Robot-Assisted Adaptive Training: Custom Force Fields for Teaching Movement Patterns

James L. Patton*, *Member, IEEE* and Ferdinando A. Mussa-Ivaldi

***Abstract*—Based on recent studies of neuro-adaptive control, we tested a new iterative algorithm to generate custom training forces to "trick" subjects into altering their target-directed reaching movements to a prechosen movement as an after-effect of adaptation. The prechosen movement goal, a sinusoidal-shaped path from start to end point, was never explicitly conveyed to the subject. We hypothesized that the adaptation would cause an alteration in the feedforward command that would result in the prechosen movement. Our results showed that when forces were suddenly removed after a training period of 330 movements, trajectories were significantly shifted toward the prechosen movement. However, de-adaptation occurred (i.e., the after-effect "washed out") in the 50–75 movements that followed the removal of the training forces. A second experiment suppressed vision of hand location and found a detectable reduction in the washout of after-effects, suggesting that visual feedback of error critically influences learning. A final experiment demonstrated that after-effects were also present in the neighborhood of training—44% of original directional shift was seen in adjacent, unpracticed movement directions to targets that were 60° different from the targets used for training. These results demonstrate the potential for these methods for teaching motor skills and for neuro-rehabilitation of brain-injured patients. This is a form of "implicit learning," because unlike explicit training methods, subjects learn movements with minimal instructions, no knowledge of, and little attention to the trajectory.**

***Index Terms*—Adaptation, control, force fields, haptics, human, human-machine interface, motor learning, robotic neurorehabilitation, robots, teaching.**

I. Introduction

ROBOT-ASSISTED training has great potential for refining motor skills and for rehabilitation. Robotic training can be highly accurate, can be sustained for very long periods of time, can measure progress automatically, and can produce a wide range of forces or motions. An important question is the *method* with which forces or motions are to be applied so as to best facilitate the goals of producing new movements. Being such a novel field of study, the most effective training algorithm has yet to be determined. Our motivational goal is to develop a robot-controlled training procedure that could facilitate the recovery of multi-joint arm coordination lost to brain injury, such as stroke or other neuromotor disabilities. However, the benefits of such procedures could be extended to the training of motor skills in the workplace and in sport. This paper investigates the initial problem of making healthy subjects move in a prechosen way, without them knowing.

This study explored one possible method in the general case of robotic teaching—one that capitalizes on the natural capacity of the nervous system to adapt to altered mechanical conditions. Rather than correcting the erroneous movements of disabled subjects, this initial study investigates the robot's ability to modify healthy subjects' hand movements to a prechosen trajectory, exploiting recent knowledge about neuromotor adaptation.

Recent studies have demonstrated that the effects of neuromotor adaptation are quite dramatic. When people are repeatedly exposed to a robot-generated force field that systematically disturbs arm motion, they are able to recover their original kinematic patterns [1]. Subjects do this by canceling the disturbances with a preplanned pattern of forces. This feedforward control is revealed by characteristic *after-effects*: when the disturbing force field is unexpectedly removed, subjects make erroneous movements in directions opposite to the perturbing forces. Adaptation and its after-effects have been demonstrated for a number of different types of disturbances, ranging from simple position-, velocity-, and acceleration-dependent force fields [1]–[5] to Coriolis forces arising in a rotating room [6] to skew-symmetric "curl" fields that produce forces in a direction perpendicular to the velocity [4]. Similar results have also been seen with visual disturbances [7]–[10]. More recent studies suggest that subjects learn an *internal model* of the perturbing force field rather than simply learning an appropriate temporal sequence of muscle activations [3], [4].

While conventional skill learning, such as learning to play a musical instrument, requires more conscious attention in order to achieve a goal, neuromotor adaptation has been argued to be closely related to *procedural learning* and a form of *implicit learning* [11]. Hence these learning mechanisms may offer a beneficial alternative to conventonal rehabilitation. Implicit learning takes place without awareness of what has been learned [12], and often does not require complete conscious attention. One example is procedural motor learning of a motor sequence that is embedded in a seemingly random set of movements [13]–[15]. Another example is sensory-motor adaptation observed in force field paradigms [11]. If one could appropriately design a set of training forces for a given subject, a new strategy for robotic training emerges.

One important feature of neuromotor adaptation is that *after-effects are predictable* via simple models. Subjects anticipate

Manuscript received March 4, 2002; revised July 25, 2003. This work was supported in part by the National Institutes of Health (NIH) under Grant 5 T32 HD07418, Grant 5 RO1 NS 35673, and Grant 1 F32 HD08658-01. *Asterisk indicates corresponding author.*

*J. L. Patton is with the Sensory Motor Performance Program, Rehabilitation Institute of Chicago, Physical Medicine and Rehabilitation, Mechanical & Biomedical Engineering Northwestern University, 345 East Superior St., Room 1406, Chicago, IL 60611 USA (e-mail: j-patton@northwestern.edu).

F. A. Mussa-Ivaldi is with the Sensory Motor Performance Program, Rehabilitation Institute of Chicago, Physiology, Physical Medicine and Rehabilitation, Institute for Neuroscience, Mechanical & Biomedical Engineering, Northwestern University, Chicago, IL 60611 (e-mail: sandro@northwestern.edu).

Digital Object Identifier 10.1109/TBME.2003.821035

the forces with an adjustment in the feedforward joint torques [1]. The adjustment cancels the forces of the field, resulting in a normal-looking movement when the force field is on, and an abnormal movement that reflects the adjustment when the force field is removed. Accordingly, we hypothesize that one can reverse-engineer this process and design force fields that will ultimately result in prechosen, "desired" after-effects.

In fact, a great number of potential approaches may solve this force field design problem. Initially, we addressed this problem using a novel technique that uses modeling to predict the behavior of the biomechanics and neural control of the movement based on systems identification of a parametric model [16]. The current investigation explores a simpler and more elegant approach. Specifically, we record the forces necessary to shift the subjects' movements to a "desired" trajectory. Then we train them by applying forces of the same amplitude but opposite direction with the intent of obtaining after effects that resemble the desired movements. In Sections II–IV, we present a general framework and then test it using a two-joint manipulandum. This *implicit teaching* technique is novel because only the experimenter knows the desired movement; therefore, it may become a valuable tool to be added to the repertoire of robotic techniques for training and therapy.

II. Methods

A. Robotic Training Algorithm

The goal of adaptive training is to determine the robot-applied forces that lead to the execution of the desired movement $\mathbf{x_D}(t)$ as an after-effect of adaptation. Here, we test the idea using the simplest possible method of adaptive training. Our approach hinges on some assumptions. First, we assume that during the early part of a short, rapid movement, the joint torques produced by a subject, $\boldsymbol{\tau}(t)$, are generated as a function of time by a feedforward controller [3], [4], [17], [18]. Second, we assumed in this initial study that peak movement speeds were roughly constant. We controlled for speed by providing a tone after the completion of each movement, indicating that the movement was either too fast (high pitch), too slow (low pitch), or just right (medium pitch, between 0.3 and 0.4 m/s). Third, we assume that following training, the adapted torques are an alteration in the feedforward plan that cancels out externally applied forces from the robot, $\boldsymbol{\tau} + \boldsymbol{\Delta\tau}$. When the endpoint forces were suddenly removed, the adaptation $\boldsymbol{\Delta\tau}$ causes the arm to move along the new after-effect trajectory. The adapted torques $\boldsymbol{\Delta\tau}$ must be equivalent to the forces required to shift the trajectory to the after-effect trajectory. Therefore, an intermediate step is to determine the forces required to shift the subject's expected trajectory to the desired trajectory.

