Biol. Cybern. 67, 491–500 (1992)

Biological Cybernetics

Vector field approximation: a computational paradigm for motor control and learning

Ferdinando A. Mussa-Ivaldi and Simon F. Giszter

Department of Brain and Cognitive Sciences, Massachussetts Institute of Technology, Cambridge MA 02139, USA

Received June 1, 1992/Accepted in revised form June 1, 1992

Abstract. Recent experiments in the spinalized frog (Bizzi et al. 1991) have shown that focal microstimulation of a site in the premotor layers in the lumbar grey matter of the spinal cord results in a field of forces acting on the frog's ankle and converging to a single equilibrium position. These experiments suggested that the neural circuits in the spinal cord are organized in a set of control modules that "store" a few limb postures in the form of convergent force fields acting on the limb's end-point. Here, we investigate how such postural modules can be combined by the central nervous system for generating and representing a wider repertoire of control patterns. Our work is related to some recent investigations by Poggio and Girosi (1990a, b) who have proposed to represent the task of learning scalar maps as a problem of surface approximation. Consistent both with this view and with our experimental findings in the spinal frog, we regard the issue of generating motor repertoires as a problem of vector-field approximation. To this end, we characterize the output of a control module as a "basis field" (Mussa-Ivaldi 1992), that is as the vectorial equivalent of a basis function. Our theoretical findings indicate that by combining basis fields, the central nervous system may achieve a number of goals such as *(1)* the generation of a wide repertoire of control patterns and *(2)* the representation of these control patterns with a set of coefficients that are invariant under coordinate transformations.

1 Introduction

In a recent set of experiments (Mussa-Ivaldi et al. 1990; Giszter et al. 1991; Bizzi et al. 1991) in collaboration with E. Bizzi, we have investigated the organization of the motor output in the spinal frog. Briefly, we stimulated with a microelectrode a number of sites in the premotor layers of the gray matter. For each stimulation site, we measured the isometric forces produced by the muscles of the leg at different ankle locations (Fig. 1). Three major results emerged from this study.

1. As shown in Fig. 1D, the microstimulation of a site in the premotor layers of the lumbar gray matter resulted in a convergent field of forces at the ipsilateral ankle. The force vector vanished at a single equilibrium point within the limb workspace. We attributed this convergent pattern of forces to the balanced recruitment of a group of agonist and antagonist muscles acting on the hip and knee joint. Muscles are known to behave as tunable spring-like elements (Rack and Westbury 1969). It seems reasonable to expect the combination of the forces generated by a group of such elements to be a field of elastic forces acting on the ankle and converging to an equilibrium location.
2. Through microstimulation of different spinal interneuronal regions, we elicited different force fields with different equilibrium points.
3. The convergent field added vectorially. When we applied two simultaneous microstimulations to two different spinal sites we obtained a force field that was proportional to the vector sum of the fields generated by the independent stimulation of each site.

These results strongly suggest that the premotor circuitry in the frog's spinal cord is organized in a number of distinct modules. Each module implements a control law which corresponds to the specification of a limb's posture. This control law is expressed by a field of forces with a single equilibrium position. As one adopts this point of view, it is natural to ask if and how the central nervous system may combine a set of output fields for generating other postures as well as more complex motor behaviors.

In the following sections we consider how the vectorial summation of a few force fields can generate a variety of control patterns. We address this issue from the point of view of approximation theory. We begin by defining a component of a desired controller's behavior as a pattern of force vectors over a limb's configuration

Corresondence to: F. A. Mussa-Ivaldi

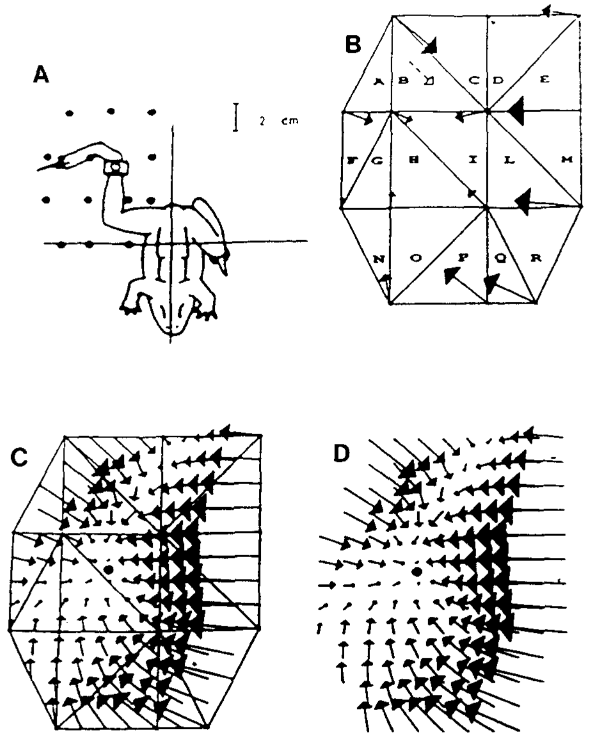

Fig. 1A–D. Force fields obtained from the microstimulation of the frog's spinal cord. **A** The workspace locations at which the ankle forces were measured are indicated by the black dots. **B** The ankle's workspace was partitioned into a set of non-overlapping triangles ($A, B, C, \ldots$). The tested locations shown in A are at the vertices of each triangle. The arrows at the same vertices represent the force vectors measured at different points and at the same latency from the onset of the stimulus. These are actual data. The dashed arrow within the triangle B was derived by linear interpolation from the vectors at the vertices of the same triangle. **C** Interpolated field. The equilibrium point is indicated by a filled circle. **D** Same field as C, without the interpolation triangles. (From Mussa-Ivaldi et al. 1990)

space. In a closely related investigation, Mussa-Ivaldi (1992) has shown that the relevant features of many patterns of vectors can be captured by representing a continuous field as a combination of *basis fields*. Basis fields are the vectorial counterpart of the local basis functions used for reconstructing a scalar map from a set of numerical examples (Powell 1987). Poggio and Girosi (1990a, b) have used the theory of basis functions to characterize the behavior of a wide class of neural networks performing multivariate association maps. In accordance with their views, here we suggest to use basis fields for representing the operation of a network of elementary control modules. A biological example of such an elementary control module may be established by the pattern of connections between a spinal interneuron and a set of motoneurons innervating a group of different muscles. This pattern of connections between spinal interneurons and spring-like muscles has been suggested to be the underlying reason for the observation of convergent force fields after stimulation of the premotor layers in the frog's spinal cord (Bizzi et al. 1991).

