Proceedings of the 2000 IEEE
International Conference on Robotics & Automation
San Francisco, CA • April 2000

Robots can teach people how to move their arm

F. A. Mussa-Ivaldi

Sensory Motor Performance Program
Rehabilitation Institute of Chicago
Physiology, Northwestern University
Chicago, IL 60611

J. L. Patton

Sensory Motor Performance Program
Rehabilitation Institute of Chicago
PM&R, Northwestern University
Chicago, IL 60611

Abstract

We describe a new theoretical framework for robot-aided training of arm movements. This framework is based on recent studies of motor adaptation in human subjects and on general considerations about adaptive control of artificial and biological systems. We propose to take advantage of the adaptive processes through which subjects, when exposed to a perturbing field, develop an internal model of the field as a relation between experienced limb states and forces. The problem of teaching new movements is then reduced to the problem of designing force fields capable of inducing the desired movements as after-effects of the adaptation triggered by prolonged exposure to the fields. This approach is an alternative to more standard training methods based on the explicit specification of the desired movement to the learner. Unlike these methods, the adaptive process does not require explicit awareness of the desired movement as adaptation is uniquely concerned with restoring a preexisting kinematic pattern after a change in dynamical environment.

1 Introduction

Robot-aided training promises to be an important application of robotics, particularly in the field of rehabilitation [1]. One critical problem is to develop the most effective training algorithm. Although there are many possible approaches to robot-aided training, in this paper we describe a novel approach that draws on the current state of knowledge of how humans learn to control movement.

An increasing number of investigations both in robotics [2, 3] and in neurobiology [4-7] have pointed out the necessity and the role of internal representations of dynamics in the control and learning of complex motor tasks. If past motor experience is used to develop a model of a limb's dynamics, then the motor system may take advantage of this model to generate new motor commands that may handle situations not yet encountered.

Strong experimental evidence for the biological relevance of internal models has been offered by a set of recent experiments that involved the adaptation of arm movements to the perturbing force fields generated by an instrumented manipulandum [5, 8-12]. The major findings of these studies are as follows: 1) when exposed to a complex but deterministic field of velocity-dependent forces, arm movements are first distorted but, after prolonged practice, the initial kinematics are recovered; 2) if, at the end of adaptation, the field is suddenly removed, after effects are visible as approximate mirror images of the initial perturbations; 3) adaptation is achieved by the CNS through the formation of a local map that associates the states (positions and velocities) visited during the training period with the expected force; 4) following adaptation this map - that is the internal model of the field undergoes a process of consolidation as revealed by evidence of retention during subsequent exposure to the same field and of retrograde interference if a different field was presented shortly after the first adaptation.

In this paper we discuss the possibility of exploiting these biological mechanisms underlying motor adaptation and the formation of internal models for assisting human motor learning. Our goal is to develop a new paradigm for facilitating the recovery of motor skills lost to stroke or other neuromotor disabilities. In designing a training procedure, we wish to take advantage of the adaptive process that is automatically triggered by changes in the dynamics of the environment. As adaptation leads to the recovery of the initial "intended" movement kinematics, the motor system must update the controller that generates the commands to the neuromuscular apparatus. In the following sections we present a general framework for guiding the subjects toward a predetermined kinematic goal that is unknown to the subject. This framework effectively defines a unified approach to adaptation (that aims at restoring previously learned behaviors in novel environments) and skill learning (that aims at acquiring new behaviors in an unchanged environment).

