Proceedings of the 2005 IEEE
9th International Conference on Rehabilitation Robotics
June 28 - July 1, 2005, Chicago, IL, USA

Arm Movement Experiments with Joint Space Force Fields using an Exoskeleton Robot*

Michael Mistry[1], Peyman Mohajerian[2], Stefan Schaal[1,3,4]
Departments of [1]Computer Science, [2]Biomedical Engineering, [3]Neuroscience
University of Southern California, 3641 Watt Way, Los Angeles, CA 90089-2520, USA
[4]*ATR Computational Neuroscience Laboratories, Kyoto, Japan*
{*mmistry,mohajeri,sschaal*}*@usc.edu*

***Abstract*— A new experimental platform permits us to study a novel variety of issues of human motor control, particularly full 3-D movements involving the major seven degrees-of-freedom (DOF) of the human arm. We incorporate a seven DOF robot exoskeleton, and can minimize weight and inertia through gravity, Coriolis, and inertia compensation, such that subjects' arm movements are largely unaffected by the manipulandum. Torque perturbations can be individually applied to any or all seven joints of the human arm, thus creating novel dynamic environments, or force fields, for subjects to respond and adapt to. Our first study investigates a joint space force field where the shoulder velocity drives a disturbing force in the elbow joint. Results demonstrate that subjects learn to compensate for the force field within about 100 trials, and from the strong presence of aftereffects when removing the field in some randomized catch trials, that an inverse dynamics, or internal model, of the force field is formed by the nervous system. Interestingly, while post-learning hand trajectories return to baseline, joint space trajectories remained changed in response to the field, indicating that besides learning a model of the force field, the nervous system also chose to exploit the null space to minimize the effects of the force field on the realization of the endpoint trajectory plan. Further applications for our apparatus include studies in motor system redundancy resolution and inverse kinematics, as well as rehabilitation.**

I. Introduction

Robot manipulanda have become a common tool for aiding scientists interested in the study of human motor control [1],[2],[3]. In particular, force field experiments are a popular technique used for determining the mechanisms underlying motor planning, execution, and learning. In these experiments, robotic manipulanda apply controlled, extraneous forces/torques either at the hand [4],[5],[6],[7] or individual joints [10] while the subjects carry out movement tasks, such as point-to-point reaching movements [4],[5],[11] or continuous patterns [6]. However, because of the mechanical constraints of the manipulanda used, these experiments have been limited to two degrees-of-freedom (DOF) movements, focusing on shoulder and elbow joints, and thus not allowing for any spatial redundancy in a movement. A new experimental platform allows us to explore a wider variety of movements, including movements in full 3-D space using the major seven degrees of freedom of the human arm. By applying torques on any or all seven joints, we can create joint-space force fields and attempt to answer questions relevant to human motor control, such as: In which coordinate system does motor planning take place? How does the brain solve the inverse kinematic problem? Does the motor system utilize any optimization? In addition to applications for human motor system understanding, the exoskeleton has the potential use for arm rehabilitation, where 3-D movements can be executed while applying different impedance characteristics for each joint.

In this paper we first describe our robot platform and discuss the control laws required for human exoskeleton use. Next we describe the force field experiment, and show results of one such experiment. Finally we conclude with a discussion of our behavioral results, and examine the exoskeleton's potential use for arm rehabilitation.

II. Methods

A. Apparatus

The experimental platform is a seven DOF hydraulically actuated exoskeleton robot arm (Sarcos Master Arm, Sarcos, Inc., Salt Lake City) (Fig. 1). Its anthropomorphic design mimics the major seven DOFs of the human arm[1], such that any joint movement of the user's arm will be approximately reflected by a similar joint movement of the exoskeleton. Conversely, torques applied by an exoskeleton joint will be reflected to the corresponding joint of the user's arm. The exoskeleton's most proximal joint is mounted to a fixed platform, and the user wields the device by holding on to a handle at the most distal joint and by a strapping of the forearm to the equivalent link of the robot just before the elbow. The shoulder remains unconstrained, but is positioned such that the three shoulder rotation axes of

*This research was supported in part by National Science Foundation grants ECS-0325383, IIS-0312802, IIS-0082995, ECS-0326095, ANI-0224419, a NASA grant AC#98-516, an AFOSR grant on Intelligent Control, the ERATO Kawato Dynamic Brain Project funded by the Japanese Science and Technology Agency, and the ATR Computational Neuroscience Laboratories.