More formally, we consider a network of control modules implementing mechanical behaviors which are described by basis fields. The total output of such a network is given by the vectorial summation of these basis fields. Our findings indicate that by combining basis fields a distributed control system may achieve a number of important goals such as *(1)* the generation of a wide repertoire of control patterns and *(2)* the representation of these patterns with a set of coefficients that are invariant under coordinate transformations.

In a sense, we present an extension of the equilibrium-point hypothesis (Feldman 1966; Bizzi et al. 1984; Hogan 1984) to a broader context. According to the equilibrium-point hypothesis, the central nervous system generates and represents the movement of a limb as a temporal sequence of equilibrium positions. One may observe that the notion of equilibrium position corresponds to a particular feature of a field of forces: a stable equilbrium position is a point in space surrounded by a pattern of attractive forces. Here we consider how a variety of force patterns, including (but not limited to) such stable equilibria can be obtained by superimposing the outputs of simple control modules. In particular, our results indicate that the pattern of output forces corresponding to the implementation of an ideal force controller can be generated by combining a number of equilibrium-point controllers. This finding is important because it shows that motor tasks as diverse as moving a limb and manipulating objects can be reduced to a single computational framework based upon the vectorial combination of convergent force fields.

2 Control, planning and vector-field approximation

Let us consider a simplified system consisting of a set of K independent control modules acting in parallel upon the muscles of a multi-joint limb. Each control module establishes the viscoelastic properties of a group of muslces. As a result, each module generates at the interface between the limb and the environment a force-field

$$F^i = \Phi^i(x, u_i)\,. \qquad (1)$$

In the above expression, the variables F^i, x and u_i indicate respectively the N-dimensional output force generated by the controlled muscles, the state of the limb/environment interface and the controller's command variable[1]. We limit our discussion to the static impedance, that is to the relation between force and position at steady state. Our main results however apply to the dynamic impedence as well.

We also make the simplifying assumption that the dimensions of the limb's configuration space do not exceed the dimensions of the position variable, x. In this case, the net force field, $F(x, u_1, \ldots, u_K)$, generated

[1] In this paper, we adopt the convention of using superscripts to indicate vector and matrix *objects* and subscripts to indicate different vector and matrix *components*. For example, in this notation the component of the force generated by the n-th controller along the m-th direction is F^n_m

by the ensemble of all the control modules acting simultaneously and independently on the limb is the sum of all the controller fields, that is:

$$F(x, u_1, \ldots, u_K) = \sum_{i=1}^{I} \Phi^i(x, u_i) . \quad (2)$$

With a redundant serial mechanism, the force field generated by the controllers at the end-point is still well-defined and can be computed analytically (Mussa-Ivaldi and Hogan 1991). In this case however, the net endpoint field must be more generally modeled as a non-linear combination of the controller fields and Eq. (2) does not necessarily characterize the end-point behavior[2]. In contrast, since the muscles operate upon the joints in a parallel arrangement (that is according to a generalized "common position" constraint), the torque fields generated by a set of independent controllers may be assumed to add linearly in configuration space (that is Eq. (2) can still describe the net torque as a funciton of joint configuration). Which rules govern the combination of control modules at the end-point of a kinematically redundant biological limb remains at present an interesting and open question.

We define the *repertoire* of a control network as the set of all possible net fields with all possible values of the control variables. If we denote by $\mathcal{U}_i$ the set of the admissible values for the command variable u_i, the repertoire of the control network is the set:

$$X = \left\{ \sum_{i=1}^{K} \Phi^i(x, u_i) \middle| u_i \in \mathcal{U}_i \right\}$$

Vector fields provide not only a description of the control processes but also a framework for specifying the planning of a desired behavior. In robotics some researchers have proposed to specify the planning of complex tasks by force fields defined over a manipulator's workspace. For example, Kathib (1986) used potential fields to represent the planning of reaching motions within obstacle-ridden environments. He suggested to represent several simultaneous goals by combining potential functions. The surface of an obstacle to be avoided corresponds to a repulsive potential "hill", and a target to be reached corresponds to an attractive potential "valley". Once each goal is encoded by a specific potential function, all these potential functions can be added to derive a total field. Then, the planning of an optimal trajectory becomes equivalent to a problem of gradient descent within this potential-energy landscape. Kathib (1986) and Hogan (1984) suggested to implement the planned potential functions by a set of controllers operating in parallel. According to this view, each controller derives the forces induced by a target or by an obstacle by directly computing the gradient of the corresponding potential field. Thus, the field specified by the planner is faithfully implemented by the controllers.

In this paper we take a different approach. First, we represent the planning as well as the execution of motor tasks by means of vector fields instead of scalar potential fields. A vectorial representation is indeed more general than a scalar representation. It is always possible to derive the former from the latter, whereas the converse is not true. A generic vector field may not be reducible to a potential function[3]. Second, we assume that the biological system has a repertoire, X, generated by the vector combination of independent control modules. This repertoire is limited by the predefined properties of the controllers and of the actuators. In other words we do not assume that any goal is accurately mapped into a corresponding control primitive. Instead, we believe that control primitives are established a-priori by the connections between muscles and neural circuits. Following this view, the execution of an arbitrary motor plan is equivalent to a field-approximation problem: given a planned field, $P(x)$, the problem is to find a field, $F(x) \in X$, which minimizes some norm

$$\|P(x) - F(x)\|^2$$

suitably defined over the limb's state space (see also the section on motor control in Poggio (1990)).

The field-approximation problem assumes a more tractable form if a planned behavior is specified by a finite set of M force vectors, $\{P^1, P^2, \ldots, P^M\}$, defined at M points, $\{x^1, x^2, \ldots, x^M\}$, rather than by a continuous field, $P(x)$. We consider these vectors as samples (or "examples") of a field to be filled in (or "completed") by an appropriate choice of control parameters. The concept of planning by a finite set of vectors is schematically illustrated by the navigation example shown in Fig. 2. Here, the task consists in moving towards a target while avoiding two obstacles. A set of repulsive forces is associated with each obstacle while other force vectors converge toward the target location.

