Journal of Experimental Psychology:
Human Perception and Performance
1995, Vol. 21, No. 5, 1174-1198

Copyright 1995 by the American Psychological Association, Inc.
0096-1523/95/$3.00

Internal Representations of the Motor Apparatus: Implications From Generalization in Visuomotor Learning

Hiroshi Imamizu, Yoji Uno, and Mitsuo Kawato
ATR Human Information Processing Research Laboratories

Recent computational studies have proposed that the motor system acquires internal models of kinematic transformations, dynamic transformations, or both by learning. Computationally, internal models can be characterized by 2 extreme representations: structured and tabular (C. G. Atkeson, 1989). Tabular models do not need prior knowledge about the structure of the motor apparatus, but they lack the capability to generalize learned movements. Structured models, on the other hand, can generalize learned movements, but they require an analytical description of the motor apparatus. In investigating humans' capacity to generalize kinematic transformations, we examined which type of representation humans' motor system might use. Results suggest that internal representations are nonstructured and nontabular. Findings may be due to a neural network model with a medium number of neurons and synapses.

The process whereby experience in one activity leads to improved performance in another is referred to as *generalization.* For example, it takes a long time for a beginner in table tennis to learn how to return the ball precisely to the corner of one side of the opposite court. However, after this skill is learned successfully, less time is needed to return the ball to any spot on the other side. Human motor learning is characterized by its great generalization ability, in which the experience gained for a particular posture or movement can improve a subsequent posture or movement. In this research, we studied generalization ability in visuomotor learning (i.e., learning an aiming task under kinematically transformed visual feedback) to understand how the internal model of kinematic transformations is represented in the human central nervous system (CNS).

Internal Models of Motor Apparatus

For an individual to grasp a cup placed in front of him or her, the CNS must solve several problems. First, it must select the best trajectory, out of an infinite number of possible trajectories, from the hand to the cup. These trajectories are thought to be planned in visuospatial coordinates because, under normal circumstances, information about the position of the cup is obtained visually. Second, visuospatial coordinates of the desired trajectory must be converted into arm–joint angles. This transformation is called *inverse kinematics* in robotics. Accordingly, the task to choose the appropriate joint angles given a desired hand position is called an *inverse kinematics problem.* Third, motor commands must be generated to coordinate the activities of many muscles so that the desired trajectory is realized. To drive the limb along the desired trajectory, the appropriate torque must be fed to the joint. The transformation from a desired pattern of motion to actuator commands is referred to as *inverse dynamics.* Thus, at least three major problems must be resolved: trajectory planning, inverse kinematics, and inverse dynamics (Hollerbach, 1982; Kawato, Furukawa, & Suzuki, 1987; Saltzman, 1979, 1987).

How does the human motor control system solve these problems so quickly? There are two general strategies for controlling the motor apparatus: feedback control and feedforward control. For example, although a line is carefully being traced with a pencil, the CNS continuously uses proprioceptive feedback from the limbs and visual feedback of the hand position. However, movements cannot be rapid and smooth by feedback control because there are substantial delays in the feedback loop (about 200 ms for visual feedback and 100 ms for proprioceptive feedback; Keele, 1981). For fast arm movements in the range of 500–600 ms, these loop delays are too long to serve the role of an efficient feedback controller (Hollerbach, 1982). According to some computational theories of motor control, feedforward control is assumed to be achieved by building and refining internal models of kinematic transformations, dynamic transformations, or both (Atkeson, 1989; Kawato & Gomi, 1992). Once the internal models are acquired, feedforward control can be executed, and a rapid smooth motion satisfying the aforementioned goal can be generated.

Apart from the context of modern robotics and neural networks, two approaches to motor skill representation have been proposed by psychologists. One is the *schema abstraction model* and the other is the *specific exemplar model* (Koh & Meyer, 1991). Schmidt (1975) proposed the schema theory of sensorimotor learning, in which four kinds of

Hiroshi Imamizu, Yoji Uno, and Mitsuo Kawato, ATR Human Information Processing Research Laboratories, Kyoto, Japan.

We thank Shinsuke Shimojo, Stefan Schaal, Christopher G. Atkeson, and David A. Rosenbaum for their helpful suggestions. We are also grateful to Frank E. Pollick for developing the method for detecting the end point of the first ballistic movement.

Correspondence concerning this article should be addressed to Hiroshi Imamizu, ATR Human Information Processing Research Laboratories, 2-2, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-02, Japan. Electronic mail may be sent via Internet to imamizu@hip.atr.co.jp.

information are available after a movement: the movement outcome, the response specification, the initial conditions, and the proprioceptive signal. Sensorimotor learning involves changing a "recall schema" in such a way that the movement outcome and initial conditions can be used in future trials to retrieve the corresponding response specifications. The basic idea behind this model is that a central tendency (the prototype) is abstracted from the experience gained with a number of examples for a specific category. On the other hand, the exemplar model maintains that specific exemplars are stored in memory. It determines the output signal for a novel exemplar calculating the distance between the novel exemplar and a stored exemplar (e.g., Chamberlin & Magill, 1992).

Methods for Representing Internal Models

Many researchers of robotics or neural networks have proposed methods to implement inverse kinematics models, inverse dynamics models, or both. Among them, two extreme approaches have been proposed (Atkeson, 1989; Hollerbach, 1982). One approach is to build an idealized physical model of the structure of the motor apparatus by estimating physical parameters. The basic ideas behind this model are common to that of the schema theory in that the central tendency is abstracted from a number of examples and that what is learned is not some particular movement but the capacity to generalize (Schmidt, 1988). The other approach is to construct a tabular representation with no analytical description of the motor apparatus. On the basis of a thorough review by Atkeson (1989), we briefly explain the differences between these two approaches. Of course, it is impossible to classify all methods into these two extreme groups, and several "hybrid," or "intermediate," approaches have been proposed. Among them, some recent approaches using multilayer feedforward neural network models have greatly broadened the range of possible representations of transformations such as kinematic and dynamic transformations (e.g., Jordan, 1990; Kawato, 1990).

