dimensions. *Experimental Brain Research*, **90**, pp. 415-426.

Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, **42**, pp. 223-227.

Mussa-Ivaldi, F.A. (1988). Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neuroscience Letters*, **91**, pp. 106-111.

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

Mussa-Ivaldi, F.A., & Hogan, N. (1991). Integrable solutions of kinematic redundancy via impedance control. *International Journal of Robotics Research*, **10**, pp. 481-491.

Mussa-Ivaldi, F.A., Morasso, P., & Zaccaria, R. (1988). Kinematic networks: A distributed model for representing and regularizing motor redundancy. *Biological Cybernetics*, **60**, pp. 1-16.

Pellionisz, A., & Llinás, R. (1985). Tensor network theory of the metaorganization of functional geometries in the central nervous system. *Neuroscience*, **16**(2), pp. 245-273.

Perkell, J.S., Matthies, M.L., Svirsky, M.A., & Jordan, M.I. (1993), Trading relations between tongue-body raising and lip rounding in production of the vowel /u/: A pilot 'motor equivalence' study. *Journal of the Acoustical Society of America*, **93**, pp. 2948-2961.

Poggio, T., & Girosi, F. (1989). A theory of networks for approximation and learning. AI Memo No. 1140, Massachusetts Institute of Technology.

Raibert, M.H. (1977). Motor control and learning by the state space model. Technical Report AI-M-351, Massachusetts Institute of Technology.

Rosenbaum, D.A., Engelbrecht, S.E., Bushe, M.M., & Loukopoulos, L.D. (1993). Knowledge model for selecting and producing reaching movements. *Journal of Motor Behavior*, **25**, pp. 217-227.

Sabes, P.N., Wolpert, D.M., & Jordan, M.I. (in preparation). Obstacle avoidance in human reaching movements and an inertial criteria for path planning. Massachusetts Institute of Technology Center for Biological and Computational Learning Technical Report.

Saltzman, E.L., & Kelso, J.A.S. (1987). Skilled actions: A task dynamic approach. *Psychological Review*, **94**, pp. 84-106.

Saltzman, E.L., & Munhall, K.G. (1989). A dynamical approach to gestural patterning in speech production. *Ecological Psychology*, **1**, pp. 333-382.

Sanger, T.D. (1994). Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation*, **6**, pp. 29-37.

Savariaux, C., Perrier, P., & Orliaguet, J.P. (1995). Compensation strategies for the perturbation of the rounded vowel [u] using a lip tube: A study of the control space in speech production. *Journal of the Acoustical Society of America*, **98**, pp. 2428-2442.

Scott, S.H., & Kalaska, J.F. (1995). Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *Journal of Neurophysiology*, **73**, pp. 2563-2567.

Soechting, J.F., Buneo, C.A., Herrmann, U., & Flanders, M. (1995). Moving effortlessly in three dimensions: Does Donders' law apply to arm movement? *Journal of Neuroscience*, **15**, pp. 6271-6280.

Stein, P.S.G., Mortin, L.I., & Robertson, G.A. (1986). The forms of a task and their blends. In S. Grillner, P.S.G. Stein, D.G. Stuart, H. Forssberg, & R.M. Herman (eds.), *Neurobiology of Vertebrate Locomotion* (pp. 201-216). London: Macmillan.

Stokbro, K., Umberger, D.K., & Hertz, J.A. (1990). Exploiting neurons with localized receptive fields to learn chaos. *Complex Systems*, **4**, pp. 603-622.

Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement. *Biological Cybernetics*, **61**, pp. 89-101.

Whitney, D.E. (1969). resolved motion rate control of manipulators and human prostheses. *IEEE Transactions on Man-Machine Systems*, **MMS-10**(2), pp. 47-53.

Wolpert, D.M., Ghahramani, Z., & Jordan, M.I. (1994). Perceptual distortion contributes to the curvature of human reaching movements. *Experimental Brain Research*, **98**, pp. 153-156.

Wolpert, D.M., Ghahramani, Z., & Jordan, M.I. (1995). Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Experimental Brain Research*, **103**, pp. 460-470.

Wright, C.E. (1990). Generalized motor programs: Reexamining claims of effector independence in writing. In Jeannerod, M. (ed.), *Attention and Performance XIII: Motor Representation and Control* (pp. 294-320). Hillsdale, NJ: Erlbaum.

Yoshikawa, T. (1985). Manipulability of robotic mechanisms. *International Journal of Robotics Research*, **4**(2), pp. 3-9.