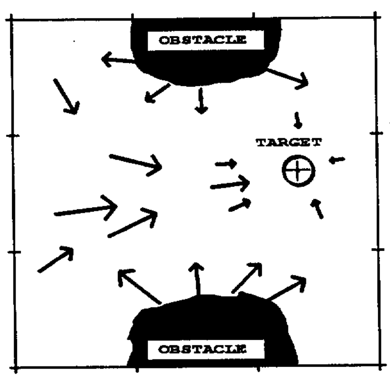

Fig. 2. Planning of a navigation task by a finite set of vectors

[2] The non-linearity induced by serial redundancy can be intuitively understood by considering that the net stiffness, K, of two springs in series with stiffnesses k_1 and k_2 is given by the (nonlinear) geometric mean of k_1 and k_2

[3] A vector field can be derived as the gradient of a potential field if and only if its curl is zero. If this condition is satisfied the field is said to be "irrotational"

3 Symbolic planning and field features

In many instances, the planning of a motor task has been regarded as an intrinsically symbolic process in which simple goals are combined to achieve more complex objectives (Lozano-Perez 1982). Vector fields provide an excellent mathematical framework for symbolic planning. Simple motor goals can be represented by specific *features* of vector fields. Furthermore, the additive property of vector fields provide a simple way to combine such features in order to obtain more complex patterns. Here, we list some field features that can be regarded as a simple "vocabulary" for task planning in 2D. Each feature may be represented by a finite set of vectors as shown in Fig. 3.

Stable equilbrium. The vectors converge toward a point at which the force is zero.
This feature corresponds to the specification of a position-control task.

Impedance. The impedance tensor at an equilibrium point specifies the forces that the controllers must generate in response to external perturbations along different directions. This tensor is fully specified by a set of eigenvectors. Geometrically, the static component of the impedance can be related to a metric tensor (see Sect. 5.1). According to this view, the distance between two points is measured by their difference in potential energy.

Unstable equilibrium. The equilibrium point is surrounded by a set of repulsive forces. Accordingly, the impedance tensor has only positive eigenvectors. This feature is appropriate to represent the goal of avoiding an obstacle in a navigation task.

Saddle point. This is a combination of stable and unstable equilibrium along orthogonal directions.

Uniform field. A field of parallel force vectors with constant amplitude corresponds to the specification of a force-control task. The positional uncertainties associated with kinematics errors and with unpredicted perturbations do not affect the output-force vector.

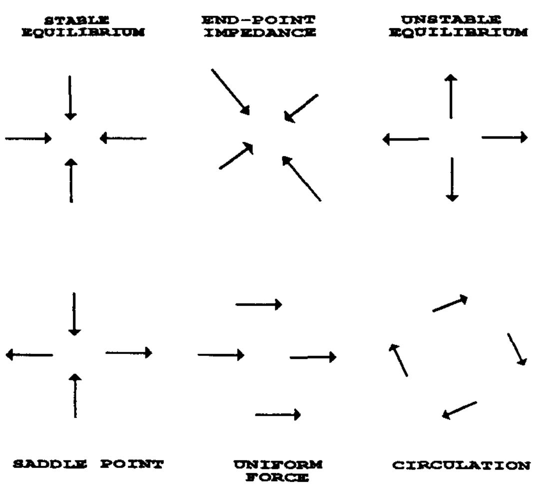

Fig. 3. Some 2*D* field features (see text)

Circulation. Circulating force patterns are induced by an antisymmetric component in the impedance tensor. This pattern, combined with the stable and the unstable equilibrium, can be used to specify a limit-cycle behavior. We also consider the circulation pattern for the sake of completeness, since it cannot be efficiently represented as the gradient of a scalar potential.
We will consider how these different features can be reproduced and combined by summing the output fields generated by a set of control networks.

4 Control modules as basis fields

Let us focus on a simplified form for the output function (1) implemented by each control module, by setting

$$\Phi^i(x, u_i) = u_i \phi^i(x)\,. \tag{3}$$

We call this form *linear tuning*: the control variable, u_i plays the role of a scaling factor applied to a (nonlinear) output field. In this case, the repertoire of the control network is the linear span:

$$X = \left\{ \sum_{i=1}^{K} u_i \phi^i(x) \middle| u_i \in \mathcal{U}_i \right\}$$

and we may apply the methods of linear algebra as outlined in Mussa-Ivaldi (1992) for approximating an arbitrary pattern of vectors. Given a set of M pattern vectors, $P^1, \ldots, P^M$, specified at M distinct locations, $x^1, \ldots, x^M$, the approximation goal is to find a set of control parameters such that

$$\sum_{i=1}^{K} u_i \phi^i(x^j) \sim P^j$$

for $j = 1, \ldots, M$. In and N-dimensional vector space, this goal is expanded into a system of MN linear equations (one per data component) in K unknowns. This system of equations can be compactly written as:

$$\Phi u = \hat{P} \tag{4}$$

where we have introduced the (unknown) control vector

$$u = (u_1, u_2, \ldots, u_K)\,,$$

the matrix

$$\begin{bmatrix} \phi_1^1(x^1) & \phi_1^2(x^1) & \ldots & \phi_1^K(x^1) \\ \phi_1^1(x^2) & \phi_1^2(x^2) & \ldots & \phi_1^K(x^2) \\ & \ldots & & \\ \phi_1^1(x^M) & \phi_1^2(x^M) & \ldots & \phi_1^K(x^M) \\ & \ldots & & \\ & \ldots & & \\ \phi_N^1(x^1) & \phi_N^2(x^1) & \ldots & \phi_N^K(x^1) \\ \phi_N^1(x^2) & \phi_N^2(x^2) & \ldots & \phi_N^K(x^2) \\ & \ldots & & \\ \phi_N^1(x^M) & \phi_N^2(x^M) & \ldots & \phi_N^K(x^M) \end{bmatrix},$$

and the "pattern vector"

$$\hat{P} = (P_1^1, P_1^2, \ldots, P_1^M, P_2^1, P_2^2, \ldots, P_2^M, \ldots, P_N^1, P_N^2, \ldots, P_N^M) \,. \quad (5)$$