Structured Representation

In a reaching movement, the CNS has to transform visual information regarding the target position into a motor command to move the hand to the target. To illustrate this transformation, suppose that human participants move a cursor on a cathode ray tube (CRT) screen using a computer mouse (see Figure 1). The kinematic transformation comprises two successive transformations in this case: from CRT coordinates to hand coordinates and from hand coordinates to joint coordinates. The CRT and hand coordinates are Cartesian, and the first transformation can be represented as

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \cos a & \sin a \\ -\sin a & \cos a \end{pmatrix}^{-1} \begin{pmatrix} X - b_1 \\ Y - b_2 \end{pmatrix} \qquad (1)$$

where the target position is denoted by (X, Y) in the CRT coordinates and by (x, y) in the hand coordinates. a is the rotation angle and (b_1, b_2) is the translation distance between the two Cartesian coordinates.

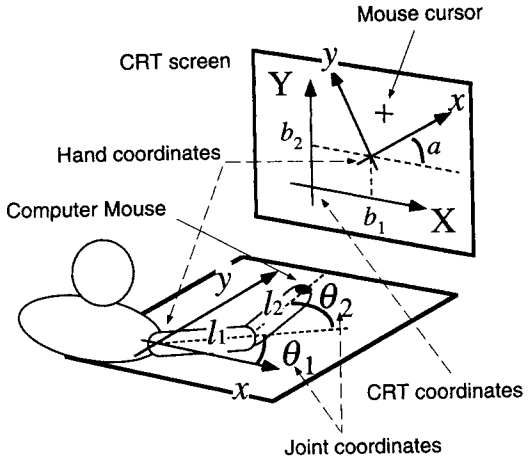

Figure 1. When we move the cursor on a CRT screen using a computer mouse, the kinematics transformation comprises two successive transformations: from CRT coordinates to hand coordinates and from hand coordinates to joint coordinates.

Suppose the arm is composed of two rigid links of length l_1 and l_2. The links rotate about vertical parallel axes that are fixed with respect to the links. The shoulder joint has a single degree of freedom measured by angle θ_1 and the elbow joint has a single degree of freedom measured by angle θ_2. The second transformation for this model can be represented as

$$\theta_2 = \arccos\left(\frac{x^2 + y^2 - l_1^2 - l_2^2}{2l_1 l_2}\right)$$

$$\theta_1 = \arctan\left(\frac{y}{x}\right) - \arctan\left(\frac{l_2 \sin(\theta_2)}{l_1 + l_2 \cos(\theta_2)}\right). \qquad (2)$$

Using these equations, the CRT coordinates of the target position (X, Y) can be converted into terms of the arm–joint angles (θ_1, θ_2). Thus, the kinematics are represented by Equations 1 and 2.

These equations include several physical parameters (i.e., a, b_1, b_2, l_1, and l_2). The representation of a transformation requiring real physical parameters is referred to as a *structured representation* or an *idealized physical model.* Unfortunately, these equations are useless by themselves for controlling the arm unless the physical parameters are estimated. Thus, the problem of building kinematic and dynamic transformations by using structured representations is how to estimate the parameters of each transformation and the motor apparatus. In the idealized case, only five parameters need to be estimated. For a practical manipulator control problem, however, the parameter estimation procedures are often complex because of the large number of parameters and the nonlinearity of the problem. Concerning the inverse dynamics transformation, An, Atkeson and Hollerbach (1988) devised a method to estimate the inertial terms in the equations representing the inverse dynamics of

a robot arm. They showed that their method can do estimation in only a small number of trials. The advantage of structured representation is that if the parameters are estimated correctly and the underlying model structure is correct, the model will be correct for all other trajectories. Thus, the generalization becomes global because the structured representations support the mappings between the input and output signals in the whole work space.

Tabular Representation

In the simplest version of a tabular representation, a kinematic transformation can be solved by constructing a table that represents the mappings between the input (e.g., desired position of the hand in the CRT coordinates) and output (e.g., arm–joint angles) signals. This tabular representation does not need prior knowledge about the structure of the motor apparatus. It computes the transformation by referring to the table rather than by solving analytic problems. Accordingly, it can avoid computational delays during real-time control. However, problems arise when simple tabular representations are applied to complex systems because the size of the table grows exponentially with the number of input variables. When the number of input variables is n and each variable takes m different values, m^n entries are required in the table. Another problem of tabular representations is their inability to generalize learned movements. Because the mappings between the input and output signals are independently determined, the experience gained for a particular posture or trajectory cannot improve the ability to attain a different posture or trajectory. Instead, generalization is restricted to a local area where a similar input leads to a similar output.

Albus (1975) proposed a more sophisticated tabular organization to represent the internal models of inverse kinematics and dynamics transformations: the cerebellar model articulation controller. In this model, a table entry is represented in a distributed fashion by a group of weights. The table size is reduced as a result of the distributed memory. The weights in the table are updated on the basis of observed input–output pairs. Because the neighboring entries of the table share common weights, the local generalization of learned movements is possible.

Connectionist Representation

This approach is inspired by the architecture of the nervous system. In this approach, the mappings between the input and output signals are achieved by a set of nonlinear computational units that have neuronlike functions (Rumelhart, 1986). This model does not directly represent mathematical formulae like Equations 1 and 2, but the kinematics parameters are implicitly represented by synaptic weights $(w_1 \ldots w_n)$. Thus, the inverse kinematics can be represented as

$$\theta_1 = f(w_1 \cdots w_n, X, Y)$$
$$\theta_2 = g(w_1 \cdots w_n, X, Y) \quad (3)$$

The functional forms of f and g depend on the precise architecture of the connectionist (neural network) model, such as whether it is a simple perceptron, a multilayer perceptron, a modular-network architecture, or a Boltzman machine. Although this approach uses parameters (synaptic weights), these parameters do not correspond to physical parameters. On this point, the connectionist representation is different from the structured representation, and it is sometimes called a "black box model." Furthermore, this approach is nontabular because it is not implemented by the local association between input and output signals. Some recent approaches have proposed modeling of kinematic and dynamic transformations using a connectionist representation (Jordan, 1990; Kawato, 1990).

Logic for Revealing Internal Representation From Different Generalizations

Suppose that people first learn to aim at a particular target in their work space and that after intensive training they are tested on their ability to aim at near and far targets in various locations. The logic for revealing the internal representation of the kinematic transformation is as follows: On the one hand, if the type of representation the human motor system uses is tabular, the effect of learning should be observed only on aiming at targets near the location of the target in the training session (local generalization). On the other hand, if the human motor system uses structured representations, the effect of learning should be observed on aiming at various targets regardless of the location of the target used in the training session (global generalization). Accordingly, researchers should be able to identify the type of representation by investigating the aiming behavior of human participants at various target locations before and after intensive training in aiming at a specific location.

