ARTICLE

Can Robots Help the Learning of Skilled Actions?

David J. Reinkensmeyer[1] *and James L. Patton*[2]

[1]Departments of Mechanical and Aerospace Engineering and Biomedical Engineering, University of California at Irvine, Irvine, CA; and [2]Department of Bioengineering, University of Illinois at Chicago, Sensory Motor Performance Program Rehabilitation Institute of Chicago, Chicago, IL

REINKENSMEYER, D.J., and J.L. PATTON. Can robots help the learning of skilled actions? *Exerc. Sport Sci. Rev.*, Vol. 37, No. 1, pp. 43–51, 2009. *Learning to move skillfully requires that the motor system adjusts muscle commands based on ongoing performance errors, a process influenced by the dynamics of the task being practiced. Recent experiments from our laboratories show how robotic devices can temporarily alter task dynamics in ways that contribute to the motor learning experience, suggesting possible applications in rehabilitation and sports training.* **Key Words:** motor control, robotics, adaptation, movement, training

INTRODUCTION

Many people have a vested interest in improving their motor skills, ranging from people with movement disorders to weekend athletes looking to improve their games and to highly skilled athletes, surgeons, and machine operators whose careers depend on their skill level. Based on our recent work (2,4,11,15,18,19), we developed the hypothesis that robotic devices can enhance the motor learning experience by temporarily altering task dynamics during movement practice. This article will critically evaluate this hypothesis. By task dynamics, we mean the physical relationship between the forces a person applies to an object and the resulting motion of the object. Altering this relationship would be expected to change the learning process because the motor system adjusts motor commands based on movement performance, and task dynamics influence performance. However, the key question is whether altering task dynamics can produce beneficial effects for learning. As we review later, our initial robotic experiments have indeed shown some benefits using this approach, but these benefits are only suggestive at this point, having typically been applied to laboratory tasks in which learning occurs relatively rapidly, and not yet having led to practical training techniques. Nevertheless, these experiments may eventually lead to new technology that improves long-term skill learning in practical tasks.

There is historical precedent for this type of approach: consider "hand-over-hand" training in rehabilitation and sports training or technology that temporarily alters the dynamics of sporting equipment with the goal of improving performance, such as overhead elastic bands that allow gymnasts to practice flips, weighted doughnuts added to baseball bats for practice swings, and training wheels for bicycles. In contrast to this simpler mechanical technology, robotic devices have the potential to alter task dynamics based on computerized sensing of ongoing movements. In addition, robots can monitor performance improvements and provide feedback to the trainee or can use sensor measurements as a basis for ongoing adjustment of training strategies. Furthermore, robots are capable of applying forces in wholly new ways not possible with human trainers or traditional equipment. Robots can even make use of computational models of motor learning mechanisms to determine appropriate forces to apply, which may open new possibilities for training.

For the purposes of this article, we will focus on two aspects of the motor skill learning process. The first aspect is *trajectory learning*, which refers to the process of learning to make motions that achieve the task at hand. For example, to play tennis well, the motor system must learn what motion of the tennis racquet will make the ball move with an appropriate spin to a desired location in the opponent's court. The second goal is *dynamic adaptation*, a term we will use to refer to the process of learning the muscular forces that achieve the target motion trajectory given the dynamics of the task. Continuing the tennis example, the player must

Address for correspondence: David J. Reinkensmeyer, Ph.D., Department of Mechanical and Aerospace Engineering, University of California, Irvine, CA 92697-3975 (E-mail: dreinken@uci.edu).

Accepted for publication: September 11, 2008.
Associate Editor: Daniel P. Ferris, Ph.D.

0091-6331/3701/43–51
Exercise and Sport Sciences Reviews

adjust his muscle activations so that they achieve the desired swing given the type of racquet the player is using and the current muscular fatigue level of the player. There are other important processes involved in motor learning, such as formation of high-level movement strategies, energy optimization, and reduction of trial-to-trial variability, and in fact, trajectory learning and dynamic adaptation are not independent processes, as for example, a desired trajectory is required to determine how dynamic adaptation should occur, but this review will focus on these two aspects because most current work has focused on them.

It would be attractive if we could exogenously transform trajectory plans or the dynamic adaptation process by electronically transferring new neural programs directly into the motor system, as imagined in Hollywood movies such as *The Matrix*. Such an approach may eventually become scientifically plausible. However, at present, robotic technology provides a straightforward way to noninvasively influence motor commands simply by manipulating task dynamics. The question that this article addresses is whether manipulating task dynamics can be useful for improving motor learning, and if so, what motor learning mechanisms account for these improvements?

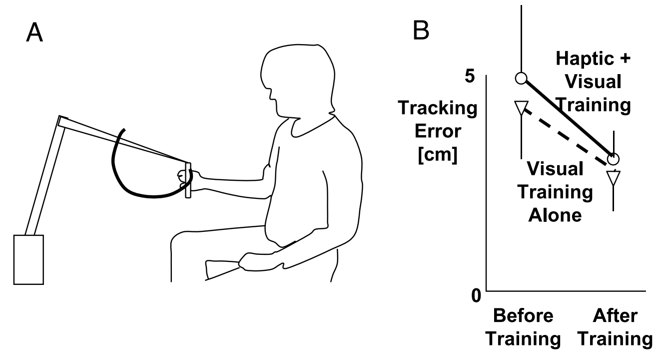

Figure 1. Effect of robot guidance on learning a desired trajectory (13). A. Subjects held the tip of a lightweight robot that guided them through a desired movement, which was a curve on the surface of a virtual sphere. B. Subjects learned to make the desired movement about as well simply by watching the robot make the desired movement (*i.e.*, with visual training episodes alone, with their hands on their laps), without being physically guided by the movement. In this experiment, there were nine training episodes of seven movements each, with test episodes of seven movements interspersed between the training episodes. Figure based on data from Liu *et al.* (13).