The approximation problem (4) can be solved for any set of pattern vectors if the output fields, $\phi^i(x)$, are linearly independent. In this case, the output fields form a basis for a K-dimensional functional space and we call them *basis fields*. It is possible to derive basis fields from scalar basis functions as detailed in (Mussa-Ivaldi 1992)[4]. Following this approach, a general representation of a continuous field is obtained by combining K_I irrotational basis fields and K_S solenoidal basis fields:

$$F(x) = \sum_{i=1}^{K_I} c_i \varphi^i(x) + \sum_{i=1}^{K_S} d_i \psi^i(x) \,.$$

An irrotational basis field, $\varphi^i(x)$, has zero curl and is the gradient of some scalar basis function, $g_i(x)$, that is

$$\varphi^i(x) = \nabla g_i(x) \quad \text{and} \quad \text{curl}(\varphi^i) = 0$$

A solenoidal basis field, $\psi^i(x)$, has zero divergence. In the Euclidean metric, an irrotational basis field is obtained from a basis function, $g_i(x)$, by applying the operator $A\nabla$, where A is an antisymmetric[5] matrix:

$$\psi^i(x) = A\nabla g_i(x) \quad \text{and} \quad \text{div}(\psi^i) = 0 \,.$$

This statement can be demonstrated for $x \in \Re^N$, with $N \geqslant 2$. In particular, with $N = 2$ and $N = 3$ an antisymmetric matrix, A, corresponds to a 90° rotation operator.

The set of control coefficients, c_i and d_i, obtained from the approximation of a field feature can be regarded as a *representation* of the corresponding planning goal within the repertoire of the controller network.

5 A simulation example

We now proceed to illustrate and develop the concepts introduced in the previous sections by applying them to a simplified arm model (Fig. 4): a two-joint planar mechanism operated by a set of K independent control networks. Each control network modulates a torque/angle relation in configuration space. Indicating by $Q = (Q_1, Q_2)$ the joint-torque vector and by $q = (q_1, q_2)$ the angular configuration of the arm, we express the torque field generated by the i-th control network as:

$$Q = u_i \tau^i(q)$$

where, u_i is a tuning parameter. Here, we do not impose any constraint on the tuning parameter, which we assume to be a real number ranging between $-\infty$ and $+\infty$. Thus, the repertoire of our model is the linear span:

$$\left\{ \sum_{i=1}^{K} u_i \tau^i(q) \,|\, -\infty < u_i < +\infty \right\}$$

We also consider the fields generated by the controllers in the cartesian coordinate system which defines the position of the hand, $r = (x, y)$ with respect to a pair of orthogonal axes centered at the shoulder (Fig. 4). In this coordinate system, the controllers generate a force vector $F = (F_x, F_y)$ for each position of the hand $x = \mathscr{L}(q)$. Since our model is not kinematically redundant, the repertoire of the control networks can be directly expressed in hand coordinates as the set of force fields

$$\left\{ \sum_{i=1}^{K} u_i \phi^i(x) \,|\, -\infty < u_i < +\infty \right\}$$

In the above expression, $\phi^i(x)$ is the hand-force field generated by the i-th controller when $u_i = 1$. This field is computed within the arm's workspace using the inverse of the Jacobian matrix $J(q) = \partial \mathscr{L} / \partial q$:

$$\phi(x) = [J(q)^{-1}]^T \tau^i(q) \,.$$

5.1 Gaussian controllers

We start by assuming that each control network generates an irrotational basis field given as the gradient of a potential function, $G_i(q)$, in configuration space. In this example, we use the multivariate Gaussian

$$G(q, q0^i) = -g_0 \exp((q - q0^i)^T K^i (q - q0^i)/2) \,. \quad (6)$$

where K^i is a negative definite symmetric matrix.

The term g_0 is a constant with the physical dimension of an energy ($[g_0] = [M][L]^2[T]^{-2}$). In the following discussion we will neglect this term by assuming that it is numerically equal to one. The torque field corresponding to (6) is:

$$\tau^i(q) = \nabla_q G(q, q0^i) = K^i(q - q0^i) G(q, q0^i) \,. \quad (7)$$

Each control network generates a field of torques converging toward an equilibrium configuration, $q0^i$. At this equilibrium configuration, the joint stiffness matrix is K^i. As shown in Fig. 5A, the torque field reaches a maximum of intensity within an annular region surrounding $q0^i$. Figure 5B illustrates the same control field in hand coordinates. The distortion introduced around the equilibrium point by the nonlinear kinematics of the model arm is revealed by the deformation of the isopotential lines.

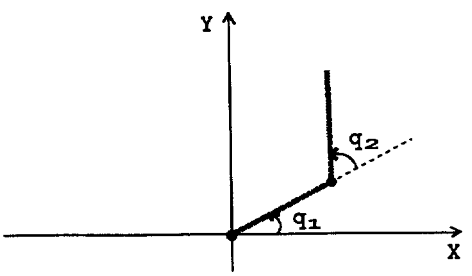

Fig. 4. A two-joint planar arm

[4] See also (Wahba 1982) for a similar generalization of splines to the approximation of vector fields on the sphere

[5] A matrix, $A = [a_{i,j}]$ is said to be antisymmetric if $a_{i,j} = -a_{ji}(a_{i,i} = 0)$

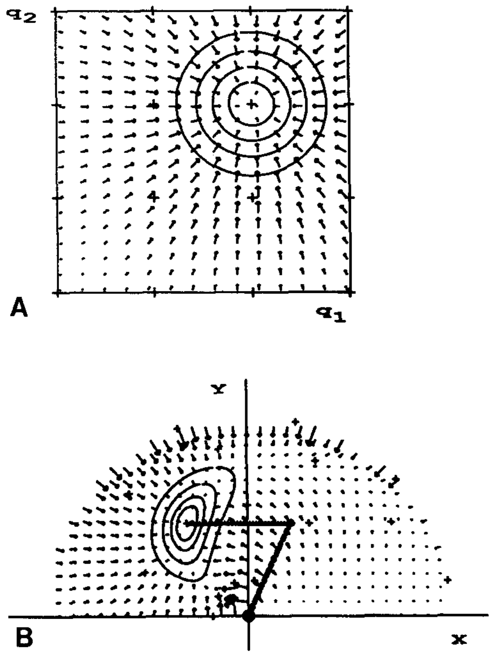

Fig. 5. A, B. Gaussian controller in joint coordinates. **A** The output field of a controller is plotted as a set of torque vectors at several joint configurations. The configuration space of the two-joint system shown in Fig. 4 has the intrinsic geometrical structure of a torus. Here, we consider a patch over the configuration space and for graphical convenience we adopt a "flattened" cartesian representation for this patch. The circles represent a set of isopotential lines. Different controllers have different equlibrium configurations. Here, the crosses indicate the equilibrium configurations implemented by 16 different controllers. These configurations are arranged in a rectangular grid. **B** The same field represented in cartesian hand-coordinates