[1]The major joints of the human arm are labeled: shoulder-flexion-extension (SFE), shoulder-abduction-adduction (SAA), humeral rotation (HR), elbow-flexion-extension (EFE), wrist-supination-pronation (WSP), wrist-flexion-extension (WFE) and wrist-abduction-adduction (WAA).

Fig. 1. Sarcos Master Arm with user.

the exoskeleton approximately intersect with the human's shoulder joint.

1) Control Architecture: Our control architecture consists of independent PD servo controllers at each joint (implemented on individual Sarcos Advanced Joint Controller analog circuit boards), with additional feed-forward torque commands computed on a centralized controller running on 2 Motorola PPC 603 parallel processors with the commercial real-time operating system vxWorks (Windriver Systems). Potentiometers and load cells at each joint are sampled at 960 Hz to provide positional and torque feedback, respectively. Joint velocities and accelerations are computed numerically by differentiating the position signal. The signals are filtered with a 2nd order Butterworth filter with cutoff frequency of 33.6 Hz for position, velocity, and torque and 4.8 Hz for acceleration. The acceleration signal requires more aggressive filtering because of the noise amplified by the numerical differentiation. The centralized controller updates the feed-forward commands and PD set points at 480 Hz.

2) Control Laws: Ideally, the user should not be burdened (unintentionally) by the exoskeleton while executing movements. Therefore our control scheme needs to compensate for the exoskeleton's gravity, Coriolis, and inertia forces. Typically, robot manipulanda use force sensors at the endeffector to measure the forces applied by the user [7],[14], and apply control laws that generate a particular impedance characteristic at the point where the user holds the manipulandum. In contrast to these previous approaches, which focused on investigations of the endeffector movement of human arms, our current robotic setup is to examine joint level effects of human motor control. Thus, we do not have a constraint in the form of a desired impedance at the endeffector, but rather impedance control for every DOF. Since we cannot reliably attach force sensors between the human arm and every link of the robot, we resorted to a model-based control approach.

a) Low Level Joint Control: The low level joint controllers are governed by the equation

$$\mathbf{u} = \mathbf{K}_D(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \mathbf{K}_P(\mathbf{q}_d - \mathbf{q}) + \mathbf{u}_{\text{ff}}, \tag{1}$$

where $\mathbf{u} \in \mathbb{R}^7$ is the vector of motor command torques, and $\mathbf{q}, \dot{\mathbf{q}} \in \mathbb{R}^7$ are the vectors of joint position and velocity, and $\mathbf{K}_D, \mathbf{K}_P$ are diagonal gain matrices. The desired joint position and velocity vectors $\mathbf{q}_d, \dot{\mathbf{q}}_d \in \mathbb{R}^7$ and the feed-forward torque command $\mathbf{u}_{\text{ff}} \in \mathbb{R}^7$, are set by the centralized controller. Ideally, we would set $\dot{\mathbf{q}}_d = \dot{\mathbf{q}}$ and $\mathbf{q}_d = \mathbf{q}$ in this controller, and compute an inverse dynamics based feed-forward command, $\mathbf{u}_{\text{ff}}$, to eliminate inertial, Coriolis, and gravity forces. However, as such a controller is neutrally stable and difficult to realize in light of inevitable modeling errors in the dynamics model, a more prudent control approach is required. In order to keep a small amount of a position reference for enhanced stability, we define

$$\mathbf{q}_d^{n+1} = \mathbf{q}_d^n + \epsilon(\mathbf{q}^n - \mathbf{q}_d^n), \tag{2}$$

where n is the discrete time step of the control loop, and ϵ is a constant factor between 0 and 1. Thus $\mathbf{q}_d$ is a filtered version of $\mathbf{q}$, and removes high frequency noise at the cost of a lag behind the current joint state. The filter parameter, ϵ, is chosen high enough such that speed of typical human arm movements is well within the bandwidth of the filter. If we choose $\epsilon = 1$, we achieve effectively the setting $\mathbf{q}_d = \mathbf{q}$. We indeed set the desired velocity $\dot{\mathbf{q}}_d = \dot{\mathbf{q}}$, but we will maintain a small amount of damping in the feed-forward command $\mathbf{u}_{\text{ff}}$.

