Journal of Motor Behavior, 1997, Vol. 29, No. 2, 147–158

Augmented Feedback Presented in a Virtual Environment Accelerates Learning of a Difficult Motor Task

Emanuel Todorov
Massachusetts Institute of Technology

Reza Shadmehr
Johns Hopkins University

Emilio Bizzi
Massachusetts Institute of Technology

ABSTRACT. One can use a number of techniques (e.g., from videotaping to computer enhancement of the environment) to augment the feedback that a subject usually receives during training on a motor task. Although some forms of augmented feedback have been shown to enhance performance on isolated isometric tasks during training, when the feedback has been removed subjects have sometimes not been able to perform as well in the "real-world" task as controls. Indeed, for realistic, nonisometric motor tasks, improved skill acquisition because of augmented feedback has not been demonstrated. In the present experiments, subjects (Experiment 1, $N = 42$; Experiment 2, $N = 21$) performed with a system that was designed for teaching a difficult multijoint movement in a table tennis environment. The system was a fairly realistic computer animation of the environment and included paddles for the teacher and subject, as well as a virtual ball. Each subject attempted to learn a difficult shot by matching the pattern of movements of the expert teacher. Augmented feedback focused the attention of the subject on a minimum set of movement details that were most relevant to the task; feedback was presented in a form that required the least perceptual processing. Effectiveness of training was determined by measuring their performance in the real task. Subjects who received the virtual environment training performed significantly better than subjects who received a comparable amount of real-task practice or coaching. Kinematic analysis indicated that practice with the expert's trajectory served as a basis for performance on the real-world task and that the movements executed after training were subject-specific modifications of the expert's trajectory. Practice with this trajectory alone was not sufficient for transfer to the real task, however: When a critical component of the virtual environment was removed, subjects showed no transfer to the real task.

Key words: augmented feedback, motor learning, table tennis, virtual environment

The acquisition of new motor skills is possible only through feedback from the environment that contains information about one's actions. The types of feedback that have been studied can be broadly categorized as knowledge of results (KR) and augmented feedback, depending on whether information normally present in the environment is used as feedback or whether extra, artificially generated information is provided. Numerous KR experiments have investigated the effects of availability of information about the achievement of a predefined criterion. The general finding has been that almost any kind of KR improves the learning rate (Bilodeau & Bilodeau, 1962), although some schedules involving trials with no KR may lead to better retention (Young & Schmidt, 1992). This has underlined the perceptuomotor system's ability to improve performance by using many kinds of naturally available information about the outcome of an action.

An interesting question of great practical significance is whether artificially generated feedback—augmented feedback—can improve skill acquisition. In contrast to the KR paradigm, the value of augmented feedback in enhancing the learning rate is an open question. For simple tasks (defined as exact replications of a prespecified movement rather than achieving a desired effect on the environment), various forms of augmented feedback have been applied with mixed effects on performance. No advantage in skill acquisition has been found when auditory or visual cues are provided in addition to error information during visual tracking (Bilodeau & Rosenquist, 1964; Cote, Williges, & Williges, 1981; Karlin & Mortimer, 1963). Young and Schmidt (1992) demonstrated that in a task involving a swing of the forearm backward and then forward, augmented feedback on the reversal position enhanced performance (compared

Correspondence address: Emanuel Todorov, M.I.T., E25-526, Cambridge, MA 02139. E-mail address: EMO@MIT.EDU

with KR alone), although information about the variability of the reversal position, or the time of reversal, had no effect. Newell and Carlton (1987) trained subjects to produce a desired finger pressure as a function of time and found that a graphical representation of the subject's response improved performance. Subjects improved even further if the target finger pressure was also visualized, but only when the graphical representation was a curve with an unfamiliar shape.

Although some forms of augmented feedback have been shown to enhance learning of simple movements, the problem is that the performance gains achieved during learning seldom transferred to the real task. Studies have found that for simple movements, augmented feedback helped performance when it was present, but subjects that trained with it actually did no better or performed worse than controls when they had to perform the task without the enhanced feedback (Lintern & Roscoe, 1980; Vander Linden, Cauraugh, & Greene, 1993). If augmented feedback is to be an effective tool for teaching a motor task, it is crucial to show that its use in training accelerates learning of the task as measured by performance in the "real world," where enhanced feedback is not available.

The usefulness of augmented feedback in teaching realistic, many-degrees-of-freedom motor tasks, has not been demonstrated. One approach to teaching complex tasks has been the design of high-fidelity virtual-environment simulators. For example, Kozak, Hancock, Arthur, and Chrysler (1993) used a simulator in which subjects were trained to place empty cans on a sequence of specified positions as fast as possible. The skill acquired in the virtual environment did not transfer to the real (and equivalent) task even though the implementation was quite realistic. The experience with flight simulators has been more promising. Lintern, Roscoe, Koonce, and Segal (1990) used augmented feedback about the desired flight path in training inexperienced military pilots to land a light aircraft. Learning transferred to the real task only when the augmented feedback was presented during periods of large deviation, rather than continuously. However, the cognitive component of this task was significantly more challenging than the required motor act.

A simpler approach has been to videotape the subject's movements and show a replay after (or in some cases during) the movement. This procedure is often used in sports instruction and is exemplified by an experiment by Kernodle and Carlton (1992), whose subjects were trained to throw a ball by viewing the record of an expert executing the movement and a replay of their own movements. Showing the replay resulted in better performance than practice alone, but only when subjects were also given verbal instructions on the aspects of the replayed movement that they needed to concentrate on or how to improve them. Carroll and Bandura (1987, 1990) compared the effects of displaying a model executing a target movement (a complicated 30-s sequence) either before or continuously with the subject's movement. Continuous viewing resulted in better performance on a retention test, but the difference disappeared if self-monitoring was allowed during the retention test. Because of the low speed and long duration of the movement sequence, this task can be argued to be mostly cognitive, as the authors demonstrated.

In summary, although various forms of augmented feedback have been shown to enhance performance on simple or isometric tasks, they have not been shown to transfer to instances in which the feedback is not available. General principles for selecting the particular form of feedback for a given task have not been identified. For realistic, nonisometric motor tasks, improved skill acquisition because of augmented feedback has not been demonstrated. Furthermore, most researchers have attempted to use augmented feedback training in addition to regular practice, not as a substitute for it (i.e., control groups have not been given extra practice balancing the training).

Our main goal in this investigation was to test the usefulness of augmented feedback in acquiring a complex motor skill. We asked whether groups trained with augmented feedback would show accelerated improvement in performance in the real task as compared with control groups who practiced without the enhanced training environment. The task we chose is one that is intermediate to a difficult table tennis shot, requiring fast and precise motion of short duration along with rapid hand–eye coordination. Effectiveness of the augmented feedback training scheme was determined by measuring the ability of subjects to hit real table tennis balls to specified targets.