The logic described earlier assumes the simplest cases in which the internal representation type is extremely structured or tabular. If the type is not so extreme or if it is a connectionist representation, intermediate generalizations should be observed.

Research Goals

It is difficult to examine the effect of learning by studying the performance of human participants under normal conditions because they have already been intensively trained in daily life in aiming at various locations in their work space. For this reason, we kinematically transformed visual feedback and asked participants to learn an aiming task under this condition. In this experiment, we studied the representation used in an internal model of kinematic transformation. Although it might be possible to investigate the representation used for a dynamic model using similar methods, our research did not focus on dynamics.

Shadmehr and Mussa-Ivaldi (1994) investigated the generalization of motor learning in a novel dynamical environment. Participants made reaching movements while holding the end-effector of a manipulandum. The experimenter

modified the dynamic properties (force field) of the manipulandum by controlling servomotors attached to it and trained the respondents in this transformed condition. There were two work spaces (i.e., right and left) in their experiment. The participants were trained in the right work space and tested in the left work space. It was found that the effect of learning was present in the left work space beyond the boundary of the trained region.

Our transformation of visual feedback followed that of Cunningham (1989). She used an aiming task under rotational visual feedback of various angles ranging from 0° to 180° and found that the aiming task was relatively easy when the rotations were 0° and 180° and more difficult when the rotations were between 90° and 120°. The difficulty of the task increased from 0° to 90° and decreased from 120° to 180°. Our pilot observations suggest that 75° rotations are relatively difficult for participants and that, through practice, they can dramatically improve performance within a few hours. We therefore adopted an aiming task under the condition of 75° rotation of visual feedback.

We investigated generalization in visuomotor learning on the basis of our predictions for different representations while using two different constraints and measurements. In Experiment 1, the task for the participant was to acquire the target as rapidly as possible. We measured the performance time (PT) from the onset of the start signal to the time when the cursor arrived at the target as an indicator of the learning effect. There were no other time constraints on the participants' performance except the instruction to acquire the target as rapidly as possible. Under these conditions, there were two phases in aiming performance: ballistic and corrective.[1] In Experiment 2, the task was to acquire the target within a short time interval (600 ms) after the start signal (a tone cue). In this case, the ballistic phase was more prominent than the corrective phase. We measured the distance between the target and the end position of the ballistic movement as an indicator of the learning effect.

In both experiments, we adopted the pretest–posttest paradigm to assess the generalization. Each experiment consisted of three sessions. The first session was the pretest, in which test targets appeared in various locations. In the second (training) session, the targets appeared in the restricted small region. The third session was the posttest, in which the test targets appeared in the same manner as the pretest. We compared the participants' performance in the pretest with that in the posttest to investigate the generalization of the training effect.

Experiment 1

Method

Participants

The participants were 5 right-handed undergraduate students from Nara Women's University or Doshisha University. They were naive as to the purpose of this experiment and were paid for their participation.

Apparatus and Stimuli

For both Experiments 1 and 2, the aiming task used a 33-in (83.82-cm) CRT monitor and a position measurement system (Northern Digital Optotrak) connected to a personal computer (Toshiba J3100). Figure 2A illustrates the arrangement of the apparatus. The participants were instructed to move their right hand above the board placed horizontally in front of them. The hand lightly touched the board during the movement so that the movement path was always within the horizontal plane. A marker (an infrared light-emitting diode) was attached to the top of the participant's hand, and its position was sampled at 100 Hz and stored in the personal computer. An occluder was placed above the board and the participant's hand to avoid direct visual feedback. The CRT's display surface was also placed nearly horizontal in front of the participant (in fact, slanted 15° from the horizontal plane toward the participant so that it was easier to observe). The room was almost completely dark.

The background color of the CRT screen was black. The start position was always located at the center of the screen, whereas the target position was located around the center, as described later. The targets were white circles 5 mm in diameter. The CRT screen displayed a movable cursor as a white circle 0.9 mm in diameter. Under normal conditions, when the participants moved their hand above the board, the cursor moved in the same direction as the hand movement. However, during both experiments, the position of the cursor was constantly rotated by 75° around the center of the CRT screen. In other words, as indicated in Figure 2B, the angular discrepancy between the moving direction of the participant's hand and that of the cursor on the screen was always 75°. We informed the participants of this before starting the experiment. The displacement gain of this transformation was 1.0 (i.e., moving the hand by X cm caused the cursor to move X cm on the CRT). The delay between the movement of the hand and that of the cursor (i.e., time for detecting the marker position and displaying it as the cursor on the CRT) was shorter than 20 ms.

As described in the introduction, we examined spatiotemporal characteristics of the participants' aiming behaviors at various target locations before and after intensive training on targets in a specified small region. To accommodate this scheme, there were two sets of targets: test targets (used in the pre- and posttests) and training targets (used in the training session).

Test targets. As shown in Figure 3, the center of each test target was located on the circumference of a circle whose center was identical to the center of the CRT screen (i.e., the starting position) and whose radius was 155 mm. They were placed every 45°, as shown in the left side of Figure 3. They were labeled T1, T2, T3, and so on, as indicated in the figure.

Training targets. In principle, each participant should have been trained on only one test target during the training session. However, training on only one target would not have met the

[1] Woodworth (1899, cited in Flowers, 1975) made a distinction between the components for voluntary movement in the aiming task: an "initial impulse phase" and a series of "secondary adjustments" made subsequently to attain the final target position. The first component is a fast, preprogrammed "ballistic" movement that brings the hand into the general area of the target. The second component comprises a number of adjustments. In this latter phase, movements are continuously monitored and adjusted according to sensory information. The first component, the "initial impulse phase" in Woodworth's terminology, can be called a "ballistic" movement and the second "current control" component can be called a "corrective" movement (Flowers, 1975).