Because our robot, as well as other similar systems used in man-machine interactions, lacks the power and speed to servo-control the arm to the desired trajectory, a simple iterative machine-learning algorithm determined the feedforward forces $\mathbf{F_{Di}}(t)$ required to shift the unsuspecting subject's trajectory to the desired trajectory $\mathbf{x_{Di}}(t)$ during the first 200 ms of the movement. Here, the subscript $\mathbf{i}$ represents the three directions of motion. Forces were presented intermittently (one in every four trials, randomly presented) to avoid inducing an adaptive response in the subject. We chose to only perturb during the first 200 ms of movement because our assumptions are that we were solely influencing the feedforward component of movement—later perturbations would be confounded by corrective submovements caused earlier errors. We began each experiment with $\mathbf{F_{Di}}(t) = \mathbf{0}$ for all times, and made incremental adjustments from one movement to the next following each trial that presented the forces using the following online learning rule:

$$\mathbf{F}_{Di}(t) = \mathbf{F_{Di}}(t) + \mu[\mathbf{x_i}(t) - \mathbf{x_{Di}}(t)] \tag{1}$$

where $\mathbf{x}(t)$ is the subject's trajectory and $\mathbf{x_{Di}}(t)$ is the desired trajectory for a given direction of movement $\mathbf{i}$. The parameter μ was heuristically chosen to be 30 Nm^{-1}. Higher values of μ led to unstable learning where the force did not converge and lower values required an excessively long time for the system to learn the forces. For safety, we also saturated the force at $\pm 14\,\mathrm{N}$. Once these forces $\mathbf{F_{Di}}(t)$ were known, we applied their vector inverse repeatedly in a training session to induce an adaptation. Note that the forces were sufficient to cause changes in both direction and velocity, but since feedback was given to keep speed roughly constant, forces tended to be almost perpendicular to movement direction.

B. Experiments

Eight healthy right-handed adults with no history of orthopedic or neurological disorders volunteered to participate. Before beginning the experiments, each subject signed a consent form that conformed to federal and university guidelines. Seated subjects held a two-degree-of-freedom manipulandum (inset, Fig. 1) described elsewhere [3], [19]. Data were collected at 100 Hz. Subjects were seated so that the center of the range of targets was anterior to the shoulder, approximately in the center of their reachable workspace (Fig. 1). A computer monitor mounted in front of the eyes provided successive visual targets for discrete movements. Each movement was in one of three randomly chosen directions but bounded by a rectangular area 38 cm wide by 28 cm tall. For all experiments, the subject's own arm was peripherally visible. Each movement was 10 cm in one of three randomly chosen directions spaced 120° apart: anterior-right, anterior-left, and posterior-center (toward the chest). To avoid fatigue, subjects could choose to rest before initiating any movement, though they rarely did. All subjects performed a total of 873 movements, broken down into the following experimental phases. The number of movements in each direction was the same for each phase.

1) *Unperturbed familiarization:* Fifteen movements (approximately 25 s) to become familiar with the system.
2) *Unperturbed baseline:* Fifteen movements (approximately 25 s) to establish a baseline pattern.
3) *Machine-learning:*Two-hundred an dninety-eight movements (approximately 450 s), with random, intermittent perturbations presented one in four trials. The computer learns the forces required to push the subject over to the "desired" trajectory.
4) *Unperturbed baseline:*Eighteen movements (approximately 27 s) to determine if the baseline pattern changed.
5) *Learning:* Three-hundred and thirty movements (approximately 500 seconds) of constant exposure to the forces.

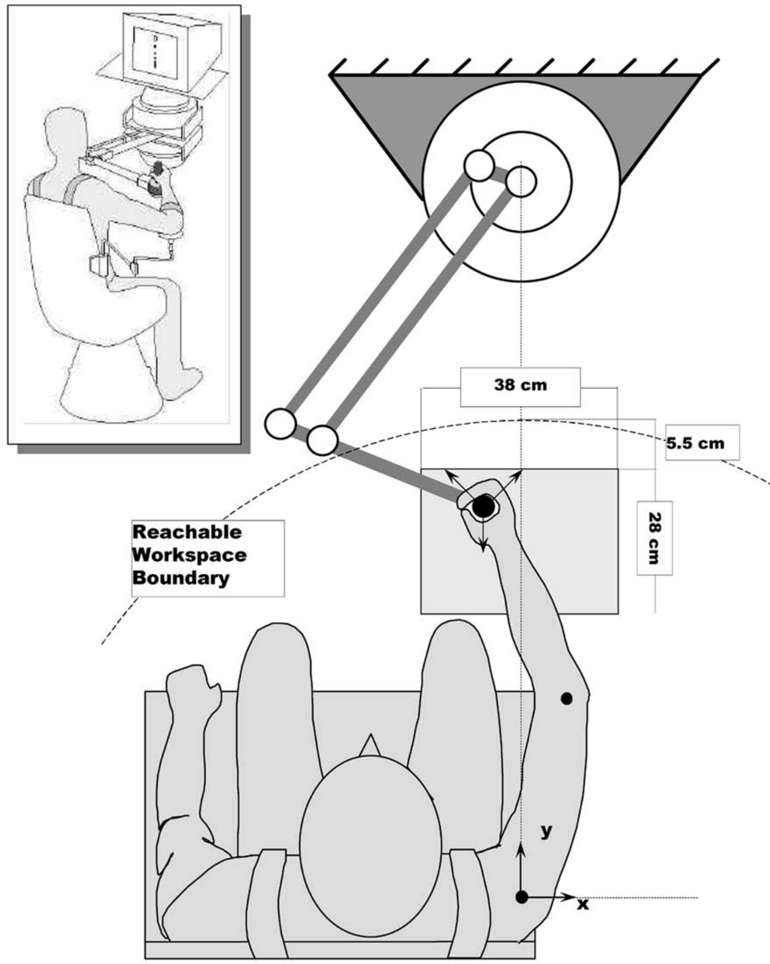

Fig. 1. Subject positioning at the experimental apparatus. Two brushed dc torque motors (PMI model JR24M4CH, Kolmorgen Motion Technologies, NY) control forces at a handle via a 4-bar linkage. Rotational digital encoders (model 25/045-NB17-TA-PPA-QAR1S, Teledyne-Gurley, Troy, NY) report absolute angular position, and a six-axis force/torque sensor (Assurance Technologies, Inc., TI F/T Gamma 30/10, Apex, NC) reports the interface kinetics. A PC acquires the signals and controls torque.

6) *After-effects:* One-hundred and twenty 120 movements (approximately 180 s), with random, intermittent removal of the force field for one in eight of the trials (*catch trials*) to determine the after-effects.
7) *Washout:* Seventy-five movements (approximately 110 s), all without forces.

The desired trajectory $\mathbf{x_D}(t)$ that we chose for all subjects was a curved, sinusoidal-shaped path from start to end point. The hand was supposed to trace this pathway with a symmetric bell-shaped velocity profile, with peak velocity midway [Fig. 3(B)]. We chose this desired trajectory because it was not considered a movement that is biomechanically or physiologically impossible to make and because it would involve forces that subjects could learn in a reasonable time. Note that our approach to the learning of this new trajectory is *implicit*—subjects were never given any knowledge of the desired trajectory, and did not receive any instructions other than to move to the target within the range of target speeds.

C. Analysis

To quantify how well we shifted subjects' movements toward the desired trajectory, we chose initial direction error, defined by measuring the error between the desired and actual initial vectors. Initial vectors were formed by connecting the start point to the point when the trajectory had moved 2.5 cm (25% of the distance to the target). We found similar results with measures used in similar studies, such as *figural error* [3] or the maximum perpendicular distance [20] and *infinity (Chebychev) norm* [19]. Initial direction error can be negative or positive. We defined positive error to correspond to a counter-clockwise rotation from the actual trajectory to the desired trajectory, and zero corresponds with the accomplishment of our goal of shifting subjects' trajectories to the desired trajectory.

All hypotheses were tested using an alpha level of 0.05. We tested for a reduction in error between baseline and after-effects, for a difference between the after effects trajectory and the desired trajectory, and for a tendency for the after effects to disappear in the washout phase.