Augmented Feedback and the Motor Task

Most real-life arm movements involve object manipulation. The goal of these movements is often defined as a desired end effect on the environment, and a given goal can typically be achieved by a large family of movements, rather than a single prespecified movement. In this context, the problem of motor learning is ill posed; that is, there are many movements that can result in the desired goal, and it is not clear how a subject arrives at one that succeeds at the task. It seems that learning to perform such a task involves the simultaneous solution of two subproblems: (A) finding the set of constraints that any successful movement must satisfy and (B) selecting a subset of movements that are easiest to produce and control and thus can be performed reliably. The common characteristics of movements satisfying both (A) and (B) will be called *task-related invariants*. The question we were concerned with was what form of augmented feedback will be most efficient in helping the perceptuomotor system solve both subproblems and achieve high performance on the task.

Given a particular task, in our case hitting a flying ball and sending it to a target, one can analytically obtain the set of constraints that any successful movement must satisfy (Problem A). Although this is often the approach taken in robotics applications, the central nervous system appears to be using a much more empirical strategy, relying substantially on trial-and-error learning. We might expect that an efficient way of teaching the task constraints would be to provide examples of reference movements that achieve the goal (and thus satisfy those constraints). Therefore, one role of augmented feedback might be to emphasize the differences between the subject's movements and the reference movement.

Selecting a reference movement that is easy to produce (Problem B) would require knowledge of what constitutes an easy or natural movement. For example, in the case of reaching movements, it has been repeatedly observed (Atkenson & Hollerbach, 1985; Morasso, 1981) that subjects produce hand trajectories that are in a straight path with a bell-shaped velocity profile, even though such movement characteristics are not in any way implied in the reaching task. Flash and Hogan (1985) have argued that the default strategy in reaching is to maximize smoothness (by minimizing the squared third derivative of position over time). Unfortunately, for tasks involving an end-effector with six degrees of freedom (i.e., position and orientation of the paddle in three-dimensional [3D] space) and a combination of spatial and temporal constraints (i.e., hitting the ball at the right time during its flight and sending it to the target), such default strategies have not been identified. However, we might expect that these (unknown) strategies will be similar across subjects, as is the case in reaching. We can then record the actual movements of an expert and use them as reference movements in augmented feedback training.

We used a training scheme in which an expert movement is recorded and the subject is instructed to attempt to replicate that movement repeatedly. Because the goal of training is not replication of the complete body movement of the expert but teaching the task-related invariants, we presented only the portion of the expert movement that was likely to contain the relevant information. Newell and Carlton (1990) suggested a similar approach: On the basis of findings in one- and two-degrees-of-freedom tasks, they proposed that the number of degrees of freedom in the feedback and in the task should match. Relying on the "ecological perception" approach (Gibson, 1979), Lintern (1991) argued that during training the task should be simplified so that the important perceptual invariants (in this case, rate of dilation of the runway) are preserved. This was demonstrated by Lintern et al. (1990): Virtual training to land an aircraft without crosswind resulted in better transfer than training with crosswind, even though the (real-world) retention test actually involved crosswind. Indirect support has also been provided by Kernodle and Carlton (1992), who found that subjects' performance improved because of video replay only when they were told details of the recording on which they were to concentrate.

In this article, we propose that in order to focus on the relevant information in nonisometric arm movement tasks, one must give augmented feedback that represents the desired movement of the end-effector, rather than the whole arm or body. There is abundant evidence that in execution of limb movements, the kinematics of the end-effector play a prominent role as the variable being controlled. For example, in writing with chalk, the end-effector will be the tip of the chalk, whether it is held by our hand or foot, or attached to the end of a lever tied to our upper arm. It is known that handwriting retains its character even if different joints of the arm, hand, or foot serve as the end-effector (Marsden, 1982). Kinematics of simple reaching movements maintain their character (straight-line trajectories of the hand with bell-shaped velocity profiles) despite radical changes in the visual feedback (Thach, Goodkin, & Keating, 1992; Wolpert, Gharamani, & Jordan, 1994) and dynamical conditions in the environment (Flash & Gurevich, 1992; Lacquaniti, Soechting, & Terzuolo, 1982; Ruitenbeek, 1984; Shadmehr & Mussa-Ivaldi, 1994). Light-emitting diodes (LEDs) attached to several joints of a person moving in the dark are enough to generate a percept of a human figure, and even contain information about some personal characteristics (Johansson, 1973). These studies have suggested that the highest level of motor planning and control maintains a goal in terms of the kinematics of the end-effector of the limb and that much of the information about movement that the visual system extracts is contained in the end-effector kinematics. In our investigation, the end-effector was the table tennis paddle.

To minimize the information processing during training even further, we presented the movements of the expert's and the subject's paddles superimposed in the same coordinate frame and concurrently with the subject's movement. This provided on-line error feedback to the subject during the movement in the same coordinate frame in which the subject's own paddle was being displayed.

In summary, we argue that augmented feedback should include the end-effector kinematics of the subject's movement and an expert movement and that animation of the two end-effectors in the same coordinate frame should be displayed concurrently with the subject's movement. We implemented this training scheme by using a virtual environment displayed on a computer screen. It must be emphasized that the experiments described in this work do not address the efficacy of virtual reality training in general. We were testing the efficacy of the particular form of augmented feedback discussed above; virtual reality is only the tool we used to implement the desired training scheme.

EXPERIMENT 1

Apparatus and Task

The experimental setup is shown in Figure 1 (the obstacle in the middle of the table was present in Experiment 2 only). We used a black wooden table (85 cm × 245 cm) and a standard net (height of 15 cm), which was attached 165 cm away from the edge closest to the subject, along with a black tape 15 cm above the net. A 40-cm square target was

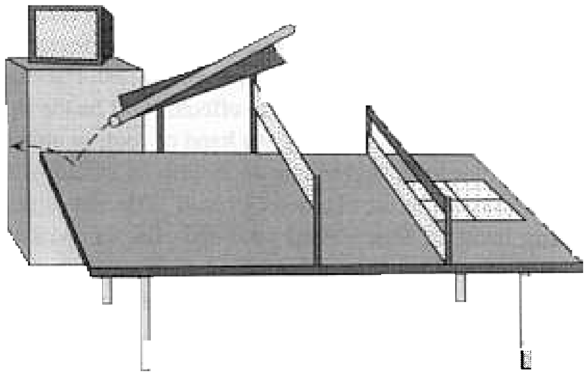

FIGURE 1. Experimental setup. Subjects held a paddle in their left hand and hit balls that were dropped through the transparent tube. The task was to send the ball between the net and the tape above it and hit the target behind the net. In Experiment 2, the horizontal obstacle in the middle was added, and subjects were asked to send the ball below it. During training, the simulation was shown on a conveniently located computer monitor; a change in body position from training to practice was not necessary.

placed horizontally on the table 20 cm behind the net. White table tennis balls were dropped (manually by the experimenter) through a transparent tube that was attached to a tilted platform on the left side of the table. The balls came out of the tube at a velocity and angle typical for an intermediate table tennis shot and bounced off the table in a highly consistent way. The subject was standing at a convenient distance in front of the table, holding a paddle in the left hand. All subjects in this study were right-handed. Therefore, our intention was to train subjects on a difficult task with their nondominant hand. The task was to let the ball bounce once, hit it with the paddle, and send it through the horizontal opening between the net and the tape such that the ball landed on the target. After every trial, the experimenter recorded the score by pressing a button. The interval between two trials was usually in the range of 5 to 10 s.