ROBOT GUIDANCE FOR TRAJECTORY LEARNING

How many weekend golfers have thought, "If only I could experience the golf swing of Tiger Woods, then I could rapidly improve my own swing?" Technically speaking, the possibility of physically experiencing complex movement trajectories is becoming increasingly feasible because of gradual improvements in robotic design techniques during the past 30 yrs. Robots can now be programmed to drive the user's limbs precisely along a desired trajectory of substantial complexity that has been prerecorded from an expert, using a variety of approaches, including position tracking control strategies, virtual channels with springy walls, and viscous force fields. One could experience the appropriate kinematics and learn the proprioceptive "feel" of the motion, and perhaps perceive nonobvious strategies. But would experiencing an expert's movement trajectory improve a person's ability to perform the trajectory?

Initial experiments suggest that robotic guidance is not a panacea for improving trajectory learning. For example, one experiment from coauthor Reinkensmeyer's laboratory compared how well healthy adult subjects learned to make a curved trajectory in a three-dimensional space with the arm when they were provided with 1) visual demonstration of the trajectory by the robot; or 2) visual and haptic demonstration of the trajectory by the robot (Fig. 1) (13). The term *haptic* refers to the human sense of touch, including cutaneous and kinesthetic sensation, so whereas in the first condition, subjects kept their hands on their laps during the robot demonstration and only saw the robot move, in the second condition, the subjects held the tip of the robot and felt the movement of the robot, thereby receiving haptic input as well as visual input. Furthermore, in this second condition, subjects were required to actively move along with the robot, rather than just remaining passive and letting the robot drag their arm through the desired motion. Blocks of demonstration trials were interspersed with blocks of test trials with no demonstration to evaluate the ongoing effect of practice with the visual and haptic demonstrations. For a study size of 20 subjects with 126 demonstration and test movements in total, there was no significant advantage to adding haptic information to the visual information. In fact, there was a nearly significant trend for learning to be worse when haptic information was added. It is unclear why the addition of haptic information tended to degrade learning, but one possible factor is that haptic perception of spatial relationships is systematically distorted compared with visual perception (10), and thus haptic information may degrade estimates of the required path shape. An earlier study with more complex trajectories that required precise timing also found no substantial benefit to haptic guidance, although subjects improved the timing of their movements slightly more when haptic guidance was available (5). Other studies of robotic (7,16) and nonrobotic (30) guidance with healthy adult volunteers have yielded similar results — no positive benefit of training with guidance.

The Guidance Hypothesis

Why would guiding movements not help people learn to make desired movements? The guidance hypothesis predicts that guidance can impair learning because it changes the dynamics of the task to be learned (30). Essentially, the task dynamics that people experience as a result of guidance are not the ones that they ultimately must learn how to perform in. Guidance also obviates the motor system from learning how to make the error-based corrections in response to the dynamics of the task. Thus, despite the intuition of weekend golfers, relying on haptic guidance to make performance temporarily better would not be predicted to automatically translate into better learning. Put another way, simply experiencing the right trajectory does not mean the motor system can translate the experience into an appropriate motor command. However, there still are reasons to use guidance to assist in motor learning.

Is the Effectiveness of Robot Guidance Task Dependent?

The usefulness of guidance for trajectory learning may depend on the task to be learned. For example, in another recent experiment in coauthor Reinkensmeyer's laboratory, guidance produced a modest learning benefit for a task in which healthy adult subjects learned how to drive a simulated wheelchair with a robotic steering wheel (15). One group of subjects trained with no guidance from the steering wheel, whereas the other group trained while receiving guidance from the steering wheel. The robotic guidance was designed to demonstrate to the subjects how they could optimally initiate and complete turns to minimize tracking errors as they followed a line through rooms. For one group of subjects that received robotic guidance, an adaptive computer algorithm reduced the firmness of the guidance as subjects learned to steer; for another group, fixed guidance was provided. Subjects who received adaptive or fixed guidance, taken as a group, not only exhibited smaller performance errors, but also learned better when to initiate their turns. Apparently, the experience of an early-initiated turn, imposed by the robotic steering wheel, caused them to adopt this strategy themselves, causing smaller errors when guidance was ultimately removed. The adaptive guidance group had a further advantage in that they did not exhibit increased error when guidance was finally removed; this is because the adaptive guidance algorithm had been slowly removing (or "fading") the guidance as performance improved. Based on these results, we hypothesized that guidance was a more effective training strategy in this steering task because the task did not require learning a precise trajectory. Specifically, a range of steering motions could result in small tracking errors, as long as the turning movements were performed at the correct time at a relatively high speed. Thus, one possible negative side effect of guidance — that it causes inappropriate changes in the magnitudes of motor commands — may have been precluded by the task. However, initial testing with a pinball-like task that requires precise timing without a precise movement trajectory did not confirm this hypothesis (14).

Robotic guidance may also be helpful for other trajectory learning tasks as well. For example, many skilled movements require coordinated motion of multiple joints while the performer keeps vision focused on an extracorporeal object (*e.g.*, most sports involving balls). In this case, guidance might be useful for teaching how to adjust joint movements while keeping the eyes on the ball. Another way guidance might be used is to guide subcomponents of a complex movement, allowing a person to focus on perfecting other components in isolation. Indeed, guidance can help free attention for performing other tasks, as shown with another driving study, in which subjects received haptic steering assistance, which found that subjects were able to depend less on vision and perform other secondary tasks better (8). The ability of guidance to free attention may thus be useful for focusing on learning of other task components. Finally, people with visual impairments might also benefit from a haptic modality of training. In summary, the range of applicability of physical guidance for enhancing trajectory learning, in terms of the types of tasks and target populations that might benefit from it, remains to be defined.