III. Results

A. Adaptive Training

Figs. 2 and 3 display data from a typical subject, while Figs. 4 and 5 illustrate the results of the group. Because of the feedback

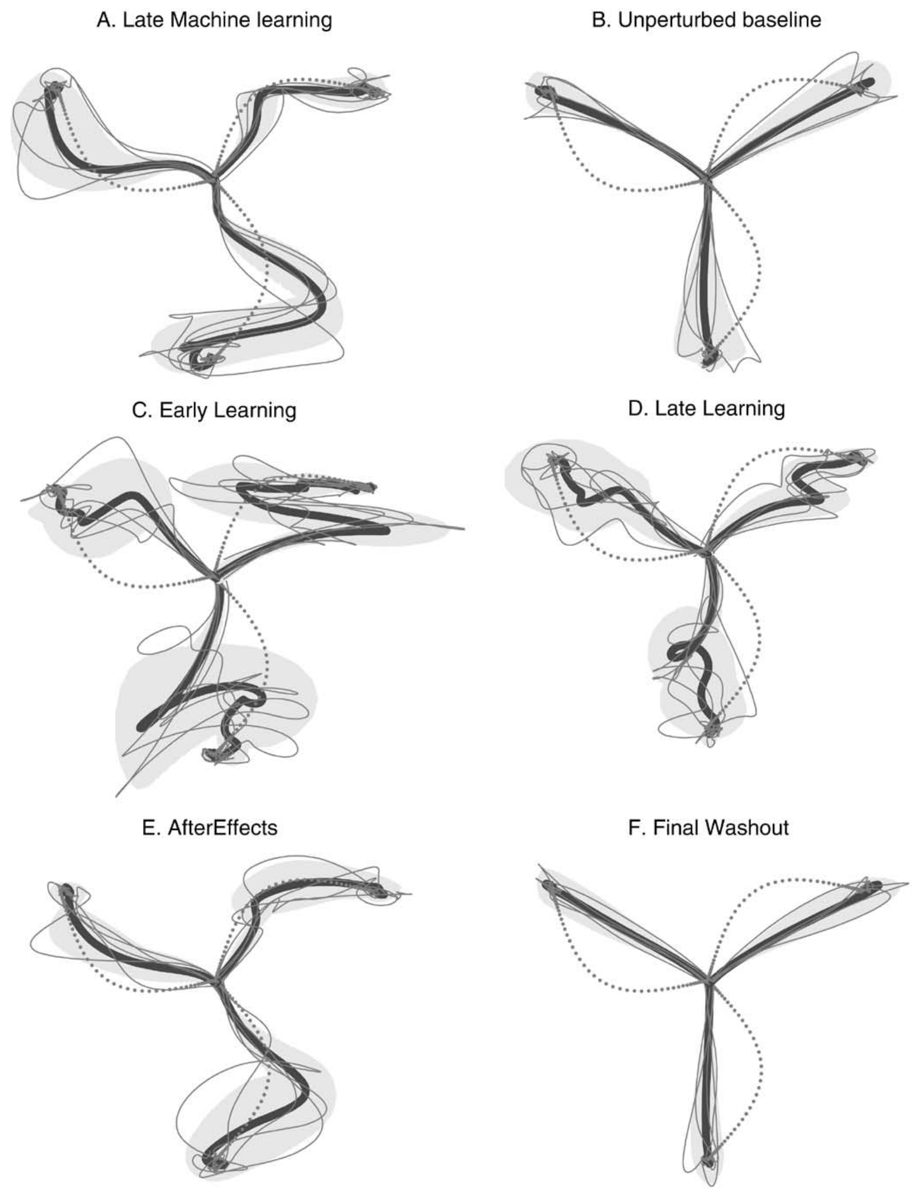

Fig. 2. Movement paths for a typical subject during successive phases of the experiment. Desired trajectories are the bold dotted lines, the average trajectories are represented as bold solid lines, individual trajectories are thin lines, and shaded areas indicate 95% confidence intervals. Online learning eventually results in a reasonable shift toward the desired trajectory (A). After unperturbed trials (B), the subject then repeatedly trains (C and D) on a force field that is the vector-opposite the forces applied in (A). The training forces are then turned off intermittently (catch trials) to test for after-effects (E). Finally, subjects move for 75 movements without forces, and the adaptation "washes out." Results from the final 15 movements of the washout phase are shown in (F).

of peak speed, this remained roughly constant across the entire experiment and did not differ from the target speed (0.404 m/s). As expected, the initial unperturbed movements were along approximately straight-line trajectories with bell-shaped-velocity profiles [Fig. 3(B)]. The unperturbed baseline trials that occurred before machine learning (not shown) did not differ significantly from the second set of unperturbed baseline trials that occurred after machine learning [Figs. 2(B) and 3(B)]. This indicated that the intermittent perturbations of the machine-learning phase did not cause any noticeable adaptation. Note that the bold dotted lines in Figs. 2(B) and 3(B) indicate the desired curved trajectory—our eventual goal. In the machine-learning phase, the robot controller gradually learned to produce forces that pushed the subjects toward the desired trajectory [Figs. 2(A) and 3(A)]. The forces resulting form the machine learning phase ranged from 0–19 N and were roughly similar across subjects (Fig. 4). On average across subjects, the interactive algorithm for machine learning converged to the desired initial direction (Fig. 5 bar 2), yet some subjects overshot and some undershot the goal.

In the learning phase, subjects were repeatedly exposed to the opposite of these forces while being asked to execute the same reaching movements. Initially, because of this perturbation their hand moved in the direction opposite to the desired path [Figs. 2(C), 3(C), and 4(C)]. All subjects adapted their control in the learning phase to recover a more normal pattern by the end of learning [Figs. 2(D) and 3(D)], although recovery was not complete [compare (D) to (B) in these Figs. 2 and 3; see also Fig. 5]. Note the changes in speed profiles from the simple bell-shaped velocity pattern in the baseline [Fig. 3(B)] to a bimodal