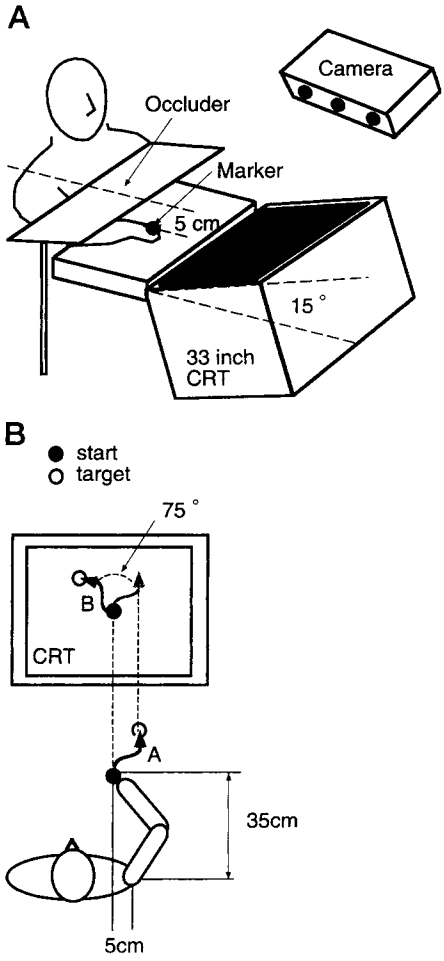

Figure 2. A: Experimental setup. B: The rearrangement of visual feedback by 75° counterclockwise rotation. The thick arrow (A) indicates the path in which the participant moved his or her hand and the arrow (B) indicates the path of the resulting cursor movement on the CRT display.

purposes of our experiment. More specifically, the participant could learn the aiming task without internal models of the kinematic transformation by simply remembering the proprioceptive feedback for a particular arm posture, but we wanted to investigate the representations used in the internal models of the kinematic transformation. Second, the values of the parameters in Equations 1 and 2 could not be determined. Because these equations include five unknown parameters, at least five different target positions (five pairs of input and output values) are required to determine the parameters. Finally, although the transformation was rotational in this experiment, it could not be distinguished from reflection on the axis whose direction diverged 37.5° from the target direction.

To avoid these problems, the participants were trained on nine targets in the small region around each test target in the manner described later. These targets were named training targets. The right side of Figure 3 illustrates the training targets around T1. The direction of the training targets was 0°, 11.25°, or −11.25°, and the distance from each of the test targets was 0, 5, or −5 mm. They were labeled *S1*, *S2*, *S3*, and so on, as indicated in the figure. S5 was always the same as each test target.

Procedure

Before the beginning of the experiment, the transformation and the task were described to the participants using a figure identical to Figure 2B as follows:

> A movable cursor represented by a small filled circle and a target represented by an open circle will appear on this screen. The position of the cursor will be constantly rotated by 75° around the center of the screen. In other words, as indicated in this figure, the angular discrepancy between the moving direction of your hand and that of the cursor on the screen will always be 75°. The ratio of hand movement to cursor movement on the screen is 1:1. The task is to move the cursor to the target as rapidly as possible after the tone cue generated by the computer. It may be difficult at first but you will find that the task will become easier with practice.

Each trial consisted of the following sequence: The participant's hand was placed at the starting position (i.e., the center of the board) by the experimenter. After a few seconds, the cursor appeared at the starting position (i.e., the center of the CRT screen). At the same time, the target also appeared in one of the possible locations. The participants observed the screen for 2 s and were not allowed to move the cursor during this period. A tone cue (click) was then generated by the computer to signal the beginning of the trial. The participant's task was to acquire the target as rapidly as possible. The trial was terminated when the cursor arrived at the target circle and stayed within it for more than 100 ms. Another tone cue signaled the end of the trial. At this same time, the target and the cursor disappeared from the screen and the trial was terminated. The time limit for arriving at the target was 8 s.

The experiment consisted of three sessions (pretest, training session, and posttest) and lasted about 2 hr. The time interval between sessions was shorter than 5 min. Four (T.F., X.O., H.K., and E.S.) of the 5 participants completed the three sessions as experimental subjects, whereas one (K.T.), as a control subject, did not participate in the training session and rested during it. The training direction was determined for each participant in the experimental group (T.F. = T1, X.O. = T7, H.K. = T1, and E.S. = T5) as shown in Figure 4.

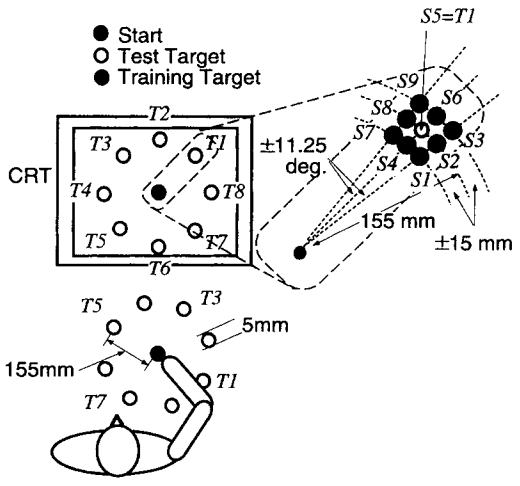

Figure 3. The left side of the figure shows the locations of the start point and test targets on the CRT screen. The right side shows the set of training targets for one of the test targets (T1).