b) Gravity, Coriolis, and Inertia Compensation: We assume the exoskeleton is governed by the well-known ridged-body dynamics model of a manipulator robot arm given by the equation

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) = \mathbf{u}, \tag{3}$$

where $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{7\times7}$ is the mass or inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^7$ denotes centrifugal and Coriolis forces, and $\mathbf{G}(\mathbf{q}) \in \mathbb{R}^7$ denotes the gravity force [16],[17]. Theoretically, we can completely cancel the dynamics of the exoskeleton robot with the feed-forward control law,

$$\mathbf{u}_{\text{ff}} = \tilde{\mathbf{M}}(\mathbf{q})(\ddot{\mathbf{q}} - \mathbf{K}'_D\dot{\mathbf{q}}) + \tilde{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}}) + \tilde{\mathbf{G}}(\mathbf{q}), \tag{4}$$

where $\tilde{\mathbf{M}}, \tilde{\mathbf{C}}, \tilde{\mathbf{G}}$ are our estimates of the robot's inertia, Coriolis, and gravity matrices, and $\mathbf{K}'_D$ is a diagonal matrix of damping gains. Assuming perfect parameter estimation, the robot's dynamics will be eliminating leaving only the damping term $\tilde{\mathbf{M}}(\mathbf{q})\mathbf{K}'_D\dot{\mathbf{q}}$ which is required to maintain stability. The damping gains can be individually tuned for each joint to apply the appropriate impedance characteristic

at each joint. Small gains can be used to keep the user's motion as unconstrained as possible. Premultiplying the damping term by the inertia matrix has the consequence of amplifying any modeling inaccuracies of the inertia matrix. Therefore, we move the damping term outside of the inertia matrix product, which empirically we find provides a greater margin of stability. Combining this control law with Equation (1) results in the equation

$$\mathbf{u} = \tilde{\mathbf{M}}(\mathbf{q})\ddot{\mathbf{q}} + \tilde{\mathbf{C}}(\mathbf{q}, \dot{\mathbf{q}}) + \tilde{\mathbf{G}}(\mathbf{q}) - \mathbf{K}'_D\dot{\mathbf{q}} + \mathbf{K}_P(\mathbf{q}_d - \mathbf{q}), \quad (5)$$

with $\mathbf{q}_d$ defined by Equation (2).

3) Parameter Identification: The estimates of our model parameters, $\tilde{\mathbf{M}}, \tilde{\mathbf{C}}, \tilde{\mathbf{G}}$, are determined with system identification techniques [15]. We recorded an hour of robot trajectories in response to sufficiently exciting desired trajectories of pseudorandom motor commands (including sine waves of various frequencies at the joint-level and discrete endpoint movements at various speeds). This data was used to regress the rigid body dynamics parameters acting on each DOF. However, we noticed that due to unmodeled nonlinearities of the robot, we obtained partially physically inconsistent rigid body parameters, e.g., non positive definite link inertia matrices and negative viscous friction coefficients. We improved this estimation with a novel nonlinear rigid body dynamics identification algorithm [13] that guaranteed a physically correct rigid body model. This model significantly increased the stability of our control system.