2 Human-robot interaction

Consider a subject interacting with a robot, as shown in Figure 1. The interaction point is the endpoint of the robot, where a force sensor is mounted. Let the position of this endpoint be described by a vector x measured with respect to some inertial frame. The coupled dynamics of subject and robot is described by the system

$$\begin{cases} E(s,\dot{s}) = -F \\ M(s,\dot{s},t) = F \end{cases} \qquad (1)$$

where $s = (x, \dot{x})$ is the state of the interaction point, F is the interaction force measured by the sensor mounted on the robot's end effector, E is the interface field generated by the robot and M is the interface field generated by the subject. These equations may be compacted into a single one when the interaction force is not explicitly considered, i.e.

$$M(s,\dot{s},t) + E(s,\dot{s}) = 0 \qquad (2)$$

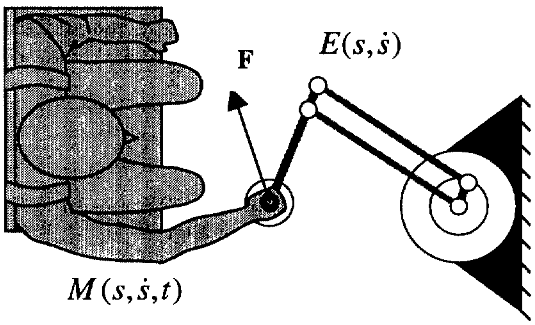

Figure 1: Human/Robot interaction. Two PMI torque motors actuate the robot via a four-bar linkage with a maximum torque of 8.8 Nm, and force at the handle is measured with a six-degree-of-freedom load sensor (ATI F/T Gamma 30/10). Digital optical encoders (Teledyne Gurley) are mounted on each motor and measure its angular position. Analog tachometers (PMI) are also mounted inside the housing of the motors and measure angular velocity. A PC acquires the signals and controls torque.

We assume that the robot field *E* is fully specified by a program. In particular, when the field is set to zero (null field), the system's behavior is fully described by the (unknown) subject's dynamics and has a solution, $\hat{s}(t)$ corresponding to the subject's *unperturbed trajectory*, which satisfies identically the equation

$$M(\hat{s}(t),\dot{\hat{s}}(t),t) = 0 \qquad (3)$$

3 Adaptation

With the above notation, the process of adapting to an external field as reported by Shadmehr and Mussa-Ivaldi [5] is summarized as follows:

1. As subjects are first exposed to a field $E(s,\dot{s})$, they generate a *perturbed trajectory* $s_P(t)$ defined by the equation:

$$M(s_P(t),\dot{s}_P(t),t) + E(s_P(t),\dot{s}_P(t)) = 0 \qquad (4)$$

2. After prolonged exposure to the robot field, subjects develop the adapted dynamics, $\overline{M}(s,\dot{s},t)$, that coupled with the robot's field restores the original trajectory, $\hat{s}(t)$:

$$\overline{M}(\hat{s}(t),\dot{\hat{s}}(t),t) + E(\hat{s}(t),\dot{\hat{s}}(t)) = 0 \qquad (5)$$

It is convenient to describe the adapted dynamics as a change of the original one:

$$\overline{M}(s,\dot{s},t) = M(s,\dot{s},t) + \Delta M(s,\dot{s},t) \qquad (6)$$

With this definition of ΔM and combining (5) with (3) one obtains the differential equation

$$\Delta M(s,\dot{s},t) + E(s,\dot{s}) = 0 \qquad (7)$$

that admits $\hat{s}(t)$ as a solution.

3. After adaptation is complete, if the robot field is suddenly reset to zero one observes the *aftereffect* trajectory, $s_A(t)$, that satisfies the equation

$$M(s_A(t),\dot{s}_A(t),t) + \Delta M(s_A(t),\dot{s}_A(t),t) = 0 \qquad (8)$$

4 Internal models

Step 2 and 3 of the adaptive process described above do not immediately imply that subjects learn an internal model of the robot field. They merely state that subjects learn to cancel the forces generated by the robot along the trajectory $\hat{s}(t)$. This would be sufficient to restore the trajectory and to observe after-effects when the field is suddenly removed.