ROBOTIC GUIDANCE FOR ASSISTING DYNAMIC ADAPTATION

The Process of Dynamic Adaptation

Even if a person knows exactly what trajectory he needs to follow to perform a motor skill successfully, he still needs to adapt his muscle forces to achieve that trajectory for the given task dynamics. Recent research has shown that the process of dynamic adaptation is driven by the experience of performance error on the last movement (3,24). A common experimental paradigm (25) asks subjects to reach to a target while a lightweight robot unexpectedly perturbs their movement with a force that depends on the movement of the hand (a "force field"). Many of the training techniques that we will discuss later, in some way, use this basic paradigm, and so we describe it in some detail now. We note that the ability to adapt to such a force field is not unique to reaching movements, as a similar pattern of adaptation has been found when a force field is applied to the leg during walking (4) or even when a force field is applied to a golf putter during putting practice (Fig. 2). We will use the golf-putting example to illustrate the basic phenomenon.

In the experiment shown in Figure 2, we attached a lightweight robot to the head of a putter and used to unexpectedly make the putter feel like it was moving through mud, using a viscous force field that was generated by applying a force with the robot proportional to the velocity of the head of the putter. Application of the force field initially slowed the putter head, causing the subject to putt short of the hole. But the subject gradually adapted to the force field, as evidenced by a return of the putt accuracy toward normal in the course of tens of movements. Removing the force field after adaptation caused an "after effect", in which the subject putted too hard, causing a large putting error in the opposite direction. The presence of this after effect is evidence that the motor system had formed a model of the new putting task dynamics and used it to predict the muscle forces needed to compensate for the viscous force field.

How does the motor system form an internal model of new task dynamics, such as a suddenly viscous golf putter? The process of internal model formation is well captured by a mathematical model, in which the motor system measures movement error resulting from the current movement attempt and increments the forces to be applied on the next movement attempt proportionally to this error in the direction that will reduce the error, so called error-based learning (3,24). The incremental modification of the forces to be applied during the next movement can also be viewed as the formation of an internal model of the task dynamics because the motor system keeps a record of how forces depend on limb position and velocity, rather than just remembering a time-based trajectory of forces (1). Often, movement trajectories do not return completely toward

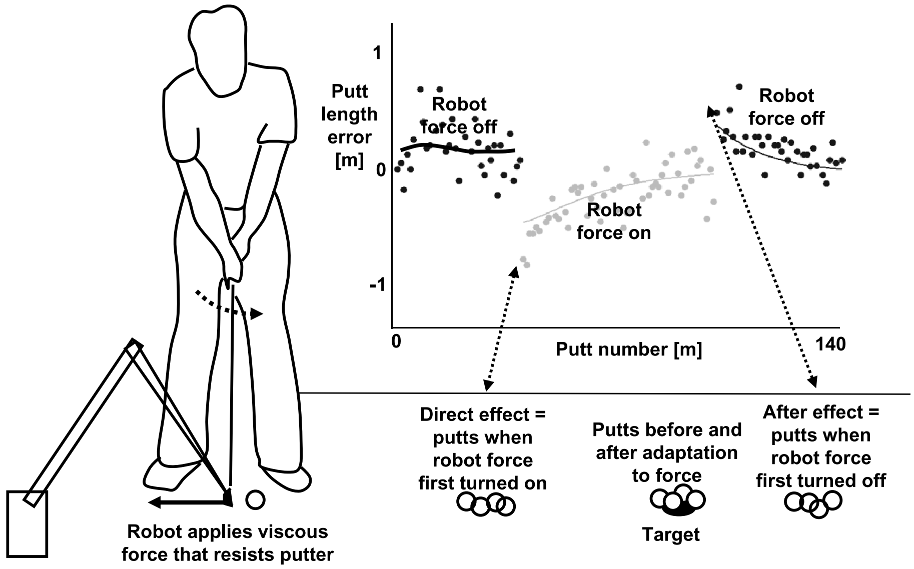

Figure 2. Example of dynamic adaptation during golf putting. A lightweight robot was attached to the head of a putter as a golfer tried to make a 1.5-m putt on a flat surface. For trials 0–40, the robot did not apply any force to the putter. After 40 practice putts, the robot resisted the putter head movement with a viscous force field for 60 putts. Initially, the subject tended to putt short of the hole (the ''direct effect'' of the field), but decreased his putt length error with practice. When the force field was unexpectedly removed at trial 101, the subject exhibited an after effect and putted too strongly, but this ''after effect'' gradually decreased in size. Data from a University of California, Irvine student project by Dr. Eric Wolbrecht, currently at University of Idaho.

normal (*i.e.*, to their preadaptation pattern) after convergence of this adaptation process, and we have shown this to be explained by a model of learning in which the motor system minimizes effort as well as trajectory errors, which is accomplished by partially reducing the previous motor command used as the base for the next trial, in a process that has been termed "slacking" or "forgetting" (3). Finally, we note that for some force fields (notably, noisy or unstable fields), the motor system relies also on impedance control to reduce error, increasing stiffness only in the directions needed (6).

Using Guidance to Aid in Dynamic Adaptation

How might robotic manipulation aid in this dynamic adaptation process? One possibility is to use robotic guidance to assist in learning difficult tasks. Specifically, some dynamic environments may be too challenging to learn because they produce too large of initial performance errors, making practice dangerous or discouraging. Computational models of motor learning further suggest that large initial errors sometimes prevent learning with biologically plausible algorithms (23). For such environments, robotic guidance might aid in learning by reducing initial performance errors. This is essentially the same intuition behind training wheels for learning bicycle riding: make the dynamics of the task easier for a time, during initial learning.

However, just like training wheels must eventually be removed, the guidance provided by the robot must eventually be removed to allow subjects to gradually experience the true target task dynamics. The question naturally arises: what is the best way to adjust the amount of physical guidance so that a person learns to move in a novel dynamic environment without ever experiencing large errors (*i.e.*, by analogy, without ever falling off the bike)?