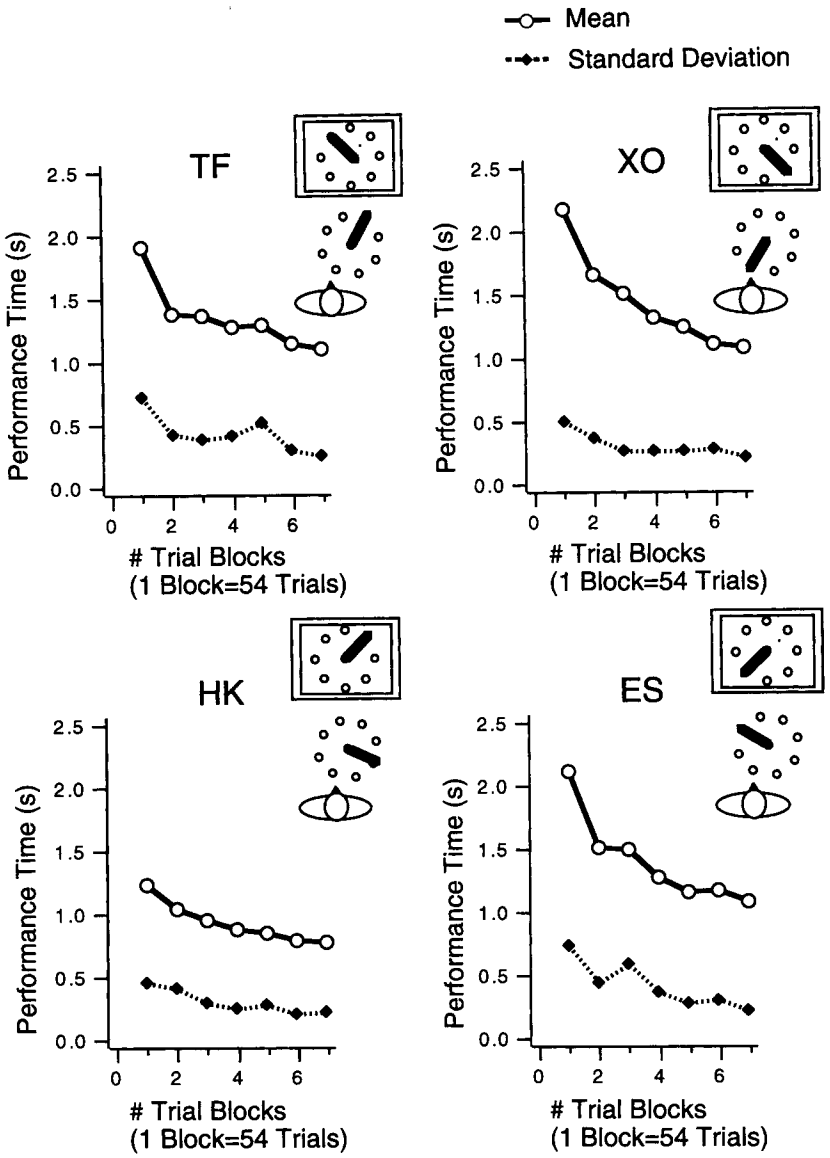

Figure 4. Learning profiles of 4 participants in the experimental group. The "clock" attached to each profile illustrates the training direction both in the hand space and on the CRT screen.

The pretest and the posttest each consisted of 56 trials and lasted about 20 min. In each session, one of eight test targets appeared seven times in random order. The time interval between trials was shorter than 5 s. The size of these two sessions (56 trials) was chosen as a compromise between two pragmatic constraints: On the one hand, it should be brief enough to minimize learning within these sessions; on the other hand, it should be long enough to eliminate noise factors and to reliably assess the participant's performance.

The training session consisted of seven blocks of trials, and each block consisted of 54 trials. In each block, one of nine training targets appeared six times in random order. A 5-min rest period was allowed between blocks. The training session lasted about 90 min.

Data Analysis

The performance time (PT) (from the onset of the start signal to the time when the cursor arrived at the target) was measured as an indicator of learning.[2] The PT consisted of the reaction time (from the onset of the start signal to the beginning of the participant's

[2] The data of the trials in which the PT exceeded 8 s (i.e., the recording time limit of our data acquisition system) were not included in this analysis. The maximum number of such trials for each target was three in the case of the pretest session, whereas there were no such trials in the case of the training and posttest sessions for any participant.

movement) and the movement time (from the beginning of the participant's movement to the time when the cursor arrived at the target).

The positions of the marker were sampled and stored at 100 Hz. The spatial precision of the position measurement system was 0.1 mm. The tangential velocities were computed by the fifth-order central difference. The computation for this difference was based on LaGrange's interpolation polynomial and is defined as

$$z_i = \frac{z_{i-2} - 8z_{i-1} + 8z_{i+1} - z_{i+2}}{12\Delta t},$$

where z_i is the ith data value, and Δt is the time interval between data values (10 ms; i.e., 100 Hz in our experiments).

The statistical analysis for participants in the experimental group was as follows: In the pretest and posttest sessions, we calculated the mean PT for each participant and obtained differential PTs by subtracting the mean PT from the original PT so that the mean of the differential PTs for each participant would be zero. Next, we conducted a two-way analysis of variance (ANOVA) on the differential PTs with the test target direction and the type of test (i.e., pretest or posttest) as variables. The test target direction was measured by the angular discrepancy from the training direction. The training direction was 0°. The next counterclockwise test target direction was 45°, and the next clockwise one was −45°. On the training session data, a two-way ANOVA was carried out with trial blocks and positions as variables after subtracting the mean PT from the original PTs of all training sessions for each participant. The statistical analysis for the control subject followed that of the experimental subjects.

Results and Discussion

Learning Process During the Training Sessions

Performance time. In Figure 4, the mean PTs and standard deviations in each block of trials are plotted separately for each of the 4 participants in the experimental group. The main effect of trial blocks was significant, $F(6, 1446) = 116.31$, $p < .0001$. We carried out post hoc comparisons (Tukey's honestly significant difference [HSD] multiple comparisons) between the trial blocks. They revealed significant differences for the first versus all other blocks and the last versus all other blocks except the sixth ($p < .05$). These results suggest that learning occurred during the training sessions.

Trajectories and velocity profiles. Parts A and B of Figure 5 show the typical trajectories of one participant (T.F.) in the first (Trials 1–9) and the last (Trials 370–378) stages of the training session. In the first stage, the trajectories are curved or bent at many points. This indicates that the participant made movement corrections many times using visual feedback. However, in the last stage of learning, the trajectories are almost linear and the recorded points are sparser. Parts C and D of Figure 5 show the velocity profiles of those trials. The peaks of the velocity profiles in the last stage of learning are much higher (about three times in the case of this participant) than those in the first stage. The profiles in the first stage have multiple peaks of low amplitude and long tails, whereas those in the last stage have a single-peak, bell-shaped appearance.