An electromagnetic sensor (IsoTrack II, Polhemus Corp.) was used to track the position and orientation of the paddle. It included a small (2 × 2 × 2 cm) receiver attached to the paddle and a transmitter placed under the table. Both the receiver and the transmitter were connected to a sensor box with a thin cable (0.3 cm diameter). The cable from the receiver was placed around the shoulders of the subject and did not restrict movements in any way. The sensor sampled the three-dimensional position and orientation of the paddle at 60 Hz and transmitted that information to a computer. During practice, the computer recorded the paddle trajectories and scores. The time when the ball was dropped or hit was not recorded.

During training, the computer displayed a realistic three-dimensional simulation of the environment (see Figure 2) on a 15-in. SVGA color monitor (in 640 × 480 pixels, 256-color mode), which was located on the left side of the table at around shoulder height of the subject. The simulated environment consisted of a graphical representation of the experimental setup, the subject's paddle, the teacher's paddle, and the ball. We used illumination, occlusion, and perspective projection to provide depth cues. The two paddles and the ball were the only moving objects in the virtual scene. We included appropriate sound effects in an attempt to replace the missing sensory information at the moment of impact between the paddle and the virtual ball. The position of the subject's paddle in the simulator was updated at 40 Hz, and the sensor delay was less than 20 ms.

The simulator could operate in three different modes. In the main training mode, the teacher's paddle executed the desired movement repeatedly (waiting for the simulated ball to come out of the tube and bounce, hitting it, and sending it to the target) while the subject was trying to move his or her paddle with that of the teacher. After every movement, a score reflecting a measure of similarity (in space and time) between the teacher's and the subject's movements was displayed. The teacher's movement was a replay of a recording of an expert player hitting a real ball and sending it to the target. In this replay, the trajectory of the ball was not recorded but was constructed from a mathematical model of the dynamics of the ball. This was possible because the ball was always delivered at the same position and velocity. This model was good enough so that the real movement recorded from the expert player hit the simulated ball and sent it in the direction of the target.

Besides the main training mode, we included two additional tools in the simulator. We could slow down the movement of the teacher to aid in the temporal matching of the student with the teacher. We also could use the teacher in a passive mode in which the teacher "followed" the subject: The subject was free to move anywhere; in each video frame, the teacher's paddle was displayed at the position along its prerecorded trajectory that was closest to the subject's current position (i.e., temporal information about the reference movement was completely ignored in this mode). This provided a way for the subject to learn the spatial component of the movement separately from the velocity profile and to adjust carefully the depth, which was more difficult to perceive than the other two dimensions.

The distances in the virtual environment were calibrated so that there was a one-to-one correspondence between the virtual environment and the physical space. Although the calibration and our model of the dynamics of the ball were good enough so that the real movement recorded from the expert player hit the simulated ball and sent it to the approximate area of the target, the fidelity of the virtual environment was not high enough so that realistic ball trajectories were simulated in response to an arbitrary paddle trajectory of the subject. In other words, we could not provide the subjects with a complete table tennis environment in which they could observe the consequences of their swings of the

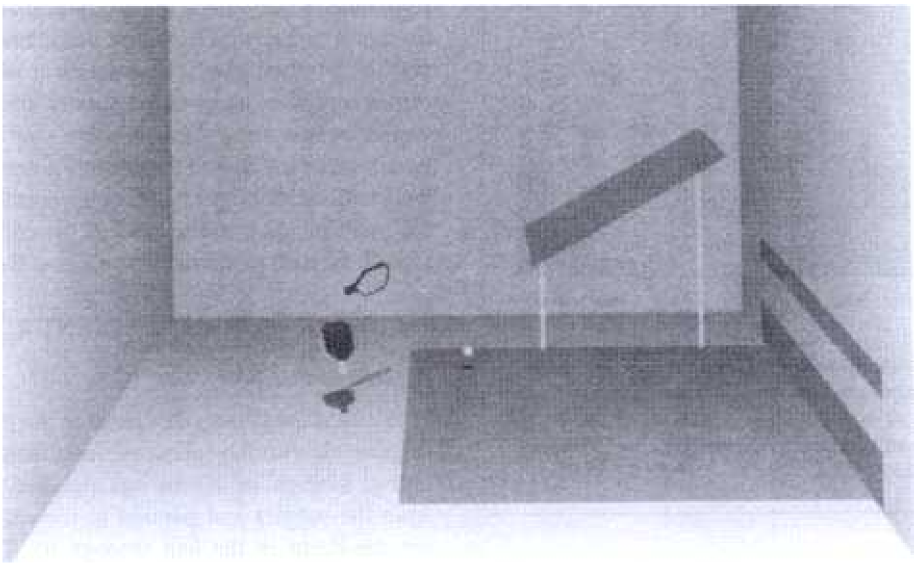

FIGURE 2. Virtual environment. A 15-in. SVGA monitor displayed a simulation of the teacher's paddle hitting the ball, superimposed on the movement of the subject's paddle (tracked by a magnetic sensor). Different colors and intensities were used for the elements in the scene, which was rendered in 648 × 480 pixels, 256-color mode.

paddle. The environment provided feedback only regarding the similarity between the subject's trajectory and the teacher's.

Procedure

Volunteer students were recruited and assigned to one of three groups: a pilot group with 13 subjects, a control group with 20 subjects, and a training group with 19 subjects. Each subject was introduced to the apparatus and the task and given 10 practice balls, the score for which was not recorded. Then, a baseline was recorded over 50 trials, which lasted approximately 10 min. Subjects in all groups were able to see where the ball landed, and therefore they knew the result of their swing, but they were not given any other feedback. After the first block, the control group was given standard coaching by the experimenter, who was an experienced player. Coaching included verbal information on the errors noticed, demonstration, and extra practice balls.