Coauthor Reinkensmeyer's laboratory answered this question by posing the problem in the form of a mathematical optimization for which the guiding robot tries to minimize the weighted sum of performance error and the assistive force it applied (2). Framing the problem this way allows for the solution of a trial-to-trial update law for robot guidance forces. Applied to the task of walking in a viscous force field, the resulting adaptive robot learning law cancelled the novel dynamic environment at the beginning of training, but then gradually allowed the subject to experience the dynamic environment as practice progressed (Fig. 3). The reduction of robotic guidance progressed based on the ongoing measurement of the subject's performance in an error-based learning paradigm, with a slacking term similar to the learning law the human motor system itself uses to adapt to force fields. Using this approach, subjects could learn to walk in the novel dynamic environment without experiencing large performance errors.

Note that although physical guidance can reduce performance errors and still permit learning, the very action of reducing performance errors may slow down learning. This is because dynamic adaptation is driven proportionally to performance errors, as previously described. An important direction for future research is therefore to define the tradeoff between the rate of learning and the amount of guidance provided during training.

Using Guidance to Make Learning Safer and Less Discouraging

To conclude the discussion of robotic guidance for movement training, we would like to highlight an important application of robotic guidance that is easy to miss. Although robotic guidance may not speed up learning or substantially improve the completeness of learning, the experiments

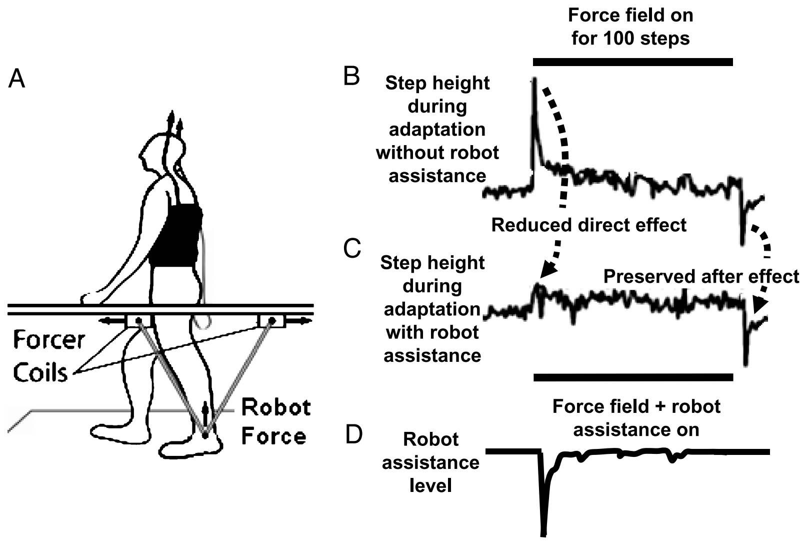

Figure 3. Robot assistance can reduce performance errors while still allowing learning. A. The robot applied a force field to the subject's ankle that forced the leg upward during swing, with a force proportional to the forward velocity of the ankle. B. Without robot assistance, application of the force field caused the subjects to step abnormally high, but they learned to step at a more normal height after several steps. Removal of the force field caused an after effect in which the subjects stepped lower than normal, indicating formation of an internal model of the robot force field. C. With robotic assistance, subjects never experienced large step height errors, and yet still learned a model of the force field, as evidenced by a similar after effect when the field was removed. D. The robot assistance started at a high level to cancel the force field but decreased from step to step according to an error-based learning law with forgetting. [Adapted from Emken JL, Benitez R , Reinkensmeyer DJ. Human-robot cooperative movement training: learning a novel sensory motor transformation during walking with robotic assistance-as-needed. *J. Neuroeng. Rehabil*. 2007;4:8. Copyright © 2007 BioMed Central. Used with permission.]

previously described show that it can permit normal amounts of learning while limiting performance errors. Thus, robotic guidance could play an important role in learning skills for which large errors are dangerous or undesirable. For example, the motivation for the driving experiment previously described was to automate powered wheelchair driver's training for people with severe movement disability. Guiding the driver's hand with a robotic joystick as he or she follows a driving course would allow the driver to practice driving safely while still learning the required skill to control the joystick. It is conceivable that a similar approach could be used for teaching people with or without movement impairments to safely drive other kinds of vehicles such as airplanes or race cars. Large performance errors can also be dangerous during sports such as gymnastics and surfing, rehabilitation tasks such as walking or transferring from a bed to a wheelchair, and surgical techniques practiced by new surgeons on actual patients. Note that although some tasks, such as driving and surgery, can be practiced safely in virtual environments without guidance, others such as walking, cannot be learned "outside" the body but still may be dangerous to practice, so guidance may be particularly useful for such tasks. Furthermore, some tasks, although not dangerous, may be frustrating to practice because of large performance errors — consider a child with poor coordination learning to write or a severely weakened stroke patient who is asked to practice reaching on his own without a therapist or caregiver by his side to verbally or physically encourage him. For such tasks, a primary benefit of robotic guidance may be to make the task less discouraging, and therefore more likely to be practiced, while still allowing learning with subsequent practice.

USING DYNAMIC ADAPTATION TO TEACH DESIRED TRAJECTORIES

Coauthor Patton and collaborators recently proposed an antithetical way to teach a desired trajectory, compared with robotic guidance, that makes use of the dynamic adaptation process itself to create afteraffects that follow the target trajectory (17). The approach was first implemented as follows for teaching a healthy person a simple curved reaching movement (Fig. 4). First, as the subject practiced reaching, the robot unexpectedly moved the hand so that it followed the desired curved reaching trajectory and measured the forces required to make the subject's hand move along this trajectory. After this machine-learning phase, the robot then applied the mirror-image of these forces during a movement training phase. Initially, this force field displaced the subject's hand to the opposite side of the desired path, increasing the error. However, with only the instruction to move to the target in a desired amount of time, the subject learned to make a nearly normal straight line reaching movement again with practice. The important result is that when the force field was then unexpectedly removed, the subject exhibited an after effect in which the hand moved along the original desired curved path. This learned movement was unexpected to the subject, but it was precisely the intent of the experimenter.