These basis fields represent a nonlinear form of position control in configuration space. Each controller is completely specified by an equilibrium configuration, $q0^i$, and by a stiffness matrix, K^i around this configuration. The control variable, u_i, changes the local stiffness, $(\partial\tau^i/\partial q|_{q0_i} = u_i K^i)$, while leaving unaffected both the equilibrium configuration and the "range of action" of the controller.

Geometrically, the matrix K^i, can be regarded as the metric tensor of a coordinate transformation:

$$q = Wq'$$

with $K^i = -(W^T W)$. In the coordinate system of q', the potential is the radial basis function:

$$G'(q', q0'^i) = -\exp(-(q' - q0'^i)^T(q' - q0'^i)/2)\,. \qquad (8)$$

In this geometrical interpretation, the stiffness matrix is equivalent to using a weighted norm for transforming the radially asymmetric potential function (6) into the corresponding symmetric form (8) (Poggio and Girosi 1990b). In the context of function approximation, using this weighted norm corresponds to assigning different degrees of importance to approximation errors along different directions. Interestingly, Poggio and Girosi (1990a, b) have derived a class of neural networks that are capable of finding the optimal weighting matrix, that is the optimal metric on the input space, from the minimization of an error functional. With respect to motor control, such a weighted-optimization scheme would lead to a set of control modules with anisotropic stiffness matrices.

5.2 Control parameters

Each Gaussian potential (6) is specified by two control parameters: the equilibrium point, $q0^i$, and the matrix K^i. These parameters could be left free to change and reach some optimal value, depending on the set of data to be approximated and on the cost functional to be minimized (Poggio and Girosi 1990a). Here we are merely concerned with the issue of generating a repertoire of combinations after these parameters have been somehow determined.

In our simulations we have placed the equilibrium points at the nodes of a rectangular grid in joint space (Fig. 5A). The spacing between nodes is given by the number of controllers and by the total angular domain covered by the grid. For example we may divide the shoulder range, Δq_1, in M steps and the elbow range, Δq_2, in N step. In this case, we obtain a lattice with $M \times N$ nodes separated by $\delta q_1 = \Delta q_1/M$ along the shoulder axis and by $\delta q_2 = \Delta q_2/N$ along the elbow axis.

Once the joint-space grid is established (that is, once we have decided the angular excursions and the number of controllers for each simulation) we determine the controller matrices K^i. We specify these matrices so that each controller generates a torque field with maximum amplitude at the four neighboring nodes; a result that is achieved by setting $K^i_{1,1} = 1/(\delta q_1)^2$, $K^i_{2,2} = 1/(\delta q_2)^2$, $K^i_{1,2} = K^i_{2,1} = 0$. Of course this is an arbitrary choice which we adopt here to deal with a particularly simple and regular array of controllers. Other interesting choices could include non-diagonal K-matrices and/or random distributions of the equilibrium points. We would like to stress that the ability of basis fields to approximate an arbitrary pattern of forces is contingent upon the existence of a significant overlap between their range of action. This overlap is obviously achieved by our choice of parameters.

5.3 Generation of field features

We now consider the task of approximating in hand coordinates the two-dimensional field features shown in Fig. 3. Each feature is expressed by five force vectors that the hand is expected to exert when placed at four distinct workspace locations. These vectors can also be regarded as "examples" needed to learn a motor task from the association between system configuration and required output force. We focus on the two force patterns labeled as "stable equilibrium" and "uniform field". These two patterns specify two complementary control prescriptions: positions control (the stable equilibrium) and force control (the uniform field).

Let us begin by considering a system of four independent control networks. Each control network generates an end-point field as shown in Fig. 6. With five

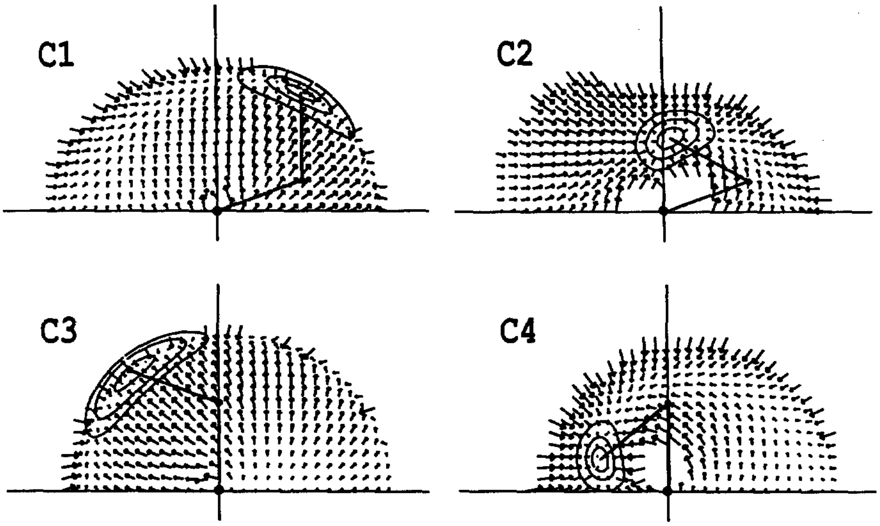

Fig. 6. Four irrotational basis fields in hand coordinates. For each basis field, the arm is plotted at the equilibrium configuration. The equilibrium point is surrounded by three isopotential lines. The equilibrium configurations in degrees at the shoulder and at the elbow are: $q0^1 = [20, 70]$; $q0^2 = [20, 130]$; $q0^3 = [90, 70]$; $q0^4 = [90, 130]$. The matrices K^i were set with $K^i_{1,1} = -0.672$, $K^i_{2,2} = -0.908$, $K^i_{1,2} = K^i_{2,1} = 0$