Subjects in the pilot and training groups were trained in the simulator. Training started with a slow version of the teacher's trajectory, followed by a passive mode in which the subject practiced the spatial component of the movement. This initial phase lasted 1–2 min; and after that, only the main mode (repeating the desired movement at the same time as the teacher) was used. The only difference between the pilot and training groups was that the pilot group subjects were not shown a ball in the virtual environment, whereas the subjects in the training group could see the teacher's paddle hitting a simulated ball. All groups spent 10 min in training or coaching, after which scores on a second block of 50 balls were recorded. Again, no additional feedback was given during the second block of trials. Subjects in the pilot and training groups were instructed that the primary objective was to hit the target rather than to execute the teacher's movement from the simulator.

Results and Discussion

For each subject, we computed the total number of target hits for the first and the second 50 trials, thus obtaining a measurement of performance before and after coaching or training. Improvement was defined for each subject as the absolute increase in performance from the first to the second block. Subjects with very low performance (less than 5 hits) in the second block of trials were excluded from the analysis. This included 3 subjects in the training group and 2 subjects in the coaching and pilot groups.

Figure 3 summarizes the performance results. Analysis of variance showed that all groups started with similar performance on the first 50 trials (no significant differences were found). The main effect of block was significant, $F(1, 42) = 31.4, p < .01$, indicating an overall improvement in the second block. There was a significant interaction effect, $F(2, 42) = 7.6, p < .01$. A Levene test for the homogeneity of variance assumption did not indicate any significant differences. The performance of the pilot group was not better than that of the control group; pilot subjects were actually worse than controls on the second block, and their improve-

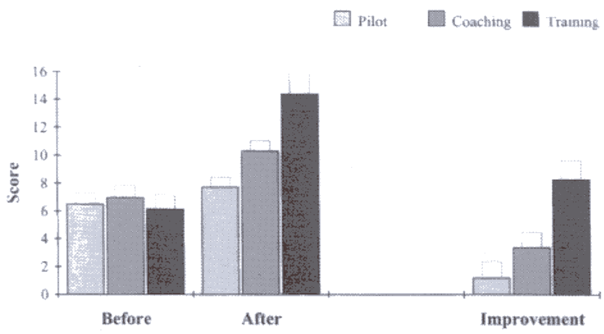

FIGURE 3. Performance in Experiment 1. The graph shows scores averaged over subjects in each group, before and after training or coaching, with standard error bars. Improvement was defined as the difference between Blocks 2 and 1 (after minus before).

ment was smaller (the differences were close to being significant). The training group, however, had significantly greater scores on the second block of trials, compared with those of both the control group (post hoc Tukey test, $p = .02$) and the pilot group (post hoc Tukey test, $p < .01$). Also, the training group showed greater improvement, compared with the control group (post hoc Tukey test, $p = .02$) and the pilot group (post hoc Tukey test, $p < .01$).

The results of Experiment 1 indicated that our augmented feedback scheme could result in better performance than a comparable amount of coaching combined with extra practice. However, the failure of the pilot experiment identified a critical component of the virtual environment: the ball. Without animation of the ball, subjects in the pilot experiment followed the teacher's trajectory without a temporal frame that related their movement to that of the ball. To send the ball to the target, however, the subject must learn not only to execute an appropriate movement but to execute it at exactly the right time (with respect to the flight of the ball). This crucial timing information was not available in the pilot training because an animation of the ball had not been used. In fact, we observed that at the beginning of the second block of 50 trials (immediately after training in the simulator), pilot subjects frequently missed the ball. This poor transfer may also have occurred because during training, pilot subjects were performing a movement that was cued by the motion of the teacher. In the real task, the movement of the arm must be cued by the motion of the ball, and there is, of course, no teacher's trajectory to follow. Indeed, when the motion of the ball was provided in the simulation, subjects performed better than controls in the real task. The failure of the pilot experiment further demonstrated that a virtual environment by itself is not a sufficient tool for training.

Although this experiment suggested that there may be advantages to using augmented feedback compared with more traditional coaching, the rapid improvements in performance indicated that the task was easier than most realistic motor tasks, which require considerably longer training periods. Furthermore, it was not clear whether the observed differences were short-term memory effects or whether they would persist over days. Finally, we felt that coaching, which was provided to the control group, could not be very well defined. It might be more appropriate to simply allow for more practice in the case of the control group. We designed the next experiment to address these issues.

EXPERIMENT 2

Apparatus and Task

We used the system described in the previous experiment (Figure 1), with the following modifications. A transparent tape at a height of 30 cm was attached to the table, 80 cm from the subject and parallel to the net (the subject could see the flight of the ball through it). The tube was made longer and less tilted, as a result of which the ball bounced closer to the end of the table, at a smaller angle and a higher velocity than before—corresponding to a more difficult shot. The task was the same as before, with the additional restriction that the ball had to pass under the transparent tape. This task was now much more difficult because it severely constrained the set of ball trajectories that could land on the target. In comparison, in the first experiment, subjects were able to choose from a wide variety of shots, including those that sent the ball in a high parabola (which is a much more controllable shot). The simulator reflected the changes in the physical setup described above and otherwise was exactly the same. A new teacher's movement was recorded, and the teacher's ball was displayed.

Procedure

A new set of students was recruited and randomly assigned to a control group with 10 subjects and a training group with 11 subjects. Each subject participated in the experiment for 3 consecutive days, two sessions per day (separated by a short break). In each session, control subjects played 30 balls, which were not scored, followed by 60 balls, which were scored—thus extra practice replaced coaching in this experiment. The experimenter provided verbal feedback on the gross errors he noticed (for control subjects only). Training subjects started each session (including the first one) with training in the simulator and then played 60 balls. Therefore, in contrast to Experiment 1, training subjects were exposed to the simulated task before the real task. In training on the first session of the 1st day, we used the same schedule in the simulator as in the previous experiment: subjects practiced for 2 min, using a slowed teacher in the passive mode, followed by training in the main mode. In all other sessions, we used only the main mode, executing the movement with the teacher at full speed. Training was terminated after the similarity score between the subject's movements and the teacher's saturated at a predefined acceptable value. As a result of this termination criterion, experimental subjects were trained longer in the first session than in any other session. The time it took the control subjects to play the first (unscored) 30 balls in each session was matched to the average training time per session for the experimental subjects.

In summary, the control and training subjects spent the same amount of time practicing. Control subjects hit 50% more real balls than the training subjects, however. The remaining practice time for the training subjects was spent in the virtual environment.

Results and Discussion

For each subject, we computed the total number of target hits for each block of 60 trials. Subjects with very low performance (no hits on the last day) were excluded from the analysis (1 subject in the control and 2 in the experimental groups, leaving 9 subjects in each group). Improvement was defined as the absolute increase in performance from the end of Day 1 to the end of Day 3.