This after-effect technique is antithetical to robotic guidance because it increases trajectory errors during learning, rather than reducing them, and thus it has been called an *error-augmentation strategy*. Using this strategy, the subject produces precisely the correct muscle activation patterns to achieve the desired trajectory during the after effect. In

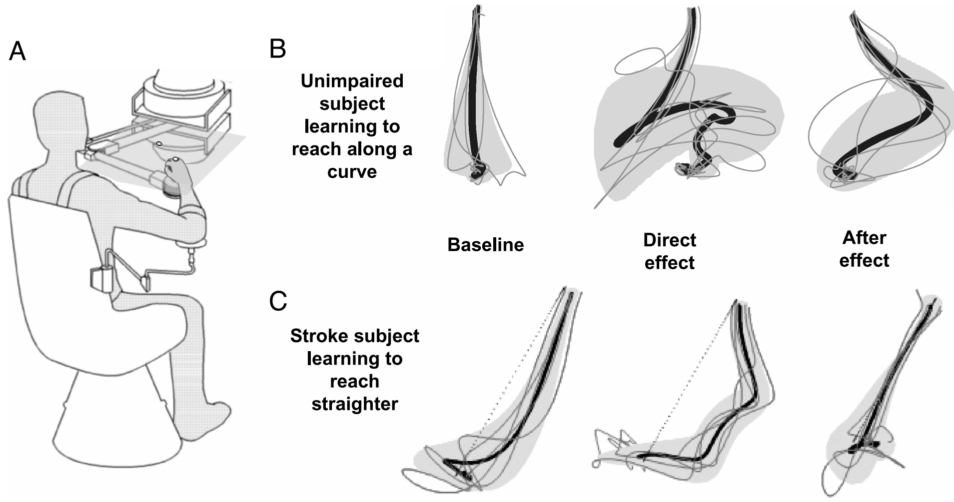

Figure 4. Using aftereffects to learn a desired trajectory. A. Subjects held the tip of a robot and reached to targets in the horizontal plane, with the arms supported. B. Using after effects to teach a desired trajectory to a healthy subject. Subjects initially reached straight to target. During training, the robot perturbed the reaching movements in the opposite direction from the desired movement using a force field. When the force field was removed, the after effects roughly followed the desired target trajectory. C. Using after effects to straighten the curved reaching path of a stroke patient. The results were similar to those in (B) for the healthy subject, except that for the stroke patient, the initial movement was curved and the after effect was straight. [Adapted from Patton J.L., and F.A. Mussa-Ivaldi. Robot-assisted adaptive training: custom force fields for teaching movement patterns. *IEEE Trans. Biomed. Eng*. 2004;51:636–646. Copyright © 2004 IEEE. Used with permission; and from Patton, J, Kovic, M, Mussa- Ivaldi F. Custom-designed haptic training for restoring reaching ability to individuals with poststroke hemiparesis. *J. Rehabil. Res. Dev.* 2006;43:643–656. Used with permission.]

contrast, when a robot enforces movement along the desired trajectory using guidance, errors are reduced and movement performance looks good, but the subject could be doing essentially anything with their muscle activations. When an aftereffect is created that follows the desired movement trajectory, the subject experiences both the correct motion and the correct muscle activations.

Implications and Limitations of Using Aftereffects to Teach Desired Movements

One intriguing aspect of trajectory learning via after effects is that it might make trajectory learning a relatively subconscious and automatic process — like the force adaptation process itself. Imagine practicing your flawed golf swing while a robot applies subtle forces to your club at key moments of your swing. You subconsciously and automatically adapt to those forces and return back to your normal flawed swing. But when the forces are then suddenly removed, you find yourself swinging with the perfect swing, for at least one swing. A key question is: would this experience of making the right movement with the right muscle activations, albeit unexpected, be enough to help you learn how to correct your swing, at least for the next round of golf?

Indeed, we observed that after effects usually wash out quickly with subsequent movement practice, but that this washout can be slowed by reducing or altering visual feedback (17). In addition, there is some evidence that washout is reduced if the after effect is desired by the operator. For example, in a study that used robotic forces to help learn visual distortions, it was hypothesized that it would be possible to teach subjects to compensate for the visual distortion using help from the robotic forces (28). Two groups of subjects learned a visual rotation, and whereas one group practiced hundreds of movements on that rotation, the other group trained instead with a special force field that was designed so that subject would make an after effect that was the appropriate movement for the visual rotation condition. Results showed that subjects reduced trajectory errors more rapidly in the force-field condition, suggesting that cutaneous and proprioceptive force feedback enhanced the learning process. Determining how long these after effects can be kept and how fast and how much people can be helped to learn with error-augmentation techniques compared with other training approaches is an important direction for future research.

Error Augmentation for Rehabilitation of Movements by Stroke Patients

Coauthor Patton and collaborators have also applied the error-augmentation technique in a rehabilitation context to improve the coordination of reaching movements made by stroke patients (Fig. 4B) (19). When stroke patients with moderate to severe motor impairment perform supported reaching movements in the horizontal plane, they typically make movements with an abnormal curvature that is repeatable. It is possible to create subject-specific force fields that amplify this curvature error using the algorithm previously described. When the stroke patients were exposed to their custom-designed force field, they automatically adapted to it, and their movement trajectories slowly adapted toward their old abnormally curved movements. We found that when the force field was unexpectedly removed, however, the after effect was such that the patients moved along a straighter path than they originally did. Interestingly, force fields that decreased errors during training had the opposite result, producing aftereffects with increased curvature

(18). Preliminary studies on stroke survivors suggest that after effects may persist longer when the after effects resemble desired movements (20). After-effect persistence might be further improved using refresher training sessions over many days to bring the performer increasingly closer to the desired outcome.