two-dimensional pattern vectors, the interpolation problem (Eq. (4)) can be exactly solved only if there are at least $K = 2 \times 5 = 10$ basis fields. With four control networks, the control vector, $u = (u_1, u_2, u_3, u_4)$, can only approximate the desired patterns. To this end, we may use the Moore-Penrose solution of Eq. (4), namely:

$$u = \Phi^+ \hat{P} = (\Phi^T \Phi)^{-1} \Phi^T \hat{P} .$$

This solution corresponds to minimizing the square error

$$\sum_{j=1}^{5} \left\| \sum_{i=1}^{4} u_i \phi^i(x^j) - \hat{P}^j \right\|^2$$

at the points, $x^1, \ldots, x^5$.

Figure 7B shows the results of the least-squares approximation for the convergent force pattern of Fig. 7A. Four target vectors are directed toward a desired equilibrium point, which is specified by a fifth vector with zero amplitude. Note that this point is different from each of the four equilibrium points implemented by the control networks. The outcome of the approximation is a field with the desired position-control features not only within the domain of the data but also in the surrounding region. In particular, the approximating field is smooth and has only one attractor. In neural-network terms, this ability to generate a satisfactory output beyond the set of data (or "examples") is usually referred to as "generalization".

An even more evident case of generalization was obtained by aproximating a parallel pattern of forces with the same controllers (Fig. 8A and B). This result demonstrates that a nearly ideal force control can be achieved within a region of the workspace by combining as few as four position controllers. Remarkably, the converse in not true. It is not possible to generate a position-control field by a simple superposition of (no matter how many) constant-force fields. This fact seems to question a conventional engineering

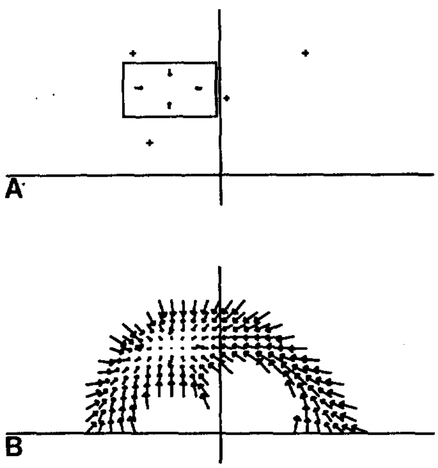

Fig. 7A, B. Approximation of a convergent force pattern. **A** The pattern vectors are included within the rectangular box. Here, as in the next figures, the crosses indicate the equilibrium points of the basis fields. **B** The approximating field. For graphical reasons we plot only the vectors whose amplitude does not exceed a predefined threshold

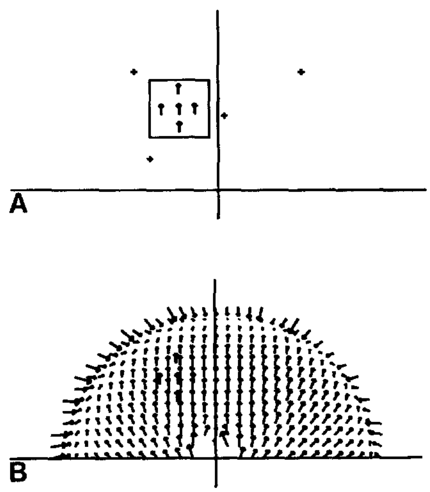

Fig. 8A, B. Approximation of a parallel force pattern. **A** force pattern and controller equilibria. **B** Approximating field wisdom according to which an ideal actuator should generate a uniform force independent of a system's state, i.e. a parallel force field. In contrast, our findings indicate that by modulating a set of spring-like actuators a network of control modules may achieve a wide repertoire of behaviors using a simple superposition mechanism.

5.4 Solenoidal controllers

With four irrotational basis fields we were also able to successfully approximate other simple vector patterns such as saddle points and unstable equilbria. However, as we expected the approximation failed when we tried to reproduce a set of circulating vectors (Fig. 9A and B). A drastic improvement, in terms both of the errors and of the smoothness of the approximating field, was obtained by replacing two of the four controllers with their solenoidal counterpart (Fig. 9C and D). To obtain

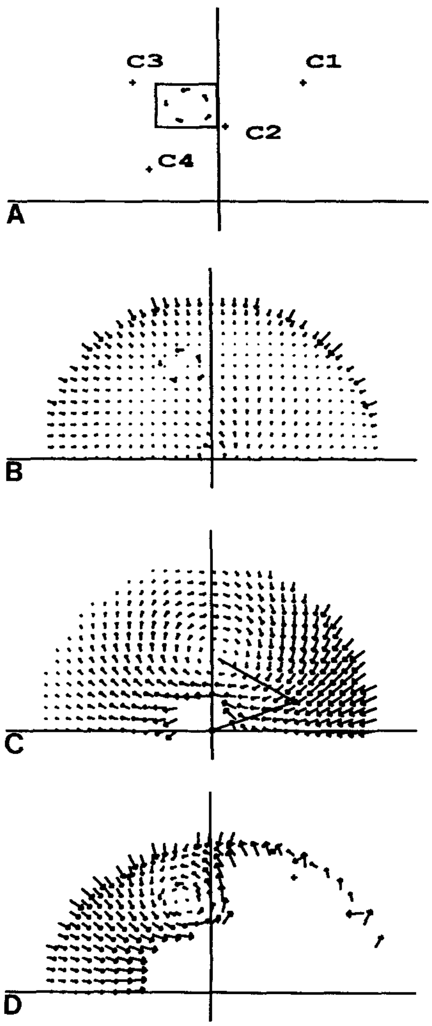

Fig. 9A–D. Approximation of circulatory pattern. **A** Force pattern and controller equilibria. **B** Result of the approximation with irrotational controllers. **C** The solenoidal controller obtained by modifying *C*2. **D** Approximating field obtained by replacing *C*2 and *C*3 with solenoidal controllers

these solenoidal controllers we multiplied the angle/torque relation (7) by the matrix[6]

$$\begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.$$

This antisymmetric matrix represents a simple feedback law in which different degrees of freedom are coupled with unbalanced gains. For example, one may consider two muscles acting on two different joints. Let us assume that the static behavior of each muscle is given by a length-tension function, $f_i(l_i, u_i)$, with $i = 1, 2$. The only way to obtain an asymmetry in the partial derivatives, $\partial f_1/\partial l_2$ and $\partial f_2/\partial l_1$, is to introduce a reflex coupling, $u_1 = \chi_1(l_2)$ and $u_2 = \chi_2(l_1)$, such that $(\partial f_1/\partial u_1)\,(\partial \chi_1/\partial l_2) \neq (\partial f_2/\partial u_2)(\partial \chi_2/\partial l_1)$.