The performance of both groups is presented in Figure 4. Analysis of variance showed that performance was increasing for both groups, $F(5, 80) = 4.0$, $p < .01$. The difference in performance between the two groups over the whole experiment was not within accepted significance levels; main effect of group, $F(1, 16) = 3.2$, $p = .09$. The reason for this was that we were looking at learning curves that were sampled often; a main effect of training, therefore, would not be expected, but rather a difference at the end of the experiment. Indeed, the score of the training group was significantly higher than the score of the control group on the last session; t test, $t(16) = 2.41$, $p = .02$. Also, the improvement of the training group was greater than the improvement of the control group; t test, $t(16) = 2.1$, $p = .05$. Interestingly, the training group was better on the very first session, $t(16) = 2.2$, $p = .04$, but the difference disappeared at the next session. The results of this second experiment suggest that the improvement caused by our training method was not a short-term effect, but that it lasted over at least a few days. In fact, the difference between the two groups increased with time.

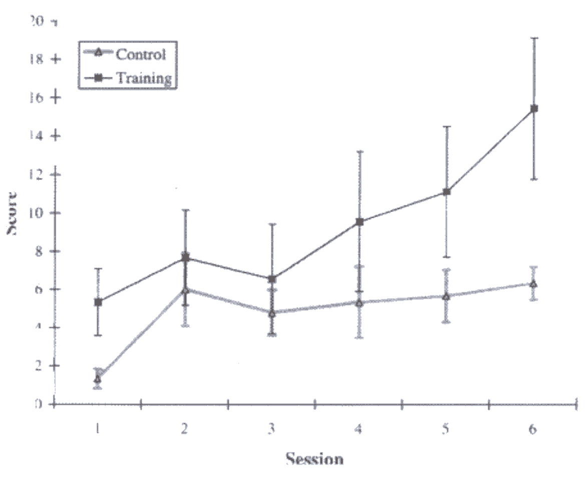

FIGURE 4. Performance in Experiment 2. Scores averaged over subjects in each group and session, with standard error bars, are shown. The score corresponded to the number of target hits (60 trials per session).

The task assigned to the subjects in Experiment 2 was much more difficult than the one used in Experiment 1, so it took more training for the difference between the two groups to become visible. A distinct feature of this experiment was that the very first session began with training in the virtual environment, before subjects had tried the real task of hitting the ball. This resulted in an immediate (short-term) difference between the two groups, which disappeared by the end of the 1st day. The reason for this probably is that we were using a very difficult task and it took a long time for the control subjects to even get close to a good trajectory, whereas the training subjects were "given" a good trajectory. It is possible that this initial difference would not have been observed if our subjects had some prior experience with the task.

Analysis of End-Effector Kinematics

In both experiments, subjects who trained in the virtual environment showed accelerated learning rates as compared with control subjects. Because training emphasized the difference in end-effector kinematics (i.e., paddle position and velocity) between a prerecorded reference movement (i.e., the teacher's movement) and the subject's movements, one would expect that learning of the reference movement would be closely related to performance. Although we argued that the kinematics of the teacher's end-effector contains the task-relevant invariants to be learned, it probably also contains expert-specific details. Thus, the subjects achieving best performance on the task will not necessarily be the ones replicating the reference movement most closely. In this section, we report on the kinematic analysis of the subject's trajectories and examine the relationship between learning the reference movement and performance.

To analyze the paddle trajectories, we had to identify (automatically) the individual movements of a subject in the continuous sensor record. The only temporal information available was the time when the experimenter recorded the score for each trial. Therefore, two consecutive scoring times specified a window that contained exactly one movement. Using the fact that subjects hit the ball very close to the point of highest velocity, we located the velocity maximum in each window and defined a fixed 1.2-s interval around it as the subject's trajectory for the trial. Visual inspection confirmed that no mistakes resulted from this simple automatic procedure and that the fixed interval was long enough to cover the complete movement. Even though the sensor sampled both position and orientation, the orientation readings were quite noisy, and we decided to exclude them from this analysis.

Figures 5, 6, and 7 show perspective projections of paddle trajectories (the viewpoint is similar to Figure 1). Each panel contains the teacher's trajectory, which is the thin tube with constant diameter (the diameter is arbitrary), and the average of all subjects' trajectories during the last 50 trials for the training, which is the thick tube. The black rings correspond to sampled positions of the paddle: Ring spacing is proportional to velocity. The ball was hit close to the point of maximum spacing between the rings (maximum velocity). The principal axis of the thick tube corresponds to the average over a set of subjects' trajectories, and the thickness at each sample point is 1/3 of the standard deviation at that point. This object was created as follows: At every sample point, the plane orthogonal to the average trajectory of all subjects was found, and the three-dimensional position along each individual trajectory at that sample point was orthogonally projected on the plane. Fitting a two-dimensional normal distribution to the points projected on each plane defined an ellipse oriented along the principal axes of the distribution. We then scaled this ellipse by a factor of 1/3 (for better visual presentation) and used the result to define the thickness of the tube at the corresponding sample point.

A)

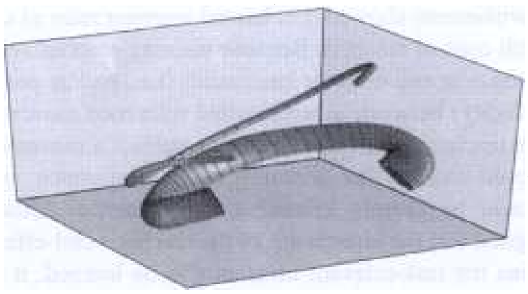

B)

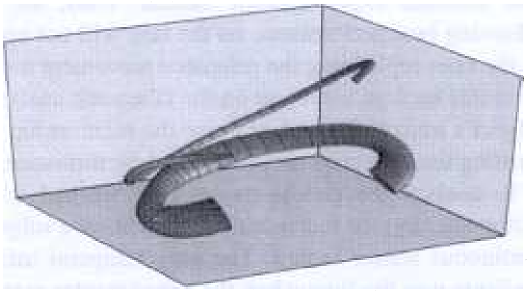

FIGURE 5. Trajectories in Experiment 1. Summarized trajectories on the last 50 trials for the training (**A**) and coaching (**B**) groups, superimposed on the teacher's trajectory. Thickness is equal to 1/3 standard deviation at each sample point; rings correspond to sample points. Viewpoint is similar to that of Figure 1. The movements started with a short backward swing, followed by a longer forward swing, and a follow-through curving backward. The point of impact was close to the point of maximum velocity (maximum ring spacing).