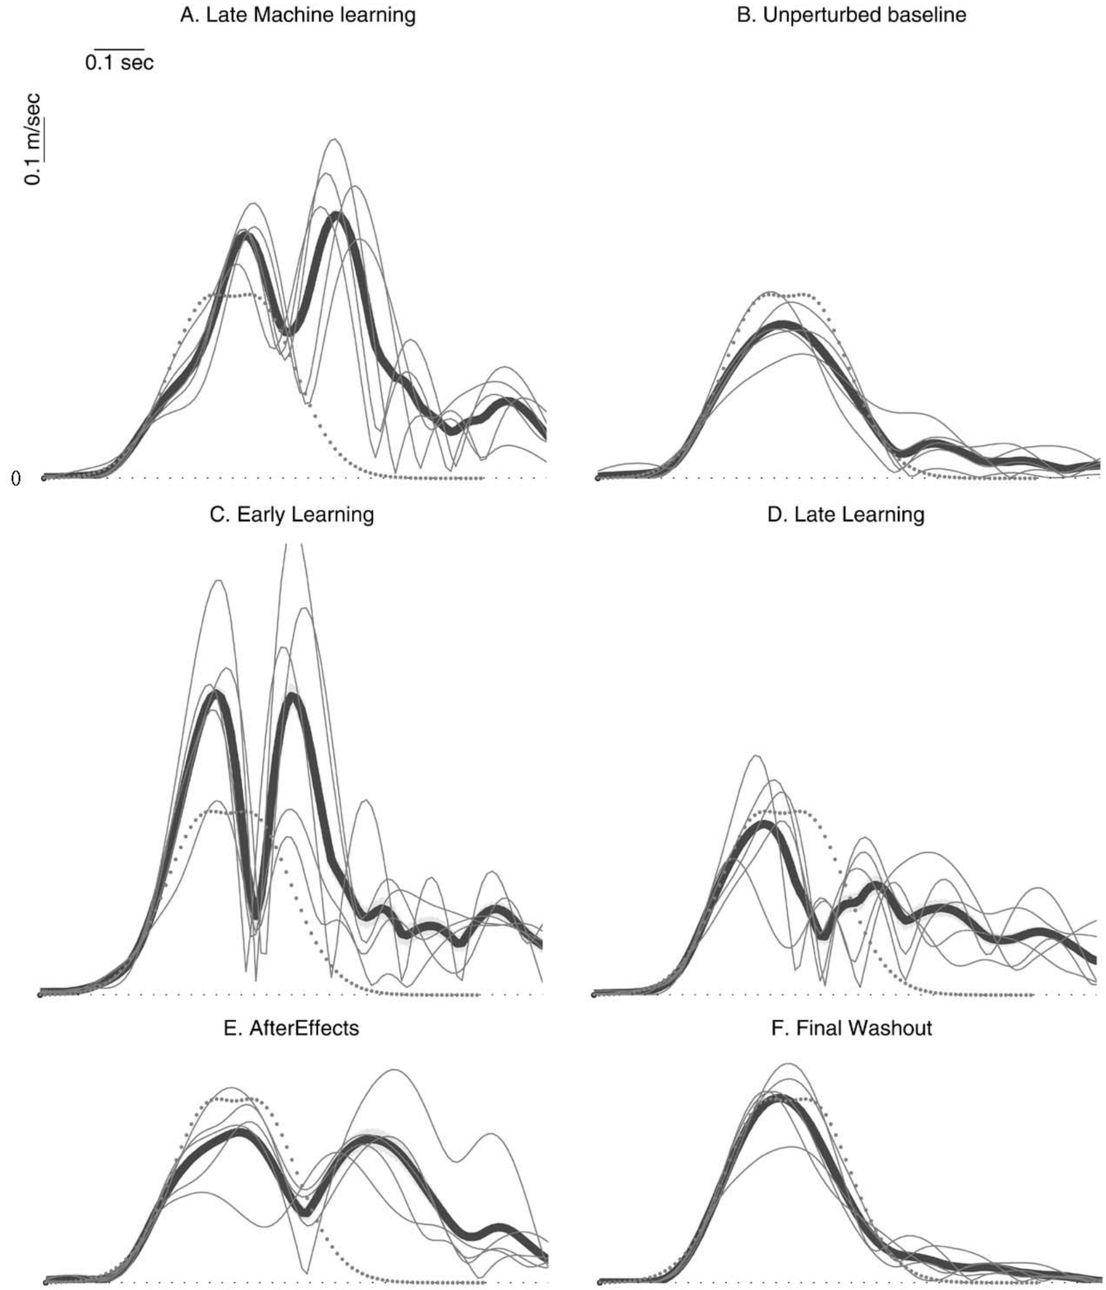

Fig. 3. Speed profiles for the same subject and experimental phases as the previous figure. Again, averages are represented as bold solid lines, individual trajectories are thin lines, and shaded areas indicate 95% confidence.

pattern when the robotic forces were on [Fig. 3(A) and (C)]. This pattern roughly corresponded to the period when the forces were on (the first 200 ms of movement) and off. By the end of the learning phase, the pattern tended to be more coalesced [Fig. 3(D)].

A critical question is whether or not we were able to cause subjects, as a group, to significantly shift movements toward the desired trajectory as an after-effect of adaptation. After-effects occurred when the robot forces were suddenly removed. The hand movements at the end of the machine-learning phase did not match exactly the desired trajectory and at the end of the learning phase, subjects did not completely recover the unperturbed kinematics. Nevertheless, subjects' after-effects trajectories were consistently shifted toward the desired trajectory [Figs. 2(E) and 3(E)]. Results for the entire group of subjects showed that this training framework was capable of causing subjects' trajectories to shift significantly $(p < 0.05)$ closer to the desired trajectory, as indicated by a reduction in error between baseline and after-effects trials (horizontal bracket at the top of Fig. 5). However, the shift was not complete. There was a significant residual difference between the desired trajectories and the after-effects, as indicated by a nonzero error in the after-effects. Moreover, the speed profiles of the after-effect trajectories never resembled the single, bell-shaped pulse of the desired trajectory [Fig. 3(E)]. The after-effects speed profiles suggest the presence of a corrective movement that began about at 300–350 ms (100–150 ms following the end of the training forces). This second speed pulse was often larger than the first. However, Fig. 3(E) suggests some amount of blending between the two components of the after-effect movement.

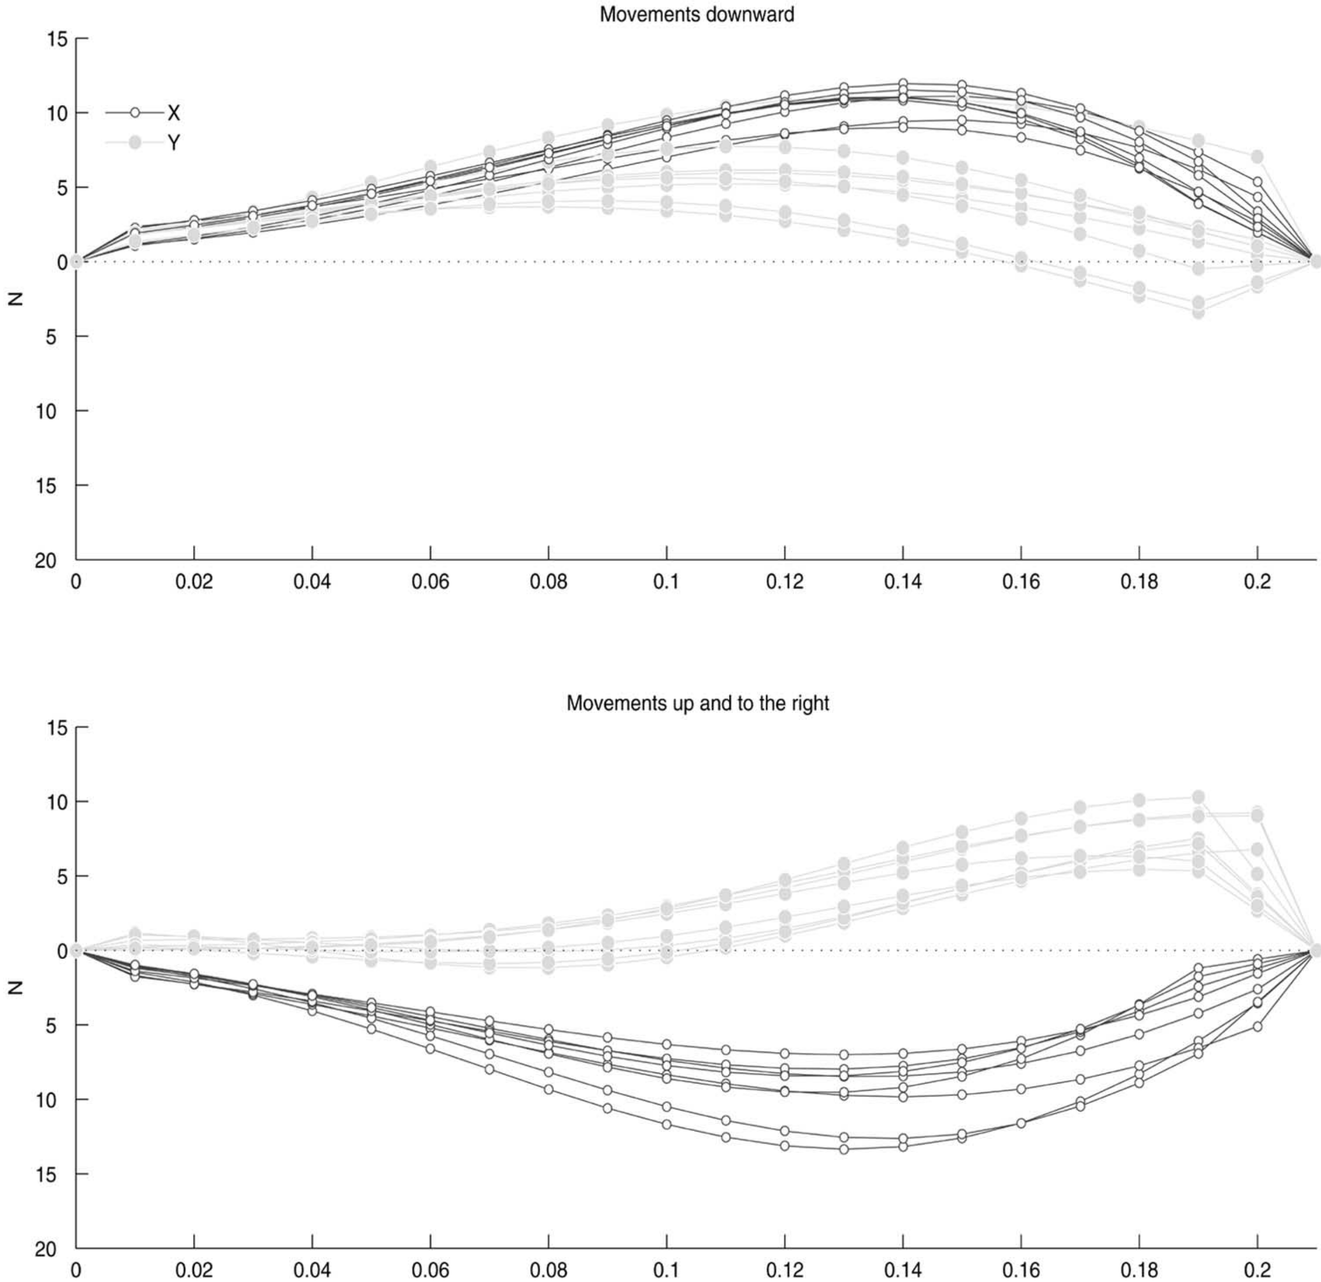

Fig. 4. Forces resulting form the machine learning algorithm for all subjects, which shifts trajectories toward the desired movement. Each figure represents a different direction of movement, and each line is a different subject's data.