Another recent study used a split-belt treadmill to alter interlimb coordination of people with stroke and also found evidence of motor adaptation resulting in beneficial after effects (22). The treadmill belt under the paretic leg was operated at a higher speed than the one under the nonparetic leg for a few minutes, causing the subjects to walk with an even greater step size asymmetry than usual, and also provoking the subjects to adapt their stepping pattern. When the treadmills were reverted to the same speed, subjects tended to walk with a more symmetric gait, at least for a few minutes.

These experiments demonstrate that people with stroke can alter abnormal limb movement patterns that appear relatively fixed and suggest that an experience of increased error might drive such adaptation. However, the reasons why people with stroke elect to use abnormal patterns in the first place (*e.g.*, are they more efficient in some sense?) and the extent to which error-augmentation techniques can produce persistent benefits to patients that are functionally meaningful remain to be identified.

AUGMENTING TASK DYNAMICS TO IMPROVE DYNAMIC ADAPTATION

Recent experiments also illustrate that it is possible to improve the rate or extent of adaptation to new task dynamics by temporarily altering those dynamics. For example, coauthor Reinkensmeyer's laboratory showed that it is mathematically possible to calculate an amount of force field amplification, which, if provided on the first trial of exposure to a force field, will induce adaptation to the field in one trial (4). This approach was applied to the task of learning to walk in a viscous force field (Fig. 5). Subjects experienced a stronger version of the force field for only the first step in the field, which caused them to increment their next muscle command with a larger increment, causing quicker adaptation.

An implication of this experiment is that transiently amplifying the dynamics of a new piece of equipment or a new task might increase the rate of learning. Of course, for real-world applications, such a technique would only make sense to apply to tasks for which it takes many trials to learn the task dynamics. In the experiment previously described, learning to walk in the force field took only approximately 10 steps to learn, and therefore the practical benefits of the error-amplification technique were small. There is no evidence yet that this technique works for other tasks that take longer to learn, although from a theoretical point of view, the technique should have at least some beneficial effect on any task in which the experience of kinematic error causes a roughly proportional change in motor command on the next movement attempt.

An alternate approach to modifying task dynamics to promote dynamic adaptation is to change the task dynamics

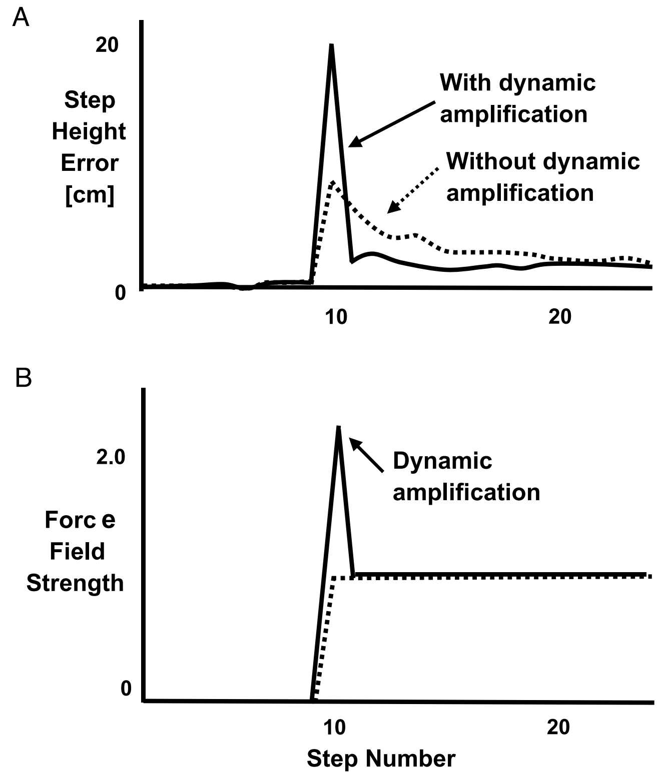

Figure 5. Accelerating motor adaptation by transiently amplifying the dynamic environment. Unimpaired subjects walked with a lightweight robot attached to their lower shank (Fig. 3A). A. At step number 10, the robot applied a force field to the lower shank that pushed the leg upward during swing with a force proportional to the forward velocity of the lower shank. Subjects adapted to this force field in about 15 steps (A, *dashed line*). By roughly doubling the force field strength on only the first step in the force field (B, *solid line*), subjects adapted more quickly to the force field (A, *solid line*). Figure based on data from Emken and Reinkensmeyer (4).

in a way that encourages a greater exploration of the task. For example, superimposing an unstable viscous force field on an anisotropic inertial force field caused subjects to adapt better to the inertial force field (11). Specifically, after subjects were given a chance to perform "free play" in the combined viscous/inertial field (*i.e.*, make whatever movements they wished), they exhibited significantly reduced error in the isolated inertial field compared with a control group that had performed a matched duration of free play but without the superposition of the unstable viscous field. Our working hypothesis to explain this result is that subjects experienced a wider range of acceleration states because of the velocity-amplifying nature of the unstable viscous force field, allowing them to build a more accurate internal model of the inertial force field (11). In addition, subjects evidently represented the viscous and inertial force fields separately because when the viscous component was removed after free play, they still maintained the inertial model, allowing them to perform well in the inertial field.

This experiment suggests that it may be possible to temporarily augment task dynamics in ways that give subjects a richer experience of the target task. If internal models are indeed modular, then the internal model for the target task may still be available for use after removal of the augmented task dynamics.

SUMMARY/CONCLUSIONS

This article described how robotic devices can temporarily alter task dynamics in ways that enhance motor learning and identified motor learning mechanisms that may account for these effects (Fig. 6). Perhaps the most obvious approach to enhancing motor learning with robots is to physically guide movements so that the practiced movements look like the ones the subject wants to learn. Indiscriminate application of this approach, however, does not seem to produce much-benefit in the amount of learning compared with visual presentation of the desired trajectory. Thus, visuomotor learning mechanisms do not appear to be much enhanced by the addition of haptic information, except for a driving task, as described above.