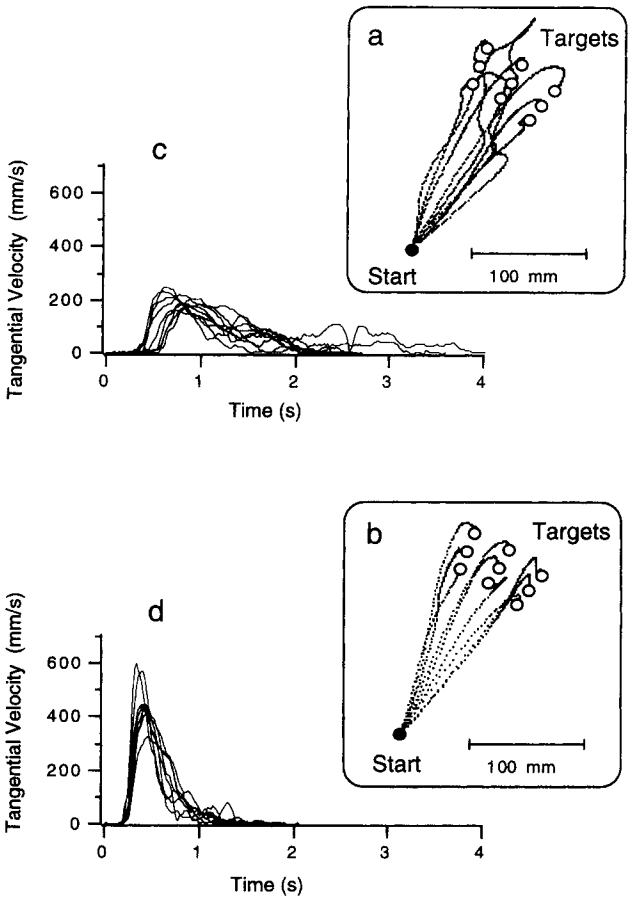

Figure 5. Typical trajectories of 1 participant (T.F.) in the first stage of the training session (Trials 1–9) superimposed on one graph (a). Those of the same participant in the last stage (Trials 370–378; b). Typical velocity profiles of the same participant in the first stage of the training session (c). Those of the same participant in the last stage (d).

It is well-known that movements tend to be performed more smoothly and gracefully after learning and practice (Georgopoulos, Kalaska, & Massey, 1981). According to Morrasso (1981), smooth movements between two points under normal conditions can be characterized by straight hand trajectories with single-peak, bell-shaped velocity profiles. Several feedforward control models of movement (e.g., Flash, 1987; Uno, Kawato, & Suzuki, 1989) have predicted trajectories of this type. In accordance with this, the modification of trajectories and their velocity profiles with practice has shown that at the beginning of a training session movement is mainly executed by feedback control, whereas at the end of the training session it is executed by feedforward control (i.e., ballistic movement).

Comparisons Between the Pretest and the Posttest

Performance time. In Figure 6, the mean PT is plotted as a function of test target position separately for each participant. The direction of the axis represents that of the target

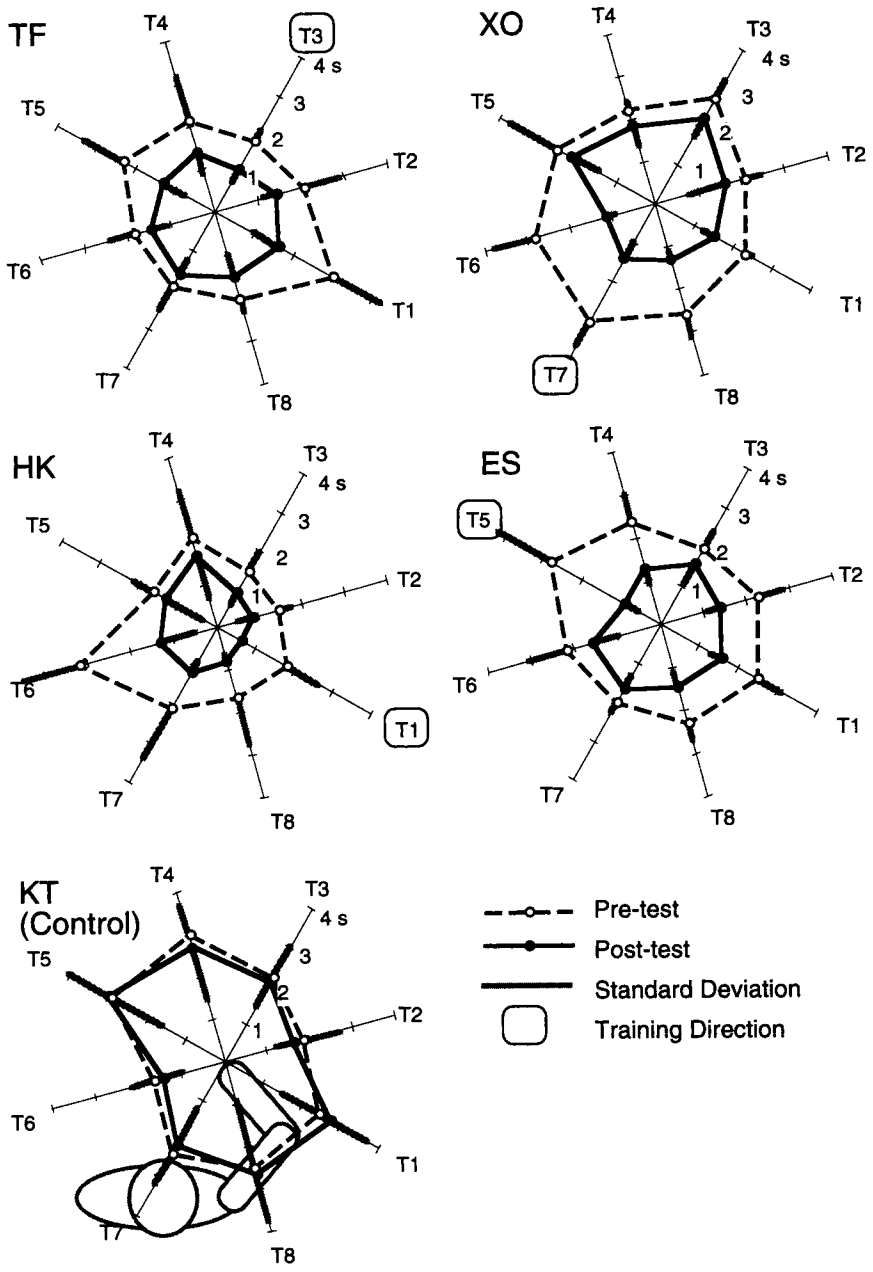

Figure 6. The mean performance times of the pretest and posttest sessions plotted as a function of test target direction in a radar chart for each participant. The direction of the axis represents that of the test target in the hand coordinates.

position in hand coordinates (not on the CRT screen). For each participant in the experimental group, the mean PTs in the posttest session (indicated by solid circles) are shorter than those in the pretest session (open circles), but for the control subject there is no consistent difference between them.