B. Persistence of After-Effects

As subjects were not given any information about the desired curved trajectory, the after-effect of adaptation gradually "washed out," with the final movement after 75 trials [Fig. 3(F); Fig. 5, final three blocks of data] being very similar to the baseline. The washout trials were subdivided into the first 15 (*early washout*), middle 45 (*middle washout*), and final 15 (*final washout*) to illustrate this outcome. In summary, although the after-effect did resemble the desired trajectory, the method used to obtain it provides no guarantee of a lasting performance. This is not surprising, as the subjects were not instructed in any way to maintain the after-effect. Actually, they were likely to interpret the after-effect as an error with respect to the intended movement and to attempt correcting this error in subsequent movements. This led us to think that if visual feedback information was removed, the washout process may be prevented or slowed down.

C. Visual Feedback and After-Effect Washout

We performed an additional experiment (Experiment 2) with four additional subjects to determine if removal of visual feedback would eliminate or attenuate the washout. Subjects repeated the experiment except for the removal of visual feedback of hand position during the movement in all but the initial, unperturbed familiarization movements. Visual feedback was available only after the movement had terminated, at which point the subject could see the cursor and move it to the appropriate starting point for the next movement. The extent and rate

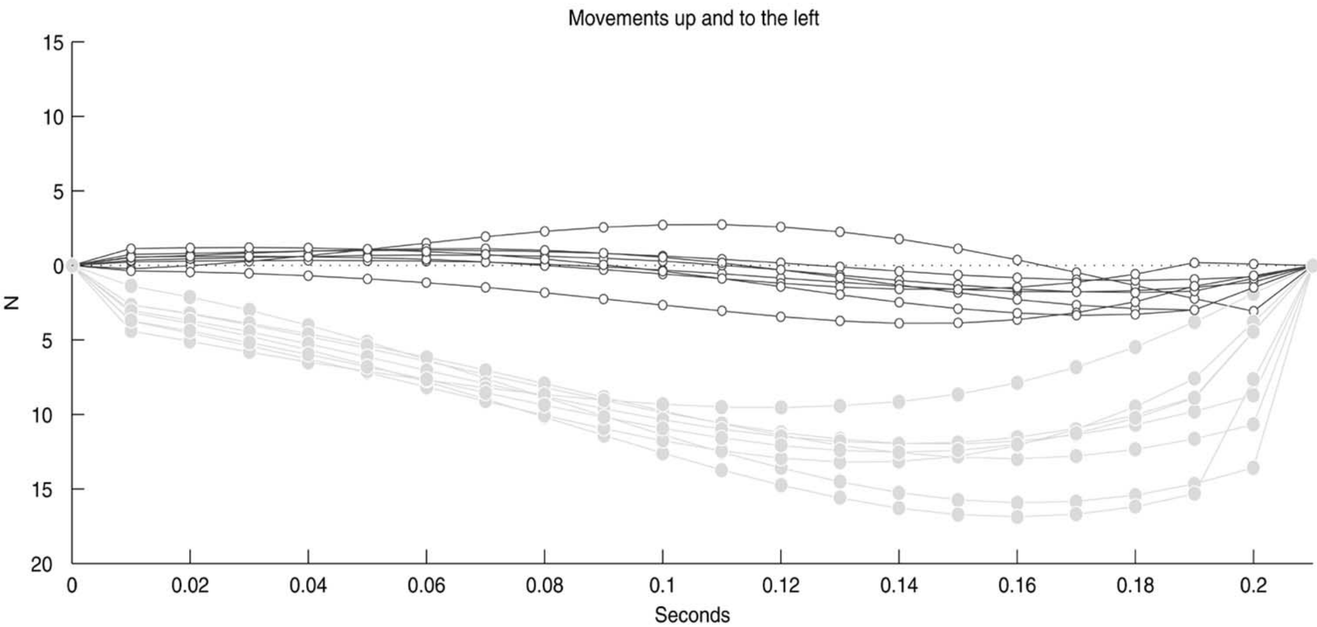

Fig. 4. (*Continued.*) Forces resulting form the machine learning algorithm for all subjects, which shifts trajectories toward the desired movement. Each figure represents a different direction of movement, and each line is a different subject's data.

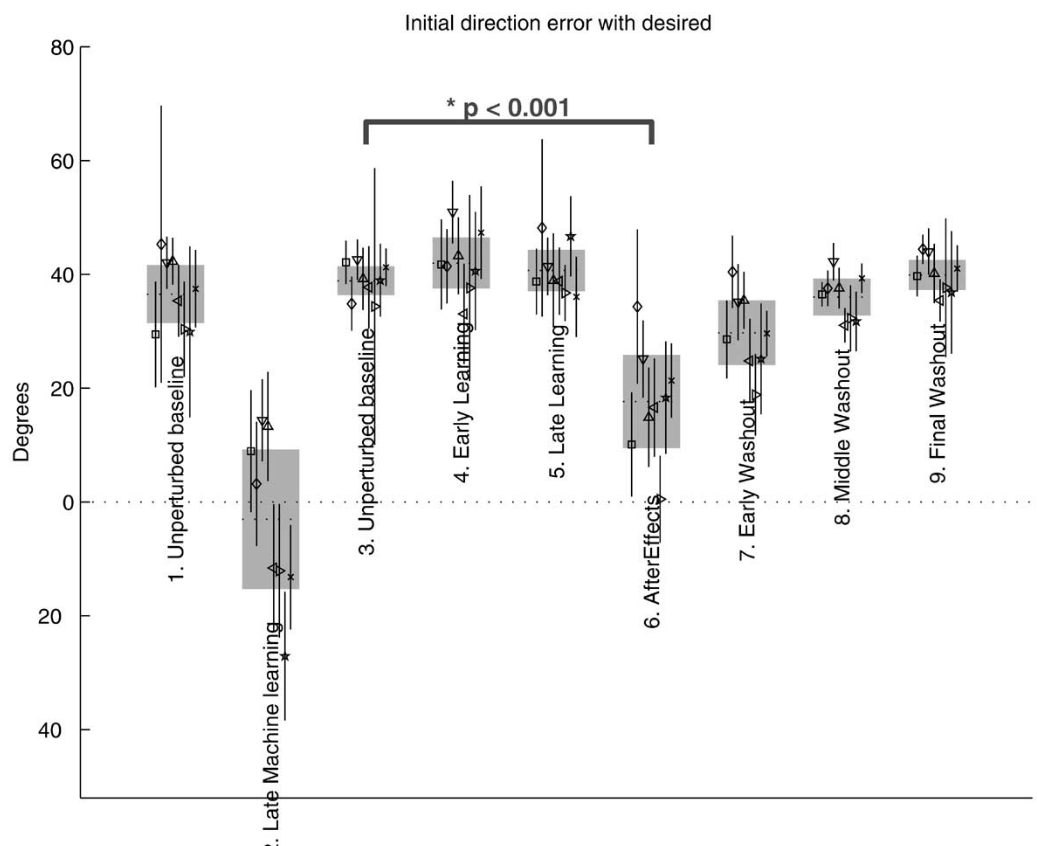

Fig. 5. Group results showing the initial direction error between the actual trajectory and the "desired" trajectory. Each block of data along the horizontal axis represents a successive phase of the experiment. Baseline is compared to after-effects to test our hypothesis that the method shifts trajectories toward the desired, resulting in a significant reduction in error. Shaded areas and dashed lines represent 95% confidence intervals and means for the group, while the symbols and bars show the individual subject data.

of the washout phase was fit to an exponential function to characterize the difference between Experiment 1 and 2

$$E = a - be^{(-t/c)} \tag{2}$$

Here, parameter a indicates an offset, while b and c represent the *amount* and the *time constant* of washout, respectively.