However, physical guidance still has an important role to play in that it can make difficult tasks doable while still allowing learning. Appropriately designed robotic guidance can gradually allow people to experience more of the actual task dynamics while limiting performance errors. Given a computational model of the learning process of interest, optimization theory provides a means to design such adaptive guidance. The mechanism of benefit for adaptive guidance in this paradigm may simply be that people are able to learn skills that they normally would not have attempted to try to learn because they are too difficult or dangerous.

An alternate approach to enhancing motor learning with robots is to use robots to augment performance errors. The motor learning mechanism that this approach seeks to manipulate is that performance errors drive motor adaptation, so techniques that change the size of performance errors should affect the rate and completeness of motor learning. We described how this approach can be used to teach desired movements using after effects, temporarily correct the coordination of people with a stroke, and to speed up motor adaptation to a novel dynamic environment. All of these techniques are currently just intriguing laboratory examples; real-world applications remain to be defined, as well as the extent to which after effects can be made permanent.

A final approach that we discussed was to temporarily augment task dynamics in a way that causes the subject to obtain a richer experience of the task. This approach is built on the premise that the motor system can extract the dynamic model required for the target task when the augmented task dynamics are removed. The putative motor learning mechanism that this technique takes advantage of is the construction of better representations of tasks when encountering a wider range of experiences.

In some applications, the approaches of physically guiding movement and augmenting performance errors may be complementary. Guidance might allow a person to begin to perform movements that are too difficult or dangerous to practice without guidance. Then, superimposing a differential amount of error augmentation on this background level of guidance may provoke faster or better learning. Guidance could then be gradually removed and error augmentation increased as learning proceeds. To give a practical example of this complementary concept: a patient who cannot support his body weight and move his legs adequately could be provided with robotic guidance to help him begin practicing walking. The same robot could be used to amplify stepping errors to provoke learning. As the person gets better at walking, the amount of guidance could be decreased, and the perturbation could be increased.

Analogs of many of these techniques can be achieved without robotic devices. For example, visual guidance can be provided by showing subjects the movement to be performed and asking them to copy it, or auditory feedback could be used to indicate performance errors. Visual demonstration can promote learning of skilled tasks (5,13,27), but it does not physically limit errors, and thus does not have the same safety benefits that physical guidance has. Nevertheless, visual and auditory feedback techniques may be a more viable approach than a robotic approach for many motor learning applications. Errors can also be enhanced using visual feedback alone (29) rather than by manipulating task dynamics, again an arguably simpler approach because it does not require a robot. The appropriate applications for haptic versus visual error augmentation remain to be defined.

So far, we have applied the techniques described to simplified laboratory tasks, with the exception being that robotic guidance is now being used in many rehabilitation clinics to help neurological patients train both upper-extremity and lower-extremity movements (9,21). Results

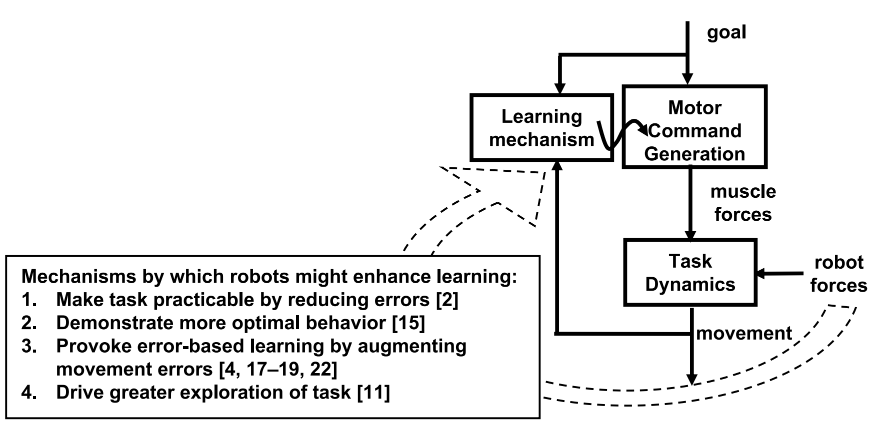

Figure 6. Block diagram illustrating physical pathway and learning mechanisms by which a robotic device might augment the learning experience. By altering task dynamics, robotic devices can make a task practicable or alter the flow of sensory information into learning centers in ways that provoke learning.

from movement rehabilitation studies that have compared a matched amount of guided or unguided practice after stroke are equivocal (12,26).

With further development, however, the principles behind the techniques described here may eventually lead to new technology for learning tasks of practical value, such as sports techniques, surgical techniques, and machine operations, as well as rehabilitation exercise. Research to date in this area has focused mainly on reaching and walking movements, but there could be applications for speech production, balance, eye movement, posture, and handwriting, among other motor systems and tasks. Research has also focused mainly on the motor component of motor training, but if these techniques influence the way that the nervous system attends to its experiences, then there is unexplored potential in application areas that involve attention deficits such as traumatic brain injury and attention deficit disorder. As previously mentioned, providing robotic guidance can reduce perceptual demands and free cognition for a secondary task, which could benefit learning of complex tasks (8). Finally, research in this area has focused exclusively on changing the mean motor behavior of subjects, averaged over many trials. However, trial-to-trial movement variability is a key factor limiting performance for many tasks, especially in sports. An important direction for future research is to define computational models of how the motor system reduces variability with motor training to determine whether manipulations of task dynamics can enhance this process as well. A detailed consideration of the fundamental computational mechanisms involved in motor skill learning will help in the development of creative ways to exploit robotic technology for encouraging the motor system to learn skilled actions.

Acknowledgments

This study was supported by the Department of Education National Institute on Disability and Rehabilitation Research grant number H133E070013, NIH R01 NS053606, and NIH N01-HD-3-3352.