For participants in the experimental group, the main effect of the test type (i.e., pretest or posttest) was significant, $F(1, 413) = 145.29$, $p < .0001$; however, in the case of the control subject, the effect was not significant, $F(1, 96) = 0.28$, $p = .59$. The interaction between the test type and the test target direction (i.e., angular discrepancy from the training direction) was significant for the participants in the experimental group, $F(7, 413) = 3.55$, $p < .001$, but not for the control subject, $F(7, 96) = 0.084$, $p = .99$. These results suggest that, in the case of the experimental group participants, an effect of learning could be identified and that the effect differed depending on the angular discrepancy from the training direction. We compared the mean PTs of the pretest and posttest sessions for each angular discrepancy

Table 1
Significance Levels of Difference Between the Mean Performance Time of the Pretest and Posttest as a Function of Angular Discrepancy From the Training Direction (Tukey's HSD Multiple Comparisons)

Angular discrepancy							
−135°	−90°	−45°	0°	45°	90°	135°	180°
ns	*ns*	$p < .01$	$p < .001$	$p < .001$	$p < .01$	*ns*	*ns*

Note. HSD = honestly significant difference.

using Tukey's HSD post hoc comparisons to study the difference in the effect of learning. Table 1 shows a significant difference between the mean PT of the pretest and that of the posttest for each angular discrepancy from the training direction. The data in the table suggest that the effect was present in the majority of cases and that it was most prominent when the angular discrepancy was relatively small (i.e., when the test target was close to that used in the training session).

Trajectories and velocity profiles. Figure 7 shows trajectories and velocity profiles of one participant (T.F.) in the pretest and posttest sessions. The trajectories are plotted in the hand coordinates. The target used in the training session was T3, as indicated in Part B of Figure 7.

For the pretest session (see Part A of Figure 7), the trajectories are curved or bent at many points, and the velocity profiles have multiple peaks of low amplitude and long tails. These spatiotemporal characteristics of the participant's performance are almost the same as those in the first stage of the training session, as shown in Parts A and C of Figure 5. For the posttest session (see Part B of Figure 7), the trajectories become straighter and the recorded points are sparser. These modifications in the trajectories were most prominent when the participant aimed at the target that had been used in the training session. The velocity profiles developed a single-peak, bell-shaped appearance; the peaks of the velocity profiles were much higher than those in the pretest. These modifications in the velocity profiles were observed not only for the target used in the training session, but also for all other targets.

Maximum velocity. We further analyzed the spatiotemporal data by using the maximum velocity as an indicator of the learning effect. As one can observe from the velocity profiles in Figures 5 and 7, the maximum value is the peak of the first ballistic movement in almost every case. If the total trajectory lengths were constant, it could be assumed that if maximum velocity values went higher, the distance traveled during the ballistic phase would become longer, and the distance during the corrective phase would become shorter. As shown in Figure 5, the peaks of the velocity profiles in the last stage of learning were much higher than those in the first stage, indicating that the strategy of motor control became more and more feedforward. Therefore, the maximum velocity reflects how dominant the feedforward control of movement (i.e., ballistic movement) is in one trial of the aiming task.

As shown in Figure 7, the peaks of the velocity profiles in the posttest were higher than those in the pretest for the untrained directions as well as for the trained direction. We compared the peaks of the velocity profiles in the pretest with those in the posttest for all participants in the experimental group and directions. In Figure 8, the mean maximum velocity is plotted as a function of the test target direction. The maximum velocity of the posttest was higher than that of the pretest. Furthermore, the difference in mean maximum velocity between the post- and pretests was almost the same for all directions, independent of the training direction. We carried out a two-way ANOVA with the test target direction and type of test (pre- or posttest) as variables and subjects as an additional variable. The main effects of the target direction, $F(7, 413) = 700.25$, $p < .0001$, and type of test, $F(1, 413) = 3.89$, $p < .0005$, were significant. However, the interaction between them was not, $F(7, 413) = 1.43$, $p = .193$. The values of maximum velocity in the posttest were higher than those in the pretest.

On the basis of these results, one can suppose that the movement for all directions was implemented much more effectively by feedforward control independent of the training direction in the posttest. However, it is unknown whether the accuracy of the feedforward control increased. Thus, in Experiment 2, we modified the aiming task so as to make the first ballistic movement become dominant and to study the accuracy of the feedforward control more directly.

There was another reason for modifying the aiming task. As shown in Part C of Figure 5, the trajectories may be curved or bent at many points in the first stage of learning. This suggests that many directions of movement are contained in one trajectory. We expected the direction of movement to be fixed during the training session, but we could not reject the possibility of a participant experiencing various directions of movement, especially in the first stage of training. Therefore, we needed some constraints on the task to make the trajectories straight.

Experiment 2

Method

Participants

The participants were 4 right-handed undergraduate students of Doshisha University. None of the participants participated in Experiment 1. They were naive about the purpose of the study and were paid for their participation.

Apparatus and Stimuli

The same equipment used in Experiment 1 was used in Experiment 2.

Design and Procedure

The design followed that of Experiment 1, but the procedure of each trial and the instructions to the participants were modified.