Our results indicated that washout was attenuated by removal of vision, although it was not completely removed. While sub-

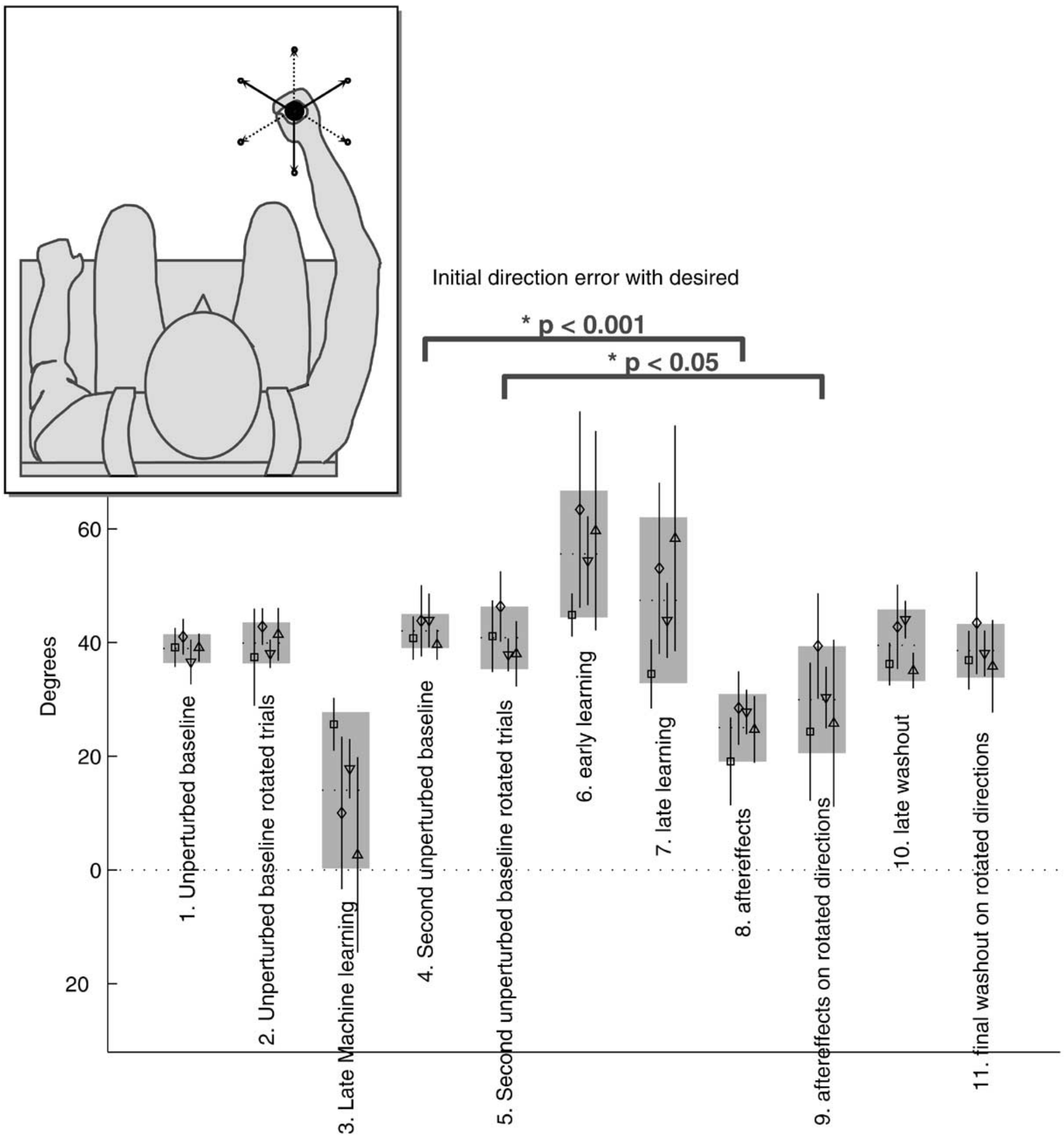

Fig. 6. Third experiment that illustrates the transfer of after-effects to adjacent, unpracticed directions (labeled rotated directions). Though these trajectories were not part of the training, there is still a significant shift toward the desired trajectory. The conventions used are the same as in Fig. 3. Inset: solid arrows indicate training directions, while the dashed arrows indicate the rotated directions.

jects appeared to learn nearly the same trajectory in both experiments, as indicated by a similar amount of washout, their after-effects decayed more slowly without vision than with vision, indicated by a significantly longer time-constant. The average time constants for decay rose from 20 movements for subjects with vision to 385 movements for subjects without vision (although the intersubject variability of the time constants also increased fifteenfold). Hence, removal of vision during the movement phase is enough to significantly decelerate the washout phenomenon seen in this study. Additionally, the no-vision data was more variable, also suggesting that visual feedback reduced the error and played a role in the control and learning of the movement.

We investigated a further question of whether adaptive training on these specific directions of movement generalized to adjacent movements that were not part of the training. We wished to test the hypothesis that the adaptive system is "broadly tuned" so that training in one set of directions can influence others [4], [21]. We evaluated four more subjects (Experiment 3) with a slightly altered protocol, which included extra trials in new directions during the baseline, after-effects and final washout sections. These trials were aimed at targets halfway between the original targets.

We found significant after-effects on these nonpracticed directions ($p < 0.05$), indicating that subjects' learning generalized to movements that were not practiced (Fig. 6). The after effects on these rotated trials were less dramatic, however—the mean reduction of error was 48%of that observed for the directions practiced. In summary, although we found significant after-effects that generalize to nearby locations these effects decay with their distance from the practiced directions, consistent with earlier adaptation results of Gandolfo *et al.* [4].

IV. Discussion

This study investigated an implicit approach to robot-assisted motor learning, which we call "adaptive training." Subjects practiced repeated reaching movements of the hand in the presence of perturbing forces. These forces were specifically designed so that when they were removed, a desired trajectory resulted as an after-effect of adaptation. Subjects showed a significant shift toward the desired trajectory, even though they were never given any instruction or knowledge of what it was. Essentially, this method "tricks" the nervous system into generating a new motor command. The generalization to movements in different directions is evidence that the subjects developed a broadly tuned internal model of the perturbing forces that influences movements in the neighborhood of the practiced directions. This study provides both verification and a practical application of what is known so far about force field adaptation and feedforward control of movement.

Although adaptive training may be considered a promising approach to teaching movements, there are possible shortcomings. First, the approach—at least in its current simple form—pays no attention to spinal reflex properties. Although we only evaluate the first 200 ms of the movement, this interval *is* long enough for spinal reflexes to play a role [22], [23]. If these contributions were more than negligible, the effectiveness of our method—which is based on adaptation of a feedforward command—would have been reduced. Movements in which feedforward control plays a smaller role (e.g., slow, large movements) may be less influenced by this method. For this reason, a more sophisticated approach may be to model the contributions and factor them into the estimates of the responses to the perturbing field [16]. However, our early attempts at such a modeling approach has proved less successful up to this point [24], mainly because the implementation is currently limited by the speed of the computers and/or by the computational efficiency of our current algorithms. As more sophisticated controllers and models are developed, the modeling approach may prove more effective.

One implicit assumption of this approach is that the forces that push the subject over to the desired trajectory $\mathbf{x_D}(t)$ can be learned correctly though they are learned along a different trajectory [near the subjects expected trajectory $\mathbf{x_E}(t)$]. Because of the nonlinear relationship between joint torques and endpoint forces, the torques applied along one trajectory are not necessarily appropriate for another [25]. However, we argue that if that the desired and expected trajectories are within a "domain of proper generalization," [16] the forces applied are a good enough approximation and can lead to desirable after effects. Recent studies have presented evidence that motor learning is broadly tuned so that training in one set of directions can influence others [1], [3], [21], [26], [27] and is even quantifiable via an adaptation model [20].