The goal achieved by adding solenoidal basis fields is to extend significantly the repertoire of the control networks. After this addition, the repertoire includes both rotational and irrotational output fields. In a biological context, it remains to be established whether or not the motor system can actually generate rotational fields. An experimental investigation of multijoint stiffness in humans (Mussa-Ivaldi et al. 1985), has suggested a negative answer: the hand stiffness measured in several subjects maintaining the hand in different workspace locations was consistently found to be symmetric. However, one should keep in mind that in these experiments (1) the fields were measured only in the proximity of the equilibrium posture of the hand and (2) the only task considered was the maintenance of hand postures. Therefore one cannot yet rule out that the biological motor system may generate force fields with non-zero curl in different tasks involving movements or forceful interactions with the environment.

5.5 Combination of field features

So far, we have shown that a number of different control features can be locally reproduced by the superposition of a few basis fields. Another crucial issue is related to the possibility of combining such local features for obtaining more complex control patterns.

As one increases the number of features to be simultaneously represented one also needs to increase the number of basis fields – that is the dimension of the functional space spanned by their combination. Figure 10 shows that the combination of field features is successfully implemented by adding controllers with different equilibrium positions. In this case, the goal was to produce a position control task in a region *A* and a force control task in a different workspace region, *B*. To achieve this goal we used 15 basis fields whose equilibrium configurations were placed on a regular 5×3 grid over the joint space. In this grid, the shoulder angle is incremented by steps of 30 deg. starting from

[6] Strictly speaking, the basis fields resulting from this operation are not solenoidal, because the metric of the joint space is not Euclidean. However, their curl is different from zero. Therefore they have at least a non-zero solenoidal component

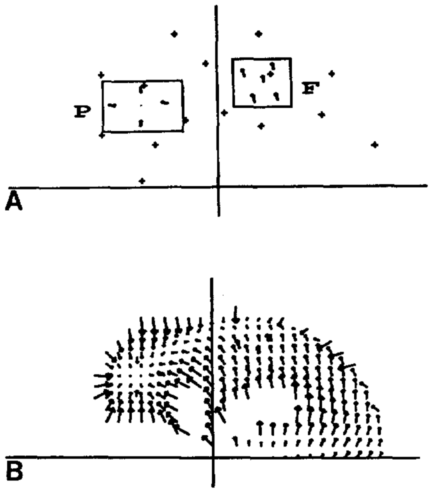

Fig. 10A, B. Combination of field features. A A convergent (P) and a parallel (F) pattern are presented simultaneously to a system with 15 controllers. B Approximating field

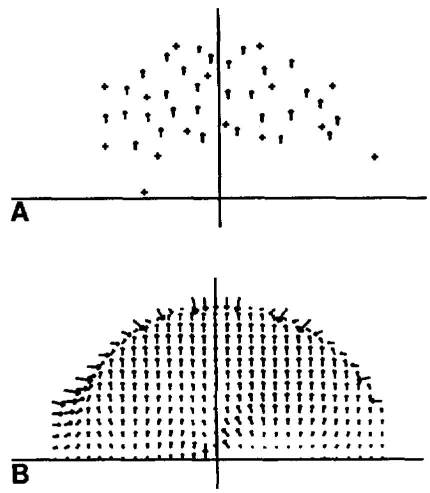

Fig. 11A, B. Approximation of a pattern defined over a wide domain. A Vector pattern and controller equilibria. The controllers are the same as those shown in Fig. 10. B The approximating field

-10 deg. and ending at 110 deg. The elbow angle is incremented by steps of 40 deg. starting from 50 deg. and ending at 130 deg. All the matrices K^i are set to

$$\begin{bmatrix} -8.2 & 0 \\ 0 & -2.06 \end{bmatrix}$$

(see Sect. 5.2). With a total of 10 pattern vectors, this set of controllers leads to a system of 20 equations in 15 unknowns. The resulting least-squares solution not only captures the patterns within the regions P and F but also generates a smooth transitions between these two regions.

5.6 Global versus local

A complementary issue to the problem of combining field features is the problem of enforcing a single feature over a larger domain. In spite of their conceptual difference, these two problems can be represented in the same way. In order to increase the domain of a feature, one may simply specify a larger set of vectors. For example, Fig. 11A shows a set of 30 parallel vectors which taken together specify a force control task over a wide region of the workspace. This pattern is successfully implemented by the superposition of only 15 controllers (Fig. 11B). This finding is somewhat surprising if one considers that *(1)* the approximation problem is highly overconstrained (60 equations in 15 unknowns) and *(2)* the desired pattern does not bear any resemblance to the basis fields generated by the controllers.

6 Discussion

Our investigation originated from some recent experimental findings (Bizzi et al. 1991) which suggested that the neural circuits in the frog's spinal cord are organized in a number of functional modules. Each module generates a field of static forces defined over a limb's workspace by recruiting a balanced group of agonist and antagonist muscles. Typically, this field is characterized by a pattern of forces converging to a single equilibrium point. From this observation we concluded that each spinal module implements a control law corresponding to a limb's equilibrium posture.

Remarkably, the experimental data have also indicated that the simultaneous activation of two modules may result in a form of vector superposition of the respective output fields. In this paper, we explored the competence of such a field-superposition mechanism to generate a variety of control patterns. Our work has been influenced by some recent investigations of Poggio and Girosi (1990b) on network learning and approximation. These authors have proposed to represent the task of learning scalar maps from sparse examples as a problem of surface reconstruction. They regarded the examples as samples of an unknown surface that can be approximated by the weighted summation of a finite set of pre-defined basis functions.