One might suggest that, instead of developing a model of the robot field, that is of the dependency of the experienced force upon s and $\dot{s}$, subjects instead learn to play back the temporal sequence of forces $\Delta M(t) = -E(\hat{s}(t),\dot{\hat{s}}(t))$ that satisfy Equation (7) along $\hat{s}(t)$. This hypothesis was contradicted by experiments of Conditt et al. [10] in which subjects were first trained to execute rectilinear reaching movements in a perturbing field. As adaptation was complete, subjects were asked to execute circular hand movements passing through the same region of space where the training had occurred. They found evidence of complete generalization: the circular movements (and their after effects) after training with reaching movements were not distinguishable from the same movements executed after training in the field with circular movements. Therefore, the simultaneous adaptation to different trajectories passing through the same region of state space is not compatible with the development of a single time-dependent compensation. This observed generalization implies that the variation of the controller *tends to approximate an algebraic identity* with the perturbing field[1]. We define the *domain of proper generalization* to be the region of state space S within which all trajectories are correctly adapted to the external field. The domain of proper generalization is the region of state space where what amounts to a model of the actual dependence of force upon state has formed. The experimental observations of Gandolfo et al. [13] suggest that this domain is limited to a small region of state space containing the trajectories that have been practiced in the field.

5 Adaptive training

The goal of adaptive training is to design a force field, $E_T(s,\dot{s})$, that leads to the execution of the *desired movement*, $s_D(t)$, as an after-effect of adaptation. This goal is concisely stated by requiring that:

$$\overline{M}(\hat{s}(t),\dot{\hat{s}}(t),t)+E_T(\hat{s}(t),\dot{\hat{s}}(t))=0 \tag{9}$$

(Equation 5)

and

$$\overline{M}(s_D(t),\dot{s}_D(t),t)=0 \tag{10}$$

(Equation 8)

After introducing the representation (6) for the dynamics change, one has an equation that describes the subject's behavior after training in the field,

$$\Delta M(\hat{s}(t),\dot{\hat{s}}(t),t)+E_T(\hat{s}(t),\dot{\hat{s}}(t))=0 \tag{11}$$

and an equation that describes the subject's behavior following removal of the field (i.e., after-effects),

$$M(s_D(t),\dot{s}_D(t),t)+\Delta M(s_D(t),\dot{s}_D(t),t)=0 \tag{12}$$

The problem of adaptive training is that of finding a force field $E_T(s,\dot{s})$, which satisfies (11) and induces a change ΔM compatible with (12). The solution of this problem involves two intermediate steps: 1) the development of a model of the unknown subject's dynamics over a domain of proper generalization S that includes both $\hat{s}(t)$ and $s_D(t)$; 2) the design of an assistive field that, when applied for the first time induces the desired movement $s_D(t)$ as a perturbation of $\hat{s}(t)$.

5.1 System Identification

This first step may be carried out in a variety of ways. A simple approach to this system identification task may focus on a limited region of state space defined by the subject's current and desired movements. Within this region, the robot applies in sparse and random order a number of perturbing fields, $E_1,E_2,\ldots,E_N$, as the subject attempts repeatedly to generate $\hat{s}(t)$. During each perturbation the resulting trajectory $s_i(t)$ as well as the interaction force $F_i(t)$ are recorded. These data are then used to set up the parameters of a dynamic model. We developed a modeling procedure based on two considerations: (a) the trajectory generated by the subject is dynamically stable [14] and (b) the geometrical structure of the subject's limb is known within a degree of uncertainty over a set of L kinematics and dynamic parameters, $p=(p_1,p_2,\ldots,p_L)$. Following these considerations, a model of arm dynamics may be formulated as a linear combination of K independent fields, $\mu_i(s,\dot{s},t)$ modulated by K real non-negative coefficients m_i

$$\hat{M}(s,\dot{s},t)=\sum_{i=1}^{K} m_i\cdot\mu_i(s,\dot{s},t) \tag{13}$$

Each field, μ_i, is a complete dynamic model (Figure 2) with particular values of dynamical parameters and is capable of generating (for example by a simple PD control mechanism) the subject's unperturbed or *intended movement*, $\hat{s}(t)$. That is:

$$\mu_i(\hat{s}(t),\dot{\hat{s}}(t),t)=0 \qquad \forall i \tag{14}$$

This insures that $\hat{M}(\hat{s}(t),\dot{\hat{s}}(t),t)=0$ and $\hat{s}(t)$ are dynamically stable [15]. The fields μ_i are constructed so that the (unknown) best parameter fit to the actual subject's dynamics lies in the convex domain spanned by their parameters. Using (13) it is straightforward to derive a linear least-squares estimate of the tuning coefficients m_i by minimizing the quadratic error,

$$\left\| \sum_{j=1}^{M}\sum_{i=1}^{K} m_i\cdot\mu_i(s_j(t),\dot{s}_j(t),t)-F_j(t)\right\|^2 \tag{15}$$

given an appropriate norm*. And indeed the purpose of this particular procedure is to use a linear least-square

[1] The simultaneous adaptation of all the trajectories, $s(t)$, within a domain, S, of the state space is compatible either with the formation of a single algebraic model $\Delta M(s,\dot{s})$ of the perturbing field or with a compensatory mechanism modulated by the trajectory $s(t)$: $\Delta M(s,\dot{s},s(t))$. However, the latter case is equivalent to state that $\Delta M(s,\dot{s},s(t))+E(s,\dot{s})=0$ is identically satisfied for all trajectories $s(t)\subset S$. That is,

$$\Delta M(s(t),\dot{s}(t),s(t))+E(s(t),\dot{s}(t))=0$$

for all $s(t)\subset S$. But then we can re-write the function $\Delta M(s(t),\dot{s}(t),s(t))$ as $\Delta M(s(t),\dot{s}(t))$ without loss of generality. The statement that the algebraic equation

$$\Delta M(s(t),\dot{s}(t))+E(s(t),\dot{s}(t))=0$$

is identically satisfied for all trajectories in a region S is equivalent to the statement that the differential equation $\Delta M(s,\dot{s})=E(s,\dot{s})$ is also an algebraic identity in S.

* A convenient norm is obtained by defining the inner product of two vector-valued time functions, a(t) and b(t) (defined over the interval [0,T]) as an extension of the Euclidean inner-product, i.e.,

$$<a,b>\equiv\int_0^T a^T(t)\cdot b(t)dt$$

method to deal with nonlinear estimation problem, avoiding time-consuming algorithms that involve iterative searches.

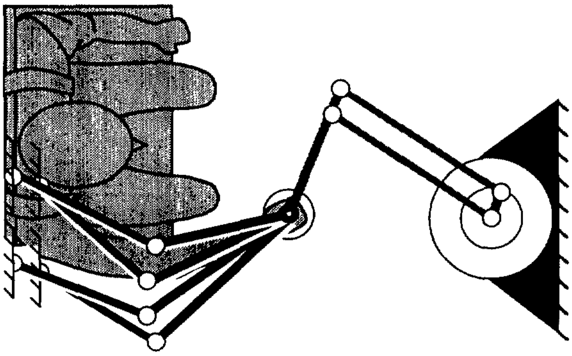

Figure 2: The subject's arm dynamics are modeled as a linear combination of K independent fields, each generated a slightly different estimate of the subject's arm and controller.