The pooled movements of subjects in the training and coaching groups in Experiment 1, second block of 50 trials, are summarized in Figures 5A and B, respectively. It is difficult to notice any differences, and statistical analysis, using a distance measure (described below), failed to identify a significant difference between the trajectories of the two groups. Furthermore, the movements of all subjects on the first block of trials (not shown here) had roughly the same shape. Therefore, whatever task-related invariants the training subjects learned in this case were not evident in the overall movement pattern. This result is not surprising, because the task implied a very obvious general shape of movements that would satisfy the criterion: a short swing of the paddle forward and up that sends the ball into an easily controlled high parabola, landing on the target. All subjects (and the teacher) used that same gross strategy, which makes trajectory differences resulting from training hard to

A)

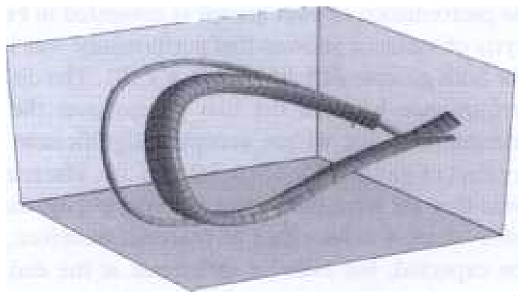

B)

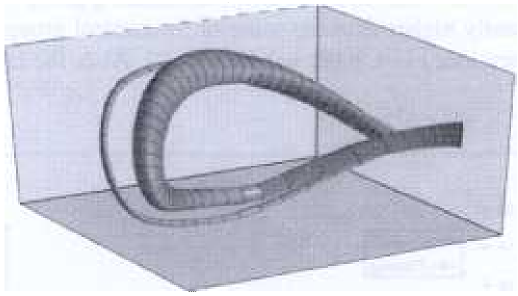

C)

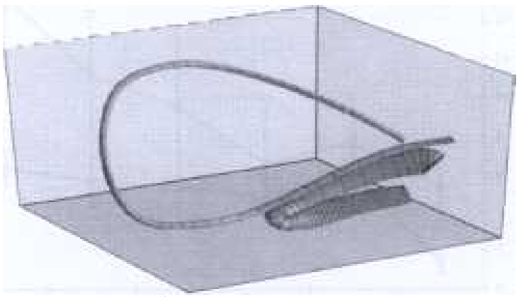

FIGURE 6. Trajectories in Experiment 2. Summarized trajectories over all six sessions for the training group in the simulator (**A**), training group in the real task (**B**), and control group (**C**), superimposed on the teacher's trajectory. The movements started with a swing backward and up, which curved down, and was followed by a forward-up swing.

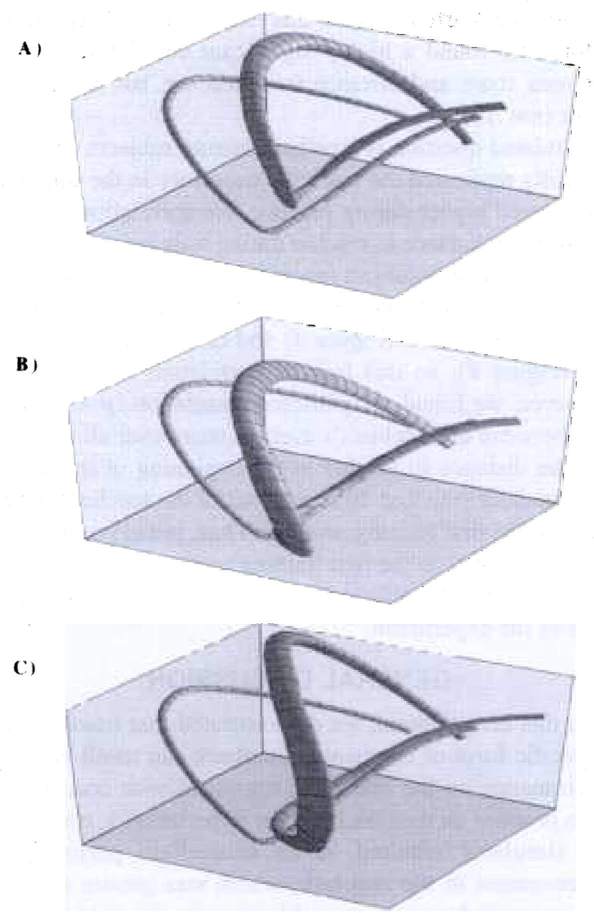

FIGURE 7. Individual trajectories in Experiment 2. Summarized trajectories for 1 training subject executing the real task on Day 1 (**A**), Day 2 (**B**), and Day 3 (**C**), superimposed on the teacher's trajectory.

isolate. Therefore, we now focus on Experiment 2, in which the task was considerably more difficult, allowing for visually different types of successful movements.

Figure 6 shows the pooled trajectories over all six sessions for the training group in the simulator (A) and on the real task (B), and for the control group (C). It is immediately obvious that the movements of the training subjects were much closer to the teacher's than the controls' movements were. Control subjects used movements (Figure 6C) very similar to those observed in Experiment 1 (Figure 5B): a short, straight swing of the paddle backward and then forward. The variability decreased around the point of impact (because we pooled trajectories, thickness was a combination of between- and within-subject variabilities). Notice that during training in the simulator, most of the variability was in depth, which is to be expected, given that subjects had some problems with depth perception. In the other two dimensions, which were parallel to the plane of projection in the simulator, subjects were more consistent during training than when playing on the real table.

Even though the training group as a whole executed movements similar to the teacher's trajectory, we observed that individual subjects modified the teacher's movement in subject-specific ways. As an example, Figure 7 shows the evolution of the movements of 1 training subject over the 3 days of the experiment. The general shape was similar to the teacher's, particularly during the forward swing, but an extra twist in depth was introduced, and it became more pronounced with time. Other subjects deviated consistently in the other two dimensions, indicating that these subject-specific modifications were not caused by poor perception of the teacher's movement in depth. In both experiments, all subjects were surprisingly consistent (even over days in Experiment 2), to the extent that it was possible to distinguish visually the trajectories of almost any 2 subjects.

To perform a quantitative analysis, we needed a numeric measure of how different two movements are. Rather than using the similarity score computed during training, which relied on the very close spatial and temporal alignment of the subject's and teacher's movements in the simulator, we defined a simpler measure of distance between two trajectories: the Euclidean distance after a translation aligning the centers of mass. For each movement of a subject, we computed the distance to the teacher's trajectory and the distance to the average trajectory of that subject over the whole session. We used the latter as a measure of consistency during each session.