We have addressed the issue of generating motor repertoires as a field-approximation problem (Poggio 1990). To this end we characterized the output of a single control module as a "basis field", that is as the vectorial equivalent of a basis function. We found that convergent force fields can be combined to generate a variety of patterns, including locally parallel patterns of forces. Functionally, the implication of this result is that a set of position controllers may be combined to obtain force control. The converse is not true: it is not possible to obtain a position-control law by combining a set of parallel force fields.

More generally, the functional form of the basis fields has a significant impact on the ability to generate and combine a set of desired field patterns. In this paper, we have considered the particular example of a set of fields generated by Gaussian potentials. The same success would not have been met if instead we had used quadratic potentials whose gradients are linear torque-angle relations. The reason is that, unlike Gaussians the quadratic functions span a low-dimensional space. Thus, contrary to intuition the purpose of linear independence is better served by a set of non-linear controllers!

We have also found that by introducing solenoidal basis fields (that is basis fields with zero divergence) a system's repertoire may become complete, including rotational and irrotational vector patterns. However, it remains to be seen whether or not such an increase of repertoire is a desirable feature from the standpoint of control. For example, Colgate (1988) and Colgate and Hogan (1989) have demonstrated that a necessary and sufficient condition for a manipulator coupled with a passive environment[7] to be stable is that the output impedance of the manipulator be passive as well. Translated in terms of vector fields, this condition implies requiring that the output field be irrotational.

Finally, we would like to make an observation regarding basis fields and coordinate transformations. As shown in Mussa-Ivaldi (1992), by combining basis fields one obtains for each field feature a set of control coefficients which are invariant under coordinate transformations. The invariance arises from the fact that in a change of coordinates, only the basis fields are transformed. Therefore, in the new coordinate system the net vector field is a linear combination of the transformed basis fields with the same set of coefficients. For instance, in our simulations we approximated a number of force patterns in end-point coordinates. That is we had derived the control coefficients by combining the end-point fields generated by the controllers to approximate the desired end-point pattern of vectors. The same control coefficients were then used to modulate torque/angle fields in joint coordinates. Thus, the computational burden associated with the existence of different coordinate systems is entirely placed in the transformation of the basis fields from one system to another. Once this transformation is carried out, the entire control repertoire of a system can be expressed by a single "dictionary" of coefficients. As a practical consequence, the problem of learning new motor tasks can be carried out equivalently either in the gemoetrical space of the controllers, after mapping the planned patterns into the controllers' coordinates, or in the geometrical space of the motor planner, after mapping the set of basis fields into planning coordinates.

[7] Loosely speaking a mechanical system is said to be passive if it cannot deliver more energy than what it has received. For a more precise definition, see Colgate (1988) and Colgate and Hogan (1989)

Acknowledgments. This work has been prompted by experiments done in collaboration with E. Bizzi. We wish to thank T. Flash, F. Girosi, N. Hogan, T. Poggio and R. Shadmehr for their useful comments. This research was supported by NIH grants NS09343 and AR26710 and by ONR grant N00014/88/K/0372.

References

Bizzi E, Accornero N, Chapple W, Hogan N (1984) Posture control and trajectory formation during arm movement. J Neuorsci 4:2738–2744

Bizzi E, Mussa-Ivaldi FA, Giszter SF (1991) Computations underlying the execution of movement: A biological perspective. Science 253:287–291

Colgate E, Hogan N (1989) An analysis of contact instability in terms of passive physical equivalents. IEEE Proc Int Conf Robot Automat :404–409

Colgate E (1988) The control of dynamically interacting systems. PhD thesis, M.I.T. Department of Mechanical Engineering

Feldman AG (1966) Functional tuning of the nervous system during control of movement or maintenance of a steady posture. ii. Controllable parameters of the muscles. iii. Mechanographic analysis of the execution by man of the simplest motor task. Biophysics 11:565–578, 766–775

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1991) Equilibrium-point mechanisms in the spinal frog. In: Arbib MA, Ewert JP (eds) Visual structures and integrated functions. Springer, Berlin Heidelberg New York, pp 223–237

Hogan N (1984) An organizing principle for a class of voluntary movements. J Neurosci 4:2745–2754

Kathib O (1986) Real-time obstacle avoidance for manipulators and mobile robots. Int Robot Res 5:90–99

Lozano-Perez T (1982) Task planning. In: Brady M, Hollerbach JM, Johnson TL, Lozano-Perez T, Mason MT (eds) Robot motion: planning and control. MIT Press, Cambridge

Mussa-Ivaldi FA (1992) From basis functions to basis fields: Using vector primitives to capture vector patterns. Biol Cybern 67:479–489

Mussa-Ivaldi FA, Hogan N (1991) Integrable solutions of kinematic redundancy via impedance control. Int J Robot Res 10:481–491

Mussa-Ivaldi FA, Hogan N, Bizzi E (1985) Neural, mechanical and geometrical factors subserving arm posture in humans. J Neuorsci 5: 2732–2743

Mussa-Ivaldi FA, Giszter SF, Bizzi E (1990) Motor-space coding in the central nervous system. Cold Spring Harbor Symp Quant Biol The Brain 55:827–835

Poggio T (1990) A theory of how the brain might work. Cold Spring Harbor Symp Quant Biol The Brain 55:899–910

Poggio T, Girosi F (1990a) Networks for approximation and learning. Proc. IEEE 78: 1481–1497

Poggio T, Girosi F (1990b) A theory of networks for learning. Science 247:978–982

Powell MJD (1987) Radial basis functions for multivariable interpolation: a review. In: Mason JC, Cox MG (eds) Algorithms for approximation. Clarendon Press, Oxford

Rack PMH, Westbury DR (1969) The effects of length and stimulus rate on tension in the isometric cat soleus muscle. J Physiol 204: 443–460

Wahba G (1982) Vector splines on the sphere, with application to the estimation of vorticity and divergence from discrete noisy data. In: Schempp W, Zeller K (eds) Multivariate approximation theory. Birkhauser Verlag, Basel Boston Stuttgart, pp 407–429