We make no claim that it is advantageous for motor learning to withhold knowledge of the desired performance from the subject. Instead, we simply point out that no explicit goal is needed to obtain the effect. It would seem plausible that if one were to inform subjects of the force field used for training and of the goals, one could get subjects to retain their after-effects. We attempted this with two additional subjects, who were explicitly asked to retain their after-effects. These subjects were able to retain their after-effects, but we could not conclude whether they were able to do so by retaining the adapted pattern they learned or by switching their motor intention to the new movement.

The suppression of vision of hand location (Experiment 2) demonstrated learning (and unlearning) is facilitated by the visual feedback of error. This result agrees with a previous result in our laboratory that demonstrates that the absence of kinematic error leads to the persistence of after-effects [19]. In that study, kinematic errors were prevented from occurring by imposing a simulated, mechanical "channel" on the movements. Hand forces revealed that recovery from adaptation was much slower compared with when kinematic aftereffects were allowed to take place. Instead, subjects persisted in generating unnecessary forces against the constraint. These results agree with motor learning models that use feedback of error as the primary teaching signal [28]–[31]. This also provides encouraging evidence that if subjects perceive their after-effects as an improvement, the beneficial results of this type of adaptive training will persist.

In order for these methods to be applicable for robotic training, the learning effect must also generalize to nearby movements that were not practiced. It is in some ways surprising that we found any generalization at all, considering that the three time-dependent force recordings were formed independently in the machine-learning phase. Nevertheless, our results showed that when subjects moved in directions that were 60° away from the directions practiced, the trials shifted an average of 44% of the initial direction error seen in training directions (Experiment 3). The same type of after-effect has been observed in velocity-dependent field training [20], [32], where after-effects were reduced $23\pm13\%$(mean and 95% confidence interval; cf. [20, Fig. 2]). Such variance or broad tuning of learning is useful in both machine learning and in physiological systems [21], [27]. One explanation is that multiple primitives observed in the nervous system may form the "motor vocabulary" that is used for a basis for movements. Such a phenomenon can be exploited, so that exhaustive training in all possible directions is not required to accomplish after-effects in all directions.

The results that 1 suppression of feedback error leads to the persistence of after-effects and 2) that the system is capable of generalizing to adjacent trajectories suggest that this method can be practically extended to the area of training. In particular, adaptive training could be used to teach difficult and specialized movements to healthy individuals or to facilitate the relearning process following brain injuries such as stroke. Such an approach needs only to induce incrementally good approximations of the desired movements. Indeed, the advantage of robot-assisted training is that procedures may be repeated and monitored for extended periods of time, thus bringing movements ever closer to a desired one. Studies using robotics for rehabilitation, assessment and training have had some success [33], [34]. The paradigm of adaptive training may add to the arsenal of possible strategies for rehabilitation and other forms of training. In fact, conflicting methods and techniques have been advocated for treatment of stroke and other forms of brain injury. Some studies suggest assisting or reducing errors during reaching movements may contribute positively to rehabilitation [35]. Other studies suggest resisting reaching movements [36]. Though these approaches follow opposing views, the efficacy of the various techniques has not been tested objectively with a tool such as a robot. It is not clear whether assistive, error-reducing forces are more appropriate than resistive or error-augmenting forces. The technique applied in this study is simply one form of resistive or error-augmenting training.

It is not known yet whether the washout observed in healthy subjects and in the conditions of this study, will be as strong for impaired patients. For these patients, the after-effects resulting from adaptive training would eventually correspond to a more normal-looking movement. Preliminary studies in our laboratory on stroke victims have revealed that after-effects may persist longer when the after-effects resemble normal movements [37], [38]. In fact, the after-effects may become permanent if they are perceived by the subject to be an improvement with respect to the initial behavior. However, not all patients may benefit from this type of procedure, because it requires that the neural areas presiding over motor adaptation be fully functional. Subjects who show poor abilities to adapt, such as cerebellar stroke patients, may have great difficulty [39]–[41].

The implications of implicit learning using a robot are remarkable: one can learn at a nearly subconscious level with minimal attention and less motivation than more explicit types of practice like pattern tracing. This idea is now being extended in preliminary form to poststroke [37]. In our preliminary work on stroke survivors, nearly all individuals are amenable and enthusiastic about this "video game." The implicit learning process takes place without awareness of what has been learned [12], and often does not require complete conscious attention. Recalling what has been learned may also be more automatic [42]. Implicit learning may involve different neural pathways [43], [44] that have earlier evolutionary origins than explicit learning areas [45]. Consequently, the robotic training processes in this study may provide an efficient teaching alternative by bypassing areas affected by attention deficits or brain damage. As rehabilitation training typically requires a balance of repetitive practice, strengthening, and expert guidance, we believe that the implicit approach presented here may inspire new forms of guidance in areas such as rehabilitation, robotic surgery, sport training, ergonomics, and manufacturing assembly training.

ACKNOWLEDGMENT

The authors would like to thank E. Davchev, C. Malecka, and P. Shah for the assistance in data collection. They would also like to thank M. E. Phillips-Stoykov for her clinical advice. Additional information related to this paper is also available at http://www.smpp.northwestern.edu/robotLab.

REFERENCES

[1] R. Shadmehr and F. A. Mussa-Ivaldi, "Adaptive representationof dynamics during learning of a motor task," *J. Neurosci.*, vol. 14, pp. 3208–3224, 1994.

[2] O. Bock, "Load compensation in human goal-directed armmovements," *Behav. Brain Res.*, vol. 41, pp. 167–177, 1990.

[3] M. A. Conditt, F. Gandolfo, and F. A. Mussa-Ivaldi, "Themotor system does not learn the dynamics of the arm by rotememorization of past experience," *J. Neurophysiol.*, vol. 78, pp. 554–560, 1997.

[4] F. Gandolfo, F. A. Mussa-Ivaldi, and E. Bizzi, "Motorlearning by field approximation," *Proc. Nat. Acad. Sci.*, vol. 93, pp. 3843–3846, 1996.

[5] T. Flash and F. Gurevitch, "Arm movement and stiffnessadaptation to external loads," in *13th IEEE Engineering in Medicine and Biology Conf.*, Orlando, FL, 1992.

[6] J. R. Lackner and P. DiZio, "Rapid adaptation to Coriolis force perturbations of arm trajectories," *J. Neurophysiol.*, vol. 72, pp. 299–313, 1994.

[7] R. Held and S. J. Freedman, "Plasticity in human sensorimotor control," *Science*, vol. 142, pp. 455–461, 1963.

[8] Z. M. Pine, J. W. Krakauer, J. Gordon, and C. Ghez, "Learning of scaling factors and reference axes for reaching movements," *NeuroReport*, vol. 7, pp. 2357–2361, 1996.

[9] R. C. Miall, D. J. Weir, D. M. Wolpert, and J. F. Stein, "Is the cerebellum a smith predictor?," *J. Motor Behav.*, vol. 25, pp. 203–216, 1993.

[10] J. W. Krakauer, M. F. Ghilardi, and C. Ghez, "Independentlearning of internal models for kinematic and dynamic control ofreaching," *Nature Neurosci.*, vol. 2, pp. 1026–1031, 1999.

[11] H. I. Krebs, N. Hogan, W. Hening, S. V. Adamovich, and H. Poizner, "Procedural motor learning in parkinson's disease," *Exp. Brain Res.*, vol. 141, pp. 425–437, 2001.

[12] C. A. Seger, "Implicit learning," *Psychological Bull.*, vol. 115, pp. 163–196, 1994.

[13] L. R. Squire, "Mechanisms of memory," *Science*, vol. 232, pp. 1612–1619, 1986.

[14] P. M. Fitts, "Perceptual-motor skill learning," in *Categories of Human Learning*, A. W. MeHou, Ed. New York: Academic, 1964, pp. 243–245.

[15] M. G. Beldarrain, J. Grafman, A. Pascual-Leone, and J. C. Garcia-Monco, "Procedural learning is impaired in patients withprefrontal lesions," *Neurology*, vol. 52, pp. 1853–1860, 1999.

[16] F. A. Mussa-Ivaldi and J. L. Patton, "Robots can teach people how to move their arm," presented at the *IEEE Int. Conf. Robotics and Automation (ICRA)*, San Francisco, CA, 2000.