Figure 8 displays distances averaged over days for the two groups during practice and for the training group in the simulator. As Figure 6 has already demonstrated, the movements of the training group during practice were significantly closer to those of the teacher than the control group's was. The control group was more consistent over the whole experiment, as indicated by the distance to the average subject's trajectory. The consistency of the training group increased both in the simulator and on the real task. These

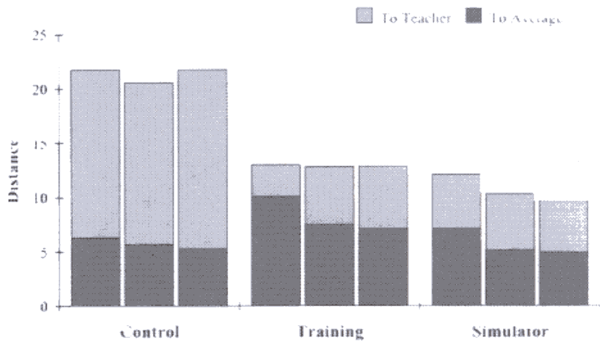

FIGURE 8. Distances to teacher's and subject's average trajectories in Experiment 2. The two distances were averaged over all trials of all subjects in each day and group (training group movements were separated into simulator and practice on the table).

consistency results were not unexpected, because the training subjects were trying not only to send the ball into the target but to do so by using a movement that was not what they would have normally tried first (as indicated by the average trajectory of the controls). More surprising was the change of the distance to the teacher's trajectory over time: Movements executed in the simulator were getting closer to the teacher's, but practice movements on the real task were not. The former effect was even more pronounced in the average time (not shown here) needed to reach the prespecified similarity level in the simulator: On the last 2 days, it was almost 3 times shorter than on Day 1. This observation suggests that the general shape of the teacher's trajectory (which is what our distance measure was most sensitive to) transferred to the real task on the very 1st day. Visual inspection of trajectory averages for each session (not shown here) confirmed this hypothesis: The trajectories executed during practice did not become more similar to the teacher's with time.

An interesting question is whether subject's trajectories changed systematically within each session. In Figure 9, the distances to the teacher's trajectory for the training group averaged over the first 5 and the last 55 trials of practice are compared. The trajectories were closer to the teacher's, immediately after training in the simulator, at the beginning of each practice session than later in the session. This finding indicates a gradual transition between the precise replication of the teacher's trajectory in the simulator and its slightly modified, subject-specific version used on the real task.

Given the indications that the distance measure defined on the whole movement did not capture all of the task-related invariants, we should not expect to find a strong relationship between distance and score on individual trials. In fact, by pooling together all trials of the training group and considering partial scores that were counted as misses before, we found a highly significant correlation ($p \sim 0$) between score and distance from teacher, but it was very weak ($r = .13$).

A related question is whether training subjects who successfully replicated the teacher's trajectory in the simulator also scored higher during practice. No correlation between the average distance to teacher during training and the average score for each subject (averages were taken over all sessions) was found. We already know that performance improved with time (Figure 4) and distance to teacher did not (Figure 8), so this lack of correlation was expected. However, we found a significant correlation ($p = .02$, $r = .74$) between each subject's average score over all sessions, and the distance to teacher at the beginning of training in the simulator (the first 30 repetitions of the teacher's movement on the first training session). Thus, initial performance in the simulator on the first training session was a good predictor of overall performance on the real task over the 3 days of the experiment.

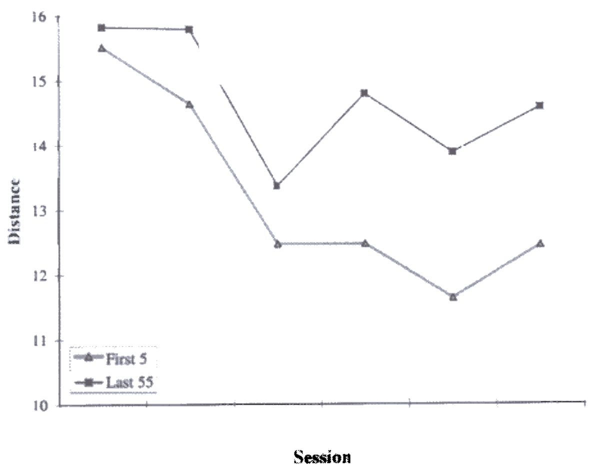

FIGURE 9. Distance from teacher for training group in Experiment 2. The first 5 and the last 55 trials of all subjects were averaged for each practice session.

GENERAL DISCUSSION

In this investigation, we demonstrated that training with a specific form of augmented feedback can result in better performance on the real task, compared with coaching or extra practice on the task itself. In Experiment 1, training in the simulator resulted in an immediate performance improvement in the real task, which was greater than the improvement from a comparable amount of coaching of the control group. More important, this initial experiment showed that the effectiveness of the augmented feedback disappears if the temporal frame of reference is removed from the training scheme: When the subjects were merely following the teacher's trajectory, they could not time their movements with respect to the motion of the ball in the real task. Indeed, compared with the control condition, training in the pilot simulator actually led to poorer performance in the real task. When the motion of the ball was added to the simulator, however, there was a significant improvement in performance of the simulator-trained subjects, compared with coached controls.

This finding supports the hypothesis that there is a significant difference between learning open (variable environment) versus closed (fixed environment) skills (Gentile, 1972). Our results do not agree with those of Schmidt (1975), who argued that subjects can learn a "schema" describing a movement in a fixed environment and then simply use KR to determine when that schema should be executed. In our study, teaching the temporal coupling between the flight of the ball and the execution of the swing was essential to the success of the training method.

We next considered a more difficult task, in Experiment 2, and found that the group receiving augmented feedback performed better than the controls, even though the controls received 50% more practice in the real task. The improvements persisted over 3 days, whereas the amount of training necessary so that a prespecified performance in the simulator could be achieved decreased with time.

Our kinematic analysis of the motion of the subjects' paddles suggests that the subjects learned the task-related invariants of the teacher's movement as presented in the virtual environment. Although such invariants were not enough to specify a unique movement, missing details were filled in, resulting in subject-specific trajectories. That filling in occurred only on the real task, however, because during training subjects were instructed to replicate the teacher. Furthermore, we observed that this filling in occurred gradually in every practice session. Because the distance measure used for kinematic analysis was not tuned to the task-related invariants, it became dominated by the filled-in details and thus failed to detect any change resulting from learning (which is presumably what happened in Experiment 1); or it detected the most obvious invariants, which transferred to the real task almost immediately (as in Experiment 2), and did not indicate any further learning, even though actual performance kept improving. Such a distance measure could correspond to actual performance only during the initial period of training when subjects learn the few invariants to which it is sensitive. We found that even this short period of sensitivity was enough for us to predict to some extent the overall performance on the real task for each subject.

An alternative hypothesis would be that only the general shape of the teacher's trajectory transfers from training in the simulator, and once it is learned, performance keeps improving because additional adjustments are made through practice only. This can account for all observations in Experiment 2 but is directly contradicted by the observations in Experiment 1, where we found a significant difference in performance for the same general shape of the subjects' trajectories.