5.2 Assistive force field design

The purpose of an assistive field is to calculate a perturbation E_A that results in the desired trajectory $s_D(t)$ *before the onset of an adaptive process*. This corresponds to the condition stated in Equation (4):

$$M(s_D(t), \dot{s}_D(t), t) + E_A(s_D(t), \dot{s}_D(t)) = 0 \quad (16)$$

A simple way to implement a perturbing field is by linearly combining P non-linear viscoelastic fields, $\varphi_i(s)$.:

$$E_A(s) = \sum_{i=1}^{P} c_i \varphi_i(s) \quad (17)$$

For the sake of implementation simplicity, one may reasonably constrain the assistive field to depend only on the state, s and not on its derivative. This choice is also motivated by the fact that earlier experimental studies of force field adaptation did only involve fields depending on hand position and velocity [5, 12]. One choice that has proven successful in our preliminary testing is the regional viscosity field,

$$\varphi_i(s) = \dot{x}e^{-xKix} \quad (18)$$

where x is the Cartesian position of the hand and K is a constant that dictates the (inverse) variance of the regional influence of the field.

A least-squares estimate of the tuning coefficients c is obtained by minimizing:

$$\left\| \sum_i c_i \varphi_i(s_D(t)) - \hat{M}(s_D(t), \dot{s}_D(t), t) \right\|^2 \quad (19)$$

5.3 Training (resistive) force field design

It is straightforward to show that for a desired trajectory that belongs to the same domain of proper generalization as the unperturbed trajectory, the training field that induces the desired movement as an after-effect is simply the opposite of the assistive field described above, i.e.:

$$E_T(s) = -E_A(s) \quad (20)$$

Indeed, with this choice, following adaptation we have:

$$\Delta M(\hat{s}(t), \dot{\hat{s}}(t), t) - E_A(\hat{s}(t)) = 0 \quad (21)$$

(Equation 11)

Within the domain of proper generalization, the field ΔM algebraically approximates the perturbing field. That is

$$\Delta M(s, \dot{s}, t) \approx \Delta\hat{M}(s) = E_A(s) \quad (22)$$

Therefore, if $s_D(t)$ is in the same domain, Equation (12) becomes

$$M(s_D(t), \dot{s}_D(t), t) + \Delta M(s_D(t), \dot{s}_D(t), t) \approx M(s_D(t), \dot{s}_D(t), t) + E_A(s_D(t)) = 0 \quad (23)$$

that is equivalent to Equation (16), which is satisfied by construction.

6 Experiments

To date, four subjects have been tested in a preliminary evaluation. We attempted to cause healthy subjects to move along a new, arbitrarily chosen desired trajectory: a curve to the right and back (dotted lines in Figure 3). No information was given about this desired movement. Four subjects held the handle of the manipulandum and made reaching movements straight outward to a to a target (then return). In all, 350 movements were made, broken down into seven phases as follows:

1) 25 unperturbed trials,
2) 25 trials with 8 intermittent viscous disturbances (used for force field design),
3) 25 more unperturbed trials,
4) 25 trials with 8 intermittent assistive field (perturbations toward the desired movement),
5) 200 trials training with the resistive field (perturbations away from the desired movement),
6) 25 trials of the resistive field, but with 8 intermittent null field ("catch") trials (testing after-effects), and
7) 25 unperturbed trials to test the retention of the after-effects. We measured position, velocity, and handle forces at 100 Hz.

This computational framework was competent to design force fields that caused the subjects' trajectories to be centered on the desired trajectory (Figure 3, dotted lines). Initial exposure to the force field resulted in trajectories that curved away from the desired trajectory (Figure 3A), but after many trials the trajectories recovered to the straight-line shape (Figure 3B). After-effects following training (Figure 3C) resembled the desired trajectory (dotted line). It is important to note that movements were along the desired trajectory even though they were never given any knowledge about the desired trajectory, and the forces from the robot were turned off. Finally, the subjects' after effects typically "washed out" within 5 trials after the removal of the force field (Figure 3D).