B. Force Fields

Since humans are so adept at motor tasks, particularly in skill acquisition and learning, it is likely that the brain includes a representation of the inverse dynamics of the limbs [4], [12]. However, this inverse dynamics model, or internal model, needs to be adaptive since in day-to-day activity humans interact with various loads and limb dynamics change during development. By studying how subjects' movements respond and adapt to novel external forces (applied by a robot manipulandum for example), scientists can draw conclusions on how planning, learning, and execution of movements occurs in humans, and how/if an internal model of dynamics is formed. Recently force field experiments have become a popular paradigm in which to explore such issues. A force field is created by the robot exerting external forces on the human subject, usually as a function of endpoint or joint positions, velocities, or accelerations. For example, [4] uses a viscous force field in the form of

$$\mathbf{f} = \mathbf{B}\dot{\mathbf{x}}, \quad (6)$$

where $\dot{\mathbf{x}}$ is the velocity of the hand, $\mathbf{f}$ is the vector of the external force to be applied at the hand, and $\mathbf{B}$ is a constant matrix which determines the strength and direction of the force. When executing point-to-point reaching movements in the absence of such fields, endpoint trajectories are typically straight with a smooth velocity profile [18]. However, upon first exposure to a novel dynamic environment (e.g. a force field), trajectories become skewed in response to the field. After several subsequent trials in the force field, endpoint trajectories actually return to their original paths. If the field is suddenly removed after learning, the trajectories once again become skewed, in the opposite direction of the initial pre-adaptation field trajectories, and are called *aftereffects* [9] of the force field adaptation. Evidence of aftereffects suggest that the brain does build an internal model of the new dynamic model of the force field, which it uses to predict and compensate for the environment such that its endpoint plan is realized. Currently these types of force field experiments have been limited to movements in a plane (e.g. a table top) using only two joints of the arm (e.g. the elbow and shoulder). Thus these experimental setups are unable to explore how redundancy plays a role in motor execution. Our experimental platform with the exoskeleton robot allows the subject to make unconstrained 3-D movements with all seven major DOFs of the arm. By applying perturbing forces at any or all of these seven joints, we are able to explore how subjects cope with intrinsic force fields and resolve redundancy during reaching tasks.

C. Experimental Procedure

For this experiment we asked three healthy right-handed male subjects to make a point-to-point reaching movement to a visual target. In each trial, the right hand begins above the shoulder (as if starting to throw a ball) and finishes with the arm slightly extended, in front of the torso (as if shaking someone's hand). The movement was chosen as such to maximize the spatial extent of the trajectory within the limited workspace of the robot. At the start of each trial, the robot servos the subject's arm to the starting location, and each trial begins with the same joint configuration. The subject is not instructed on how to execute the movement, but told that the reaching task should only involve an arm movement (not torso, hips, etc.) This way the subject's shoulder position stays approximately at the intersection of the exoskeleton's three shoulder rotation axes, and the subjects' and exoskeleton's joint angles will remain in cohesion with minimal discrepancy. The subject is also instructed to complete each trial within half a second. The force field, when applied, adds a shoulder velocity dependent torque to the elbow joint. Specifically, with the field on, the control law at the elbow becomes

$$\hat{u}_{\text{EFE}} = u_{\text{EFE}} + k_{\text{field}}\left(\dot{q}_{\text{SFE}} + \dot{q}_{\text{SAA}}\right), \quad (7)$$

where u_{EFE} is the torque command to the elbow joint generated by the control law of Equation (5), $\dot{q}_{\text{SFE}}, \dot{q}_{\text{SAA}}$ are the shoulder-flexion-extension and shoulder-adduction-abduction velocities, and k_{field} is a gain that determines the strength of the force field. Each subject completes 8 blocks of 25 trials each. The first 2 blocks are with the force field turned off (called the "null-field" condition, using only the control law of Equation (5)). By the second block the subject has sufficiently adapted to the manipulandum and the null-field condition. The third block consists of 22

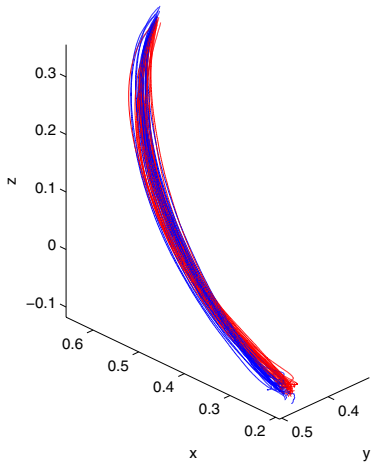

Fig. 2. Three dimensional trajectories of the right hand of Subject A. The 25 trials of block 2 (blue) are superimposed onto the 25 trials of block 7 (red). Hand trajectories return to the original null-field trajectories after adapting to the force field. Units are in meters.