Our main motivation for our experiments was to convincingly demonstrate an advantage of a specific augmented-feedback training procedure, compared with extra practice or coaching—something that has not been demonstrated so far on difficult multijoint movements. We attempted to design the optimal training scheme for the particular task by considering many possible characteristics of augmented feedback and choosing the ones that seemed most appropriate. As a result, we found significant performance benefits but did not prove that any one of the choices we made was necessary. Below we briefly summarize the important alternatives to our training scheme, which could be studied in future work:

1. Rather than focusing on the end-effector kinematics, one could visualize the movement of the arm or the entire body; note that it may be technically impossible to present all that information while the subject and reference movements are superimposed in the same frame of reference.

2. Instead of being superimposed in the same frame, the two movements can be shown on separate monitors; this will make visual comparison significantly more difficult.

3. The subject could execute the movement after, and not simultaneously with, the teacher; in this case, the comparison with the teacher's movement would rely either on memory or some abstract distance measure.

4. One can vary the reference movement, and even include as reference movements some of the successful movements of the subject; although variability can improve generalization, learning to replicate a new reference movement every few trials might have a negative effect.

5. It is remotely possible that the subject can learn entirely through observation and that executing the movement during training is not necessary; we feel that although observation may improve performance relative to a baseline, it is unlikely to result in better performance than actual practice on the task.

As we emphasized above, virtual reality was only the tool for implementing the training scheme we chose. It may be argued that we were confounding the type of training (the points discussed above) with the mode of training (virtual reality vs. a human model or a video replay). However, the distinction between type and mode of training is rather blurred when certain types of training can be implemented in only one mode of training. This was exactly the case here: Virtual reality was the only mode of training that allowed us to isolate the movement of the expert's end-effector, superimposing it on the subject's end-effector and displaying the two movements concurrently.

Although the features of the augmented feedback procedure that contribute to learning need to be clarified, we feel that the primary question to be addressed is: What exactly does the motor system learn in a difficult, nonisometric motor task? Unless more insight into this issue is gained, attempts to compare the details of complicated augmented feedback procedures may fail to provide general answers. One possibility, which we proposed above, is that learning a nonisometric task can be subdivided into learning (a) the constraints that any successful movement must satisfy and (b) the particular movement satisfying those constraints that is easiest to execute and control. To address these important issues in future work, we need to develop nonisometric laboratory tasks that preserve most of the complexity of the task studied here yet involve a simpler environment that can be experimentally manipulated.

REFERENCES

Atkenson, C. G., & Hollerbach, J. M. (1985). Kinematic features of unrestrained vertical arm movements. *Journal of Neuroscience, 5,* 2318–2330.

Bilodeau, E. A., & Bilodeau, I. McD. (1962). Motor skills learning. *Annual Review of Psychology, 13,* 243–280.

Bilodeau, I. McD., & Rosenquist, H. S. (1964). Supplementary feedback in rotary-pursuit tracking. *Journal of Experimental Psychology, 64,* 53–57.

Carroll, W. R., & Bandura, A. (1987). Translating cognition into action: The role of visual guidance in observational learning. *Journal of Motor Behavior, 19,* 385–398.

Carroll, W. R., & Bandura, A. (1990). Representational guidance of action production in observational learning: A causal analysis. *Journal of Motor Behavior, 22,* 85–97.

Cote, D. O., Williges, B. H., & Williges, R. C. (1981). Augmented feedback in adaptive motor skill training. *Human Factors, 23*(4), 505–508.

Flash, T., & Gurevich I. (1992). Arm movement and stiffness adaptation to external loads. *Proceedings of the IEEE Engineering Medical Biology Conference, Orlando FL, 13,* 885–886.

Flash, T., & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience, 5*(7), 1688–1703.

Gentile, A. M. (1972). A working model of skill acquisition with relation to teaching. *Quest, 17,* 3–23.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics, 14,* 201–211.

Karlin, L., & Mortimer, R. G. (1963). Effects of verbal, visual and auditory augmented cues on learning a complex motor skill. *Journal of Experimental Psychology, 65,* 75–79.

Kernodle, M. W., & Carlton, L. G. (1992). Information feedback and the learning of multiple-degree-of-freedom activities. *Journal of Motor Behavior, 24,* 187–196.

Kozak, J. J., Hancock, P. A., Arthur, E. J., & Chrysler, S. T. (1993). Transfer of training from virtual reality. *Ergonomics, 36*(7), 777–784.

Lacquaniti, F., Soechting, J. F., & Terzuolo, C. A. (1982). Some factors pertinent to the organization and control of arm movements. *Brain Research, 252,* 393–397.

Lintern, G. (1991). An informational perspective on skill transfer in human-machine systems. *Human Factors, 33*(3), 251–266.

Lintern, G., & Roscoe, S. N. (1980). Visual cue augmentation in contact flight simulation. In S. N. Roscoe (Ed.), *Aviation psychology* (pp. 227–238). Ames, IA: Iowa State University Press.

Lintern, G., Roscoe, S. N., Koonce, J. M., & Segal, L. D. (1990). Display principles, control dynamics, and environmental factors in pilot training and transfer. *Human Factors, 32*(3), 299–317.

Marsden, C. D. (1982). The mysterious motor function of the basal ganglia: The Robert Wartenberg lecture. *Neurology (NY), 32,* 514–539.

Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research, 42,* 223–227.

Newell, K. M., & Carlton, M. J. (1987). Augmented information and the acquisition of isometric tasks. *Journal of Motor Behavior, 19,* 4–12.

Newell, K. M., & Carlton, M. J. (1990). The interaction of criterion and feedback information in learning a drawing task. *Journal of Motor Behavior, 22,* 536–552.

Ruitenbeek, J. C. (1984). Invariants in loaded goal directed movements. *Biological Cybernetics, 51,* 11–20.

Schmidt, R. A. (1975). A schema theory of discrete motor skill learning. *Psychological Review, 82,* 225–260.

Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience, 14,* 3208–3224.

Thach, W. T., Goodkin, H. P., & Keating, J. G. (1992). The cerebellum and the adaptive coordination of movement. *Annual Review of Neuroscience, 15,* 403–442.

Vander Linden, D. W., Cauraugh, J. H., & Greene, T. A. (1993). The effect of frequency of kinetic feedback on learning an isometric force production task in nondisabled subjects. *Physical Therapy, 73,* 79–87.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Experimental Brain Research , 103*(3), 460–470..

Young, D. E., & Schmidt, R. A. (1992). Augmented kinematic feedback for motor learning. *Journal of Motor Behavior, 24,* 261–273.

Submitted August 7, 1995
Revised June 10, 1996