[17] N. Bhushan and R. Shadmehr, "Computational nature of humanadaptive control during learning of reaching movements in forcefields," Bilological Cybern., 1999, to be published.

[18] M. A. Conditt and F. A. Mussa-Ivaldi, "Central representationof time during motor learning," *Proc. Nat. Acad. Sci.*, vol. 96, pp. 11 625–630, 1999.

[19] R. A. Scheidt, D. J. Reinkensmeyer, M. A. Conditt, W. Z. Rymer, and F. A. Mussa-Ivaldi, "Persistence of motor adaptationduring constrained, multi-joint, arm movements," *J. Neurophysiol.*, vol. 84, pp. 853–862, 2000.

[20] K. A. Thoroughman and R. Shadmehr, "Learning of actionthrough adaptive combination of motor primitives," *Nature*, vol. 407, pp. 742–747, 2000.

[21] F. A. Mussa-Ivaldi, "Modular features of motor control andlearning," *Curr. Opin. Neurobiol.*, vol. 9, pp. 713–717, 1999.

[22] D. J. Dewhurst, "Neuromuscular control system," *IEEE Trans. Biomed. Eng.*, vol. BME-14, pp. 167–171, 1967.

[23] N. Hogan, "Mechanical impedance of single and multi-articularsystems," in *Multiple Muscle Systems*, S. L.-Y. Woo, Ed. New York: Springer-Verlag, 1990, pp. 149–164.

[24] J. L. Patton and F. A. Mussa-Ivaldi, "PrespecifiedAfter-effects elicited from robotic force fields," *Neurosci. Abstr., 2000*, 2000.

[25] J. M. Hollerbach and T. Flash, "Dynamic interactions betweenlimb segments during planar arm movements," *Biological Cybern.*, vol. 44, pp. 67–77, 1982.

[26] Z. Ghahramani, D. M. Wolpert, and M. I. Jordan, "Generalization to local remappings of the visuomotor coordinate transformation," *J. Neurosci.*, vol. 16, pp. 7085–7096, 1996.

[27] S. Schaal and C. G. Atkeson, "Constructive incrementallearning from only local information," *Neural Computation*, vol. 10, pp. 2047–2084, 1998.

[28] M. Kawato, "Feedback-error-learning neural network forsupervised learning," in *Advanced Neural Computers*, R. Eckmiller, Ed. Amsterdam, The Netherlands: North-Holland, 1990, pp. 365–372.

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

[30] J. Albus, "A theory of cerebellar function," *Math. Biosci.*, vol. 10, pp. 25–61, 1971.

[31] D. Marr, "A theory of cerebellar cortex," *J. Physiol.*, pp. 437–470, 1969.

[32] Y. Matsuoka, *Models of Generalization in Motor Control*. Cambridge, MA: M.I.T. Press, 1998.

[33] P. S. Lum, C. G. Burgar, D. E. Kenney, and H. F. Van der Loos, "Quantification of force abnormalities during passive andactive-assisted upper-limb reaching movements in post-strokehemiparesis," *IEEE Trans. Biomed. Eng.*, vol. 46, pp. 652–662, June 1999.

[34] H. I. Krebs, N. Hogan, M. L. Aisen, and B. T. Volpe, "Robot-aided neurorehabilitation," *IEEE Trans. Rehab. Eng.*, vol. 6, pp. 75–87, Mar. 1998.

[35] B. Bobath, *Adult Hemiplegia: Evaluation and Treatment*, 2nd ed. London, U.K.: Heinemann, 1978.

[36] D. R. Voss, M. K. Ionta, and B. J. Myers, *Proprioceptive Neuromuscular Facilitation*, 3 ed. Philadelphia, PA: Harper & Row, 1985.

[37] J. Patton, F. Mussa-Ivaldi, and W. Rymer, "Altering movement patterns in healthy and brain-injured subjects via custom designed robotic forces," presented at the *Embc2001: 23rd Annu. Int. Conf. IEEE Engineering in Medicine and Biology Society (EMBS)*, Istanbul, Turkey, 2001.

[38] C. C. Raasch, F. A. Mussa-Ivaldi, and W. Z. Rymer, "MotorLearning in reaching movements by hemiparetic subjects," *Soc. Neurosci. Abstr.*, vol. 924.17, p. 2374, 1997.

[39] M. J. Weiner, M. Hallett, and H. H. Funkenstein, "Adaptation to lateral displacement of vision in patients withlesions of the central nervous system," *Neurology*, vol. 33, pp. 766–772, 1983.

[40] J. N. Sanes, B. Dimitrov, and M. Hallett, "Motor learningin patients with cerebellar dysfunction," *Brain*, vol. 113, pp. 103–120, 1990.

[41] A. J. Bastian, T. A. Martin, J. G. Keating, and W. T.Thach, "Cerebellar ataxia: Abnormal control of interaction torquesacross multiple joints," *J. Neurophysiol.*, vol. 76, pp. 492–509, 1996.

[42] R. A. Magill, "1997 C. H. McCloy research lecture: Knowledge is more than we can talk about: Implicit learning in motor skillacquisition," *Res. Quart. Exercise Sport*, vol. 69, pp. 104–110, 1998.

[43] S. T. Grafton, E. Hazeltine, and R. B. Ivry, "Abstract andeffector-specific representations of motor sequences identified with PET," *J. Neurosci.*, vol. 18, pp. 9420–9428, 1998.

[44] H. I. Krebs, T. Brashers-Krug, S. L. Rauch, C. R. Savage, N. Hogan, R. H. Rubin, A. J. Fischman, and N. M. Alpert, "Robot-aided functional imaging: application to a motor learningstudy," *Human Brain Mapping*, vol. 6, pp. 59–72, 1998.

[45] A. S. Reber, F. F. Walkenfeld, and R. Hernstadt, "Implicitand explicit learning: Individual differences and IQ," *J. Exp. Psych.: Learning, Memory, Cognition*, vol. 17, pp. 888–896, 1991.

James L. Patton (M'01) was born in Ann Arbor MI, in 1965 and raised in Racine, WI. He received the B.S. degrees in mechanical engineering and engineering science/bioengineering from the University of Michigan, Ann Arbor, in 1989. He earned a M.S. degree in theoretical mechanics from Michigan State University, Lansing, in 1993. He received the Ph.D. degree in biomedical engineering from Northwestern University, Evanston, IL, in 1988.

He has worked for Ford Motor Company and in nuclear medicine before turning his attention to biomechanics and control of human movement. Currently, he is Associate Director of the Sensory Motor Performance Program at the Rehabilitation Institute of Chicago and at Northwestern University. His research and training blends the fields of robotics, automatic control, neurophysiology, physical therapy, and human movement science. His work has focused on haptics, modeling the human-machine interface, robotic teaching, and robotic facilitation of recovery from a brain injury.

Dr. Patton is a member of the IEEE Engineering in Medicine Biology Society and the Society for Neuroscience. He is a reviewer of IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING, *IEEE Engineering in Medicine Biology* magazine, and IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION. He also is on the program committee of upcoming IEEE conference biomedical robotics.

Ferdinando A. Mussa-Ivaldi was born in Torino, Italy, in 1953. He received the Laurea degree in physics from the University of Torino in 1978 and the Ph.D. degree in biomedical engineering from the Politecnico of Milano, Milan, Italy, in 1987.

He is Professor of Physiology, Physical Medicine and Rehabilitation and Biomedical Engineering at Northwestern University, Evanston, IL. He is Director of the Robotics Laboratory and of the Neural Engineering Initiative of the Rehabilitation Institute of Chicago. His areas of interest and expertise include robotics, neurobiology of the sensory-motor system, and computational neuroscience. Among his achievements are the first measurement of human arm multijoint impedance, the development of a technique for investigating the mechanisms of motor learning through the application of deterministic force fields, the discovery of a family of integrable generalized inverses for redundant kinematic chains, the discovery of functional modules within the spinal cord that generate a discrete family of force-fields, the development of a theoretical framework for the representation, generation, and learning of limb movements, and the development of the first neuro-robotic system in which the brainstem of a lamprey controls the behavior of a mobile-robot through a closed-loop interaction. He has authored 98 full-length publications and 66 abstracts.

Dr. Mussa-Ivaldi is a member of the Society for Neuroscience.