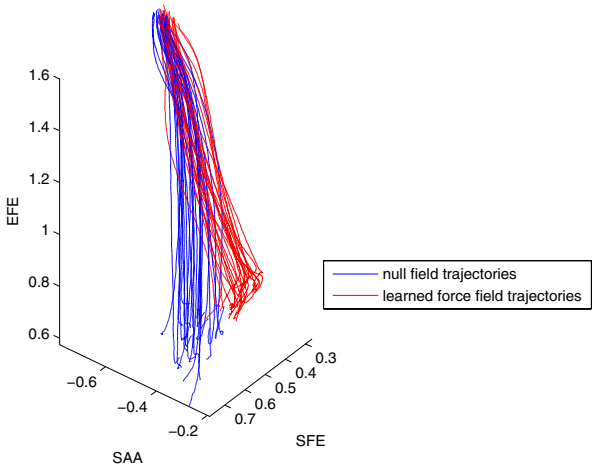

Fig. 3. Three dimensional trajectories of three joints of Subject A: shoulder-flexion-extension (SFE), shoulder-adduction-abduction (SAA), and elbow-flexion extension (EFE). The 25 block 7 joint trajectories (red) are skewed to the right of the 25 block 2 trajectories (blue). Although the hand has returned to its original trajectory after learning, the joint trajectories have been altered. Units are in radians.

trials of the null-field condition, with three randomly placed catch-trials with the force field on. Blocks 4-7 are executed with the field on, giving the subject opportunity to adapt to the field by block 7. Finally, the 8th block has 22 trials of the field on, with 3 randomly placed catch-trials with the field off.

III. Results

Figure 2 shows the 25 hand trajectories during subject A's block 2 (i.e. null-field trajectories) superimposed onto the 25 hand trajectories during block 7 (i.e. force field trajectories, after learning). Figure 3 shows the joint-space trajectories of SFE, SAA, and EFE for the same subject during the same blocks, with the force field trajectories skewed to the right in the figure. Figures 4 and 7 show the same subject's average hand and joint-space trajectories during block 2 and block 7. Additionally shown are the subject's initial trajectories when first exposed to the field (the average of the 3 catch trials of block 3) which are skewed to the right in the figure, and aftereffects (average of the 3 catch trials of block 8) which are skewed in the opposite direction. Average trajectories are found by resampling each trajectory (by applying an anti-aliasing FIR filter) to a fixed array length, and then computing the mean of each index across all trials. Figures 5, 6, 8, and 9 show endpoint and joint results for subjects B and C.

As expected, Figures 4-6 demonstrate that upon initial exposure to the force field, the subjects' hand trajectories are altered. As also demonstrated in other force field experiments, after sufficient adaptation hand trajectories return to the original null-field trajectory. Additionally, aftereffects appear when the field is suddenly removed, indicating that an internal model of the force field dynamics has been learned. Interestingly, however, Figures 7-9 show that the joint space trajectories do not return to the original null-field trajectories, and in fact converge onto an altered trajectory.

These results suggest that the motor system may be planning reaching movements in extrinsic (task space) coordinates and does not create fixed positional trajectories for the joints to follow. Rather, the motor system is able to exploit the redundancy available in the human arm (perhaps to reduce the effects of the force field) while still executing its endpoint plan.

IV. Conclusion

In this paper we demonstrated how a complex exoskeleton robot can be employed for behavioral studies of motor control. Despite that this hydraulic robot cannot compete with the quality of impedance control of very small scale haptic devices, a model-based controller with carefully tuned parameters accomplished a surprising quality of gravity, Coriolis, and inertia compensation, such that the robot did not alter the movement of a human inserted into the exoskeleton structure too much. We demonstrated the usefulness of this new experimental platform in a behavioral experiment, where human subjects were exposed to a novel dynamic environment in the form of joint space force fields. The experiment revealed that movement planning of reaching movements in humans does not employ desired joint space trajectories, but may rather be similar to inverse kinematics controllers with some form of null space optimization.