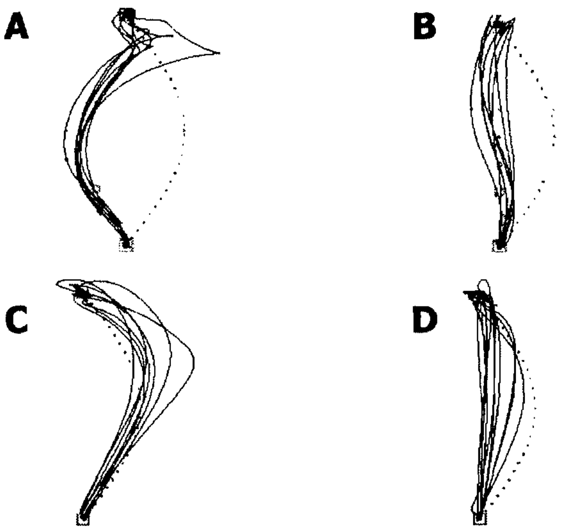

***Figure 3: 20 cm movements** outward (solid lines) from one subject that test the field design. The dotted line in each plot indicates the "desired trajectory." The later stages of the experiment are shown. A and B show the beginning and end of training. C shows the after-effects, and D shows the persistence of after-effects after the field was turned off.*

7 Summary and Conclusions

We are currently exploring this procedure as a tool for motor rehabilitation following stroke and other neurological injuries. The use of robotic devices in rehabilitation has been pursued by other researchers who have used these devices for assessing motor performance and for guiding the patients along the desired movement paths [1]. The approach presented here is different in that, initially, the field generated by the robot pushes the learner's performance away from the intended movement. In this regard, it is a resistive rather than assistive field. Although the above principles prove to be sound, it's actual therapeutic value remains to be determined. A critical step in this process is the demonstration that the improvements obtained as after effects can be made to persist and to consolidate after training. What is not clear is whether the wash-out seen in these healthy subjects (Figure 3D) will also exist in pathological situations where the result of training is a more normal movement. Preliminary studies in our laboratory on stroke subjects have revealed that after-effects did persist for many trials when the after-effects resembled healthy movements. However, not all patients may benefit from this type of procedure, which requires that the neural areas that preside over motor adaptation be fully functional. Current evidence indicates that these areas may include the cerebellum as a critical component [16]. It is then plausible to expect that patients with cerebellar lesions may not be able to adapt to externally imposed force fields.

A final consideration concerns the domain of generalization associated with motor learning. A recent computational study [17] suggested that the most efficient motor learning might occur through the combination of independent elementary controllers (or "motor primitives") with localized regions of influence (also called "receptive fields"). The important observation behind this idea is that if a motor primitive has a large receptive field, then when the primitive's parameters are modified so as to improve performance in one region of state space, unwanted interference effects are likely to occur in other regions belonging to the same receptive field. In contrast to this analysis, some studies of the biological motor system have revealed the existence of motor primitives with large receptive fields. In particular, electrophysiological studies involving the stimulation of muscles and of the spinal cord in spinalized frogs [18] indicated a) that the focal stimulation of a site in the lumbar spinal cord results into the activation of multiple muscles acting on the ipsilateral leg; and b) that synergistic muscle recruitment generates a field of viscoelastic forces over a broad region of the leg workspace. The hypothesis that interneuronal circuits in the spinal cord organize a set of well-defined muscle synergies was recently supported by a computational analysis of electromyographic (EMG) activities induced in frogs by cutaneous stimulation of the leg [19]. Taken together, these studies indicate that motor commands generated by the brain are not directed at controlling the forces of individual muscles or single joint torques. Instead, these commands modulate the viscoelastic force fields produced by specific sets of muscles [20]. These force fields have influence over broad regions of the limb state space as each muscle within a synergy contributes a significant force over a large region of state space. A risky consequence of broad tuning is that indeed human subjects display a significant degree of negative interference when adapting to a novel force field [5, 9]: after adaptation is completed in a region of workspace, aftereffects are clearly observable in different regions.

Nevertheless, this interference, characteristic of human motor learning, may have a payoff in terms of stability. The large area of influence of muscle viscoelastic properties provides an efficient means to counteract disturbances over broad regions. The approach to motor learning presented in this paper depends critically upon the presence of large regions of generalization where learning leads to the development of a correct model of the applied field. In this regard, understanding which movements can be successfully learned through the adaptive interaction with a robot will provide us with new critical data about the extent of the receptive fields of biological motor primitives.