References

1. Conditt MA, Gandolfo F, Mussa-Ivaldi FA. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J. Neurophysiol.* 1997; 78:554–60.
2. Emken JL, Benitez R, Reinkensmeyer DJ. Human-robot cooperative movement training: learning a novel sensory motor transformation during walking with robotic assistance-as-needed. *J. Neuroeng. Rehabil.* 2007; 4:8.
3. Emken JL, Benitez R, Sideris A, Bobrow JE, Reinkensmeyer DJ. Motor adaptation as a greedy optimization of error and effort. *J. Neurophysiol.* 2007; 97:3997–4006.
4. Emken JL, Reinkensmeyer DJ. Robot-enhanced motor learning: accelerating internal model formation during locomotion by transient dynamic amplification. *IEEE Trans. Neural Syst. Rehabil. Eng.* 2005; 13:33–9.
5. Feygin D, Keehner M, Tendick F. Haptic guidance: experimental evaluation of a haptic training method for a perceptual motor skill. Proc. 10th Int. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems (Haptics 2002). 2002; Orlando, FL, pp. 40–7.
6. Franklin DW, Liaw G, Milner TE, Osu R, Burdet E, Kawato M. Endpoint stiffness of the arm is directionally tuned to instability in the environment. *J. Neurosci.* 2007; 27:7705–16.
7. Gillespie B, O'Modhrain S, Tang P, Pham C, Zaretsky D. The virtual teacher. Proceedings of ASME Dynamic Systems and Control Division, Symposium on Haptic Interface for Virtual Environment and Teleoperator Systems. 1998; 64:171–8.
8. Griffiths P, Gillespie RB. Sharing control between human and automation using haptic interface: primary and secondary task performance benefits. *Hum. Factors.* 2005; 47:574–90.
9. Harwin WS, Patton JL, Edgerton VR. Challenges and opportunities for robot-mediated neurorehabilitation. *Proc. IEEE.* 2006; 94:1717–1726.
10. Henriques DYP, Flanders M, Soechting JF. Haptic synthesis of shapes and sequences. *J. Neurophysiol.* 2004; 91:1808–21.
11. Huang F, Patton JL, Mussa-Ivaldi F. Interactive priming enhanced by negative damping aids learning of an object manipulation task. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* 2007; 1:4011–4.
12. Kahn LE, Zygman ML, Rymer WZ, Reinkensmeyer DJ. Robot-assisted reaching exercise promotes arm movement recovery in chronic hemiparetic stroke: a randomized controlled pilot study. *J. Neuroeng. Neurorehabil.* 2006; 3:12.
13. Liu J, Cramer SC, Reinkensmeyer DJ. Learning to perform a new movement with robotic assistance: comparison of haptic guidance and visual demonstration. *J. Neuroeng. Rehabil.* 2006; 31:20.
14. Marchal L, Reinkensmeyer DJ. Effect of robotic guidance on motor learning of a timing task. Proceedings of the 2nd Biennial IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics, Scottsdale, AZ, USA, October 19–22, 2008, pp. 199–204.
15. Marchal Crespo L, Reinkensmeyer DJ. Haptic guidance can enhance motor learning of a steering task. *J. Mot. Behav.* 2008; 40:545–56.
16. O'Malley MK, Gupta A, Gen M, Li Y. Shared control in haptic systems for performance enhancement and training. *J. Dyn. Syst. Meas. Control.* 2006; 128:75–85.
17. Patton JL, Mussa-Ivaldi FA. Robot-assisted adaptive training: custom force fields for teaching movement patterns. *IEEE Trans. Biomed. Eng.* 2004; 51:636–46.
18. Patton JL, Phillips-Stoykov ME, Stojakovich M, Mussa-Ivaldi FA. Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Exp. Brain Res.* 2005; 168: 368–83.
19. Patton J, Kovic M, Mussa-Ivaldi F. Custom-designed haptic training for restoring reaching ability to individuals with poststroke hemiparesis. *J. Rehabil. Res. Dev.* 2006; 43:643–56.
20. Patton J, Mussa-Ivaldi F, Rymer W. Altering movement patterns in healthy and brain-injured subjects via custom designed robotic forces. Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society. 2001; 2: 1356–9.
21. Reinkensmeyer D, Emken J, Cramer S. Robotics, motor learning, and neurologic recovery. *Annu. Rev. Biomed. Eng.* 2004; 6:497–525.
22. Reisman DS, Wityk R, Silver K, Bastian AJ. Locomotor adaptation on a split-belt treadmill can improve walking symmetry post-stroke. *Brain.* 2007; 130(Pt. 7):1861–72.
23. Sanger TD. Failure of motor learning for large initial errors. *Neural Comput.* 2004; 16:1873–86.
24. Scheidt RA, Dingwell JB, Mussa-Ivaldi FA. Learning to move amid uncertainty. *J. Neurophysiol.* 2001; 86:971–85.
25. Shadmehr R, Mussa-Ivaldi FA. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 1994; 14:3208–24.
26. Takahashi CD, Der-Yeghiaian L, Le V, Motiwala RR, Cramer SC. Robot-based hand motor therapy after stroke. *Brain.* 2008; 131(Pt. 2): 425–37.
27. Todorov E, Shadmehr R, Bizzi E. Augmented feedback presented in a virtual environment accelerates learning of a difficult motor task. *J. Mot. Behav.* 1997; 29:147–158.
28. Wei Y, Patton JL. Forces that supplement visuomotor learning: a "sensory crossover" experiment. Proc. 12th Int. Symp. Haptic Interfaces for Virtual Envir. and Teleoperator Sys. 2004; 194–9.
29. Wei Y, Bajaj P, Scheidt R, Patton JL. Visual error augmentation for enhancing motor learning and rehabilitative relearning. Proceedings 2005 IEEE International Conference on Rehabilitation Robotics. 2005; 505–10.
30. Winstein CJ, Pohl PS, Lewthwaite R. Effects of physical guidance and knowledge of results on motor learning: Support for the guidance hypothesis. *Res. Q. Exerc. Sport.* 1994; 65:316–32.