In addition to human arm movement studies, the exoskeleton has the potential for novel approaches to human

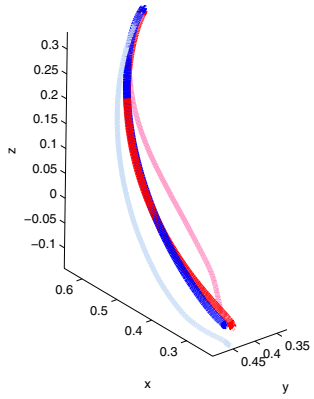

Fig. 4. Subject A's average hand trajectory of block 2 (blue) and block 7 (red). Additionally shown are the initial force field exposure trajectory (average of 3 catch trials in block 3) which is skewed to the right, and the average aftereffect trajectory (average of 3 catch trials of block 8) skewed to the left. Units are in meters.

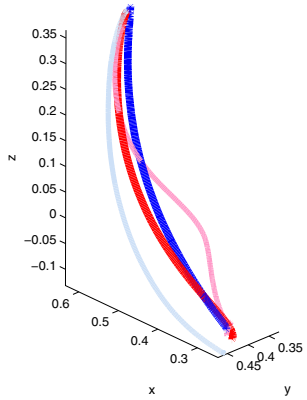

Fig. 5. Subject B's average hand trajectory of block 2 (blue) superimposed onto block 7 (red), with initial force field exposure trajectory to the right, and the average aftereffect trajectory to the left. Units are in meters.

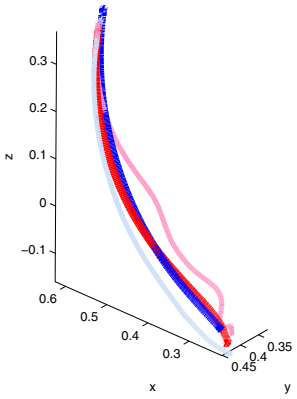

Fig. 6. Subject C's average hand trajectory of block 2 (blue) superimposed onto block 7 (red), with initial force field exposure trajectory to the right, and the average aftereffect trajectory to the left. Units are in meters.

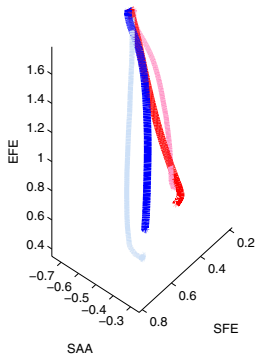

Fig. 7. Subject A's average joint trajectory of block 2 (blue) and block 7 (red, skewed to the right), with initial force field exposure trajectory (rightmost line), and the average aftereffect trajectory (leftmost line). Units are in radians.

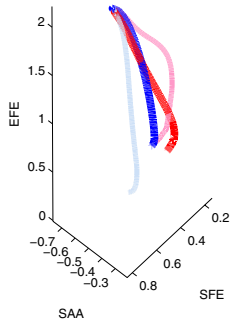

Fig. 8. Subject B's average joint trajectory of block 2 (blue) and block 7 (red, skewed to the right), with initial force field exposure trajectory (rightmost line), and the average aftereffect trajectory (leftmost line). Units are in radians.

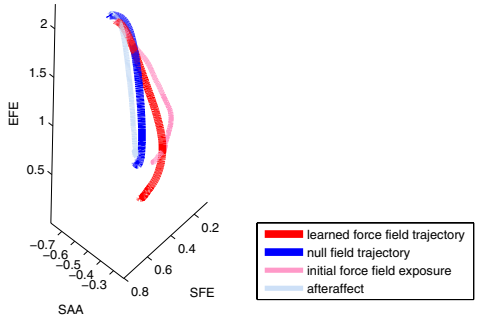

Fig. 9. Subject C's average joint trajectory of block 2 (blue) and block 7 (red, skewed to the right), with initial force field exposure trajectory (rightmost line), and the average aftereffect trajectory (leftmost line). Units are in radians.