Acknowledgements. This work was supported by NIH-NINDS (NS35673-01), NSF (9900684) and the Ralph and Marion C. Falk Research Trust Fund.

References

[1] H. Krebs, N. Hogan, M. Aisen, and B. Volpe, "Robot-aided neurorehabilitation," *IEEE Transactions on Rehabilitation Engineering*, vol. 6, pp. 75-87, 1998.

[2] C. Atkeson and S. Schaal, "Robot learning from demonstration," presented at Machine Learning: Proceedings of the fourteenth international conference (ICML '97), 1997.

[3] R. Burridge, A. Rizzi, and D. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *International Journal of Robotics Research*, vol. 18, pp. 534-555, 1999.

[4] C. Atkeson, "Learning arm kinematics and dynamics," *Annual Review of Neuroscience*, vol. 12, pp. 157-183, 1989.

[5] R. Shadmehr and F. A. Mussa-Ivaldi, "Adaptive representation of dynamics during learning of a motor task," *Journal of Neuroscience*, vol. 14, pp. 3208-3224, 1994.

[6] M. Kawato and D. Wolpert, "Internal models for motor control," *Novartis Foundation Symposium*, vol. 218, pp. 291-304, 1998.

[7] D. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Networks*, vol. 11, pp. 1317-1329, 1998.

[8] N. Bhushan and R. Shadmehr, "Computational nature of human adaptive control during learning of reaching movements in force fields.," *Biological Cybernetics*, vol. 81, pp. 39-60, 1999.

[9] T. Brashers-Krug, R. Shadmehr, and E. Bizzi, "Consolidation in human motor memory," *Nature*, vol. 382, pp. 252-255, 1996.

[10] M. A. Conditt, F. Gandolfo, and F. A. Mussa-Ivaldi, "The motor system does not learn the dynamics of the arm by rote memorization of past experience," *Journal of Neurophysiology*, vol. 78, pp. 554-560, 1997.

[11] M. A. Conditt and F. A. Mussa-Ivaldi, "Central representation of time during motor learning," *Proceedings of the National Academy of Sciences*, vol. in press, 1999.

[12] T. Flash and I. Gurevich, "Arm stiffness and movement adaptation to external loads," *Proceedings of the Annual Conference on Engineering in Medicine and*, vol. 13, pp. 885-886, 1992.

[13] F. Gandolfo, F. A. Mussa-Ivaldi, and E. Bizzi, "Motor learning by field approximation," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, pp. 3843-6, 1996.

[14] J. Won and N. Hogan, "Stability properties of human reaching movements," *Experimental Brain Research*, vol. 107, pp. 125-36, 1995.

[15] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for nonlinear systems," *Automatica*, vol. 34, pp. 683-696, 1998.

[16] W. Thach, H. Goodkin, and J. Keating, "The cerebellum and the adaptive coordination of movement," *Annual Review of Neuroscience*, vol. 15, pp. 403-442, 1992.

[17] S. Schaal and C. Atkeson, "Constructive incremental learning from only local information," *Neural Computation*, vol. 10, pp. 2047-2084, 1998.

[18] S. Giszter, F. Mussa-Ivaldi, and E. Bizzi, "Convergent force fields organized in the frog's spinal cord," *Journal of Neuroscience*, vol. 13, pp. 467-491, 1993.

[19] M. Tresch, P. Saltiel, and E. Bizzi, "The construction of movement by the spinal cord," *Nature Neuroscience*, vol. 2, pp. 162-167, 1999.

[20] F. Mussa-Ivaldi, S. Giszter, and E. Bizzi, "Linear combinations of primitives in vertebrate motor control," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 91, pp. 7534-7538, 1994.