arm rehabilitation. While wielding the apparatus, the user is able execute a wide range of discrete and periodic movements in full 3-D space, without being burdened by weight or inertia of the device. For rehabilitation procedures, the arm can be positioned in specific joint configurations, and different impedance properties can be applied to each joint. For example, wrist patients would be able to perform exercises involving entire arm movements, not just isolated wrist movements. Specific impedance characteristics or torque perturbations can be applied at the wrist only, while the remaining joints are as unconstrained as possible. Further applications in human movement science and arm rehabilitation will be explored in the future.

REFERENCES

[1] F. Mussa-Ivaldi, N. Hogan, E. Bizzi, "Neural, mechanical, and geometric factors subserving arm posture in humans," J. Neuroscience, vol. 5, pp. 2732-2743, 1985.

[2] H. Krebs, M. Aisen, B. Volpe, N. Hogan, "Robot-aided neurorehabilitation: Initial application to a stroke rehabiliatation," Proceedings of the MRCAS'95 - Second International Symposium on Medical Robotics and Computer Assisted Surgery, pp. CA260-CA262, Nov. 1995.

[3] S. Blakemore, D. Wolpert, C. Frith, "Why can't you tickle yourself?" Neuroreport, vol. 11, pp. R11-16, 2000.

[4] R. Shadmehr and F. Mussa-Ivaldi, "Adaptive Representation of Dynamics During Learning of a Motor Task," J. Neuroscience, vol. 14(5), pp. 3208-3224, May 1994.

[5] F. Gandolfo, F. Mussa-Ivaldi, and E. Bizzi, "Motor Learning by Field Approximation," Proc. Natl. Acad. Sci. USA, vol. 93, pp. 3843-3846, 1996.

[6] M. Conditt, F. Gandolfo, and F. Mussa-Ivaldi, "The Motor System Does Not Learn the Dynamics of the Arm by Rote Memorization of Past Experience," The American Physiological Society, pp. 554-559, 1997.

[7] H. Gomi and M. Kawato, "Human arm stiffness and equilibrium-point trajectory during multi-joint movement," Biol. Cybern., vol. 76, pp. 163-171, 1997.

[8] R. Shadmehr and F. Mussa-Ivaldi, "Computational Elements of the Adaptive Controller of the Human Arm," Advances in Neural Information Processing Systems, vol. 6, pp. 1077-1084, 1994.

[9] M. Jeannerod, The Neural and Behavioural Organization of Goal-Directed Movements, Oxford: Clarendon, 1988.

[10] S. Scott, "Apparatus for measuring and perturbing shoulder and elbow joint positions and torques during reaching," J. Neuroscience Methods, vol. 89, pp. 119-127, 1999.

[11] K. Singh and S. Scott, "A motor learning strategy reflects neural circuitry for limb control," Nature Neuroscience, vol.6, no. 4, pp. 399-403, 2003.

[12] M. Kawato, "Internal models for motor control and trajectory planning," Current Opinion in Neurobiology, vol. 9, pp. 718-727, 1999.

[13] S. Schaal, R. Cory, J. Nakanishi, M. Mistry, J. Peters, "Physically Consistent Nonlinear Parameter Identification of Rigid Body Dynamics," (in preparation).

[14] G. Burdea, Force and Touch Feedback for Virtual Reality, New York: John Wiley & Sons, Inc., 1996.

[15] C. An., C. Atkeson, J. Hollerbach, Model-Based Control of a Robot Manipulator, Cambridge, MA: The MIT Press, 1988.

[16] T. Yoshikawa, Foundations of Robotics: Analysis and Control, Cambridge, MA: MIT Press, 1990.

[17] L. Sciavicco, B. Siciliano, Modeling and Control of Robot Manipulators, New York: The McGraw-Hill Companies, Inc., 1996.

[18] P. Morasso, "Spatial control of arm movements," Experimental Brain Research, vol. 42, pp. 223-227, 1981.