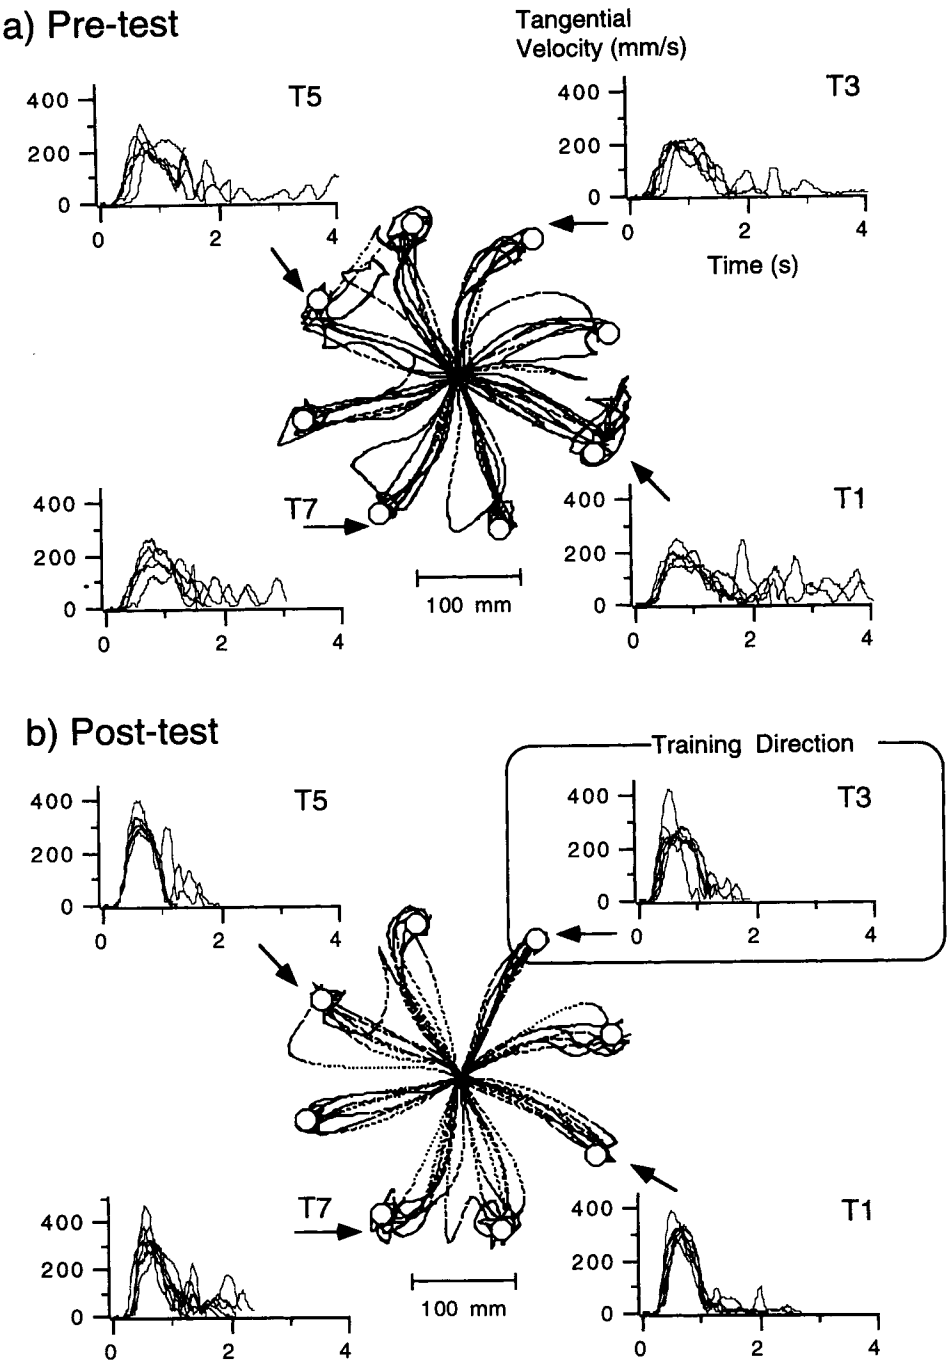

Figure 7. Typical trajectories and velocity profiles of 1 participant (T.F.) in the pretest session superimposed on one graph (a). Those of the same participant in the posttest session (b).

After describing the transformation in the same manner as Experiment 1, the experimenter instructed the participants as follows:

> After hearing the first tone cue generated by the computer, please move the cursor as close to the target as possible. Your time limit is 600 ms. At the end of this limit the computer will generate the second cue. You must freeze your movement after it.

Before the beginning of the experiment, the participants were allowed 16 trials for practice. The experimenter instructed them to remember the time interval of 600 ms during this practice period.

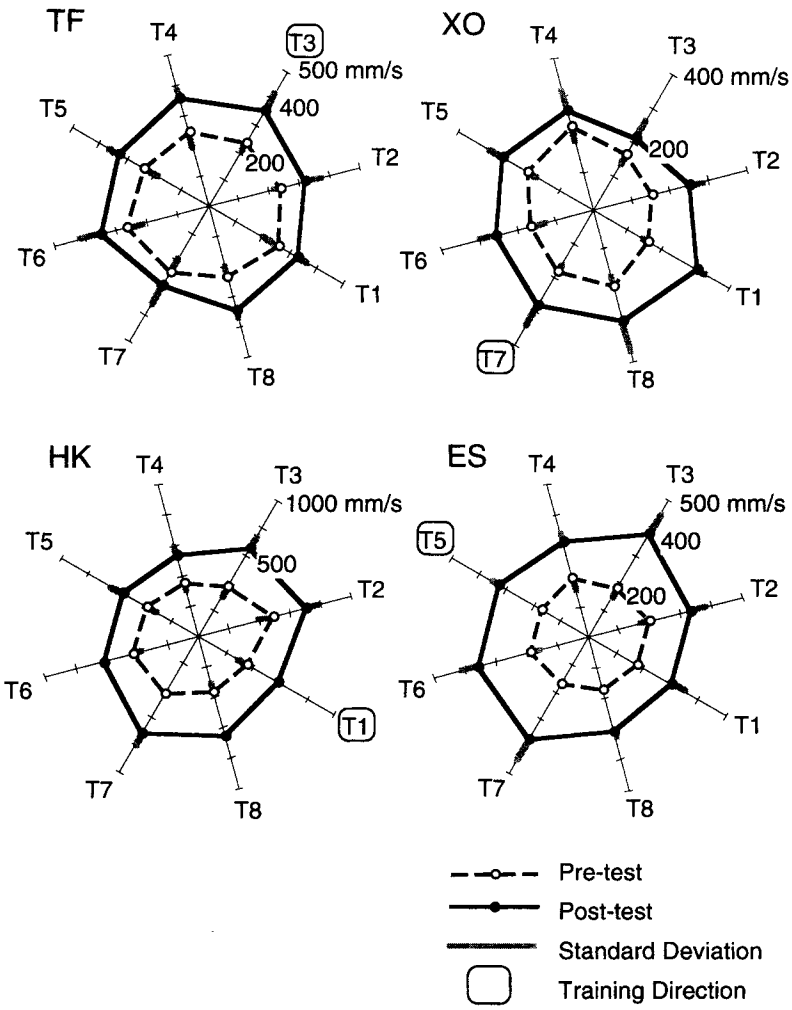

Figure 8. The mean maximum velocities of the pretest and posttest sessions plotted as a function of test target direction in the same manner as Figure 6.

The procedure before the first tone cue that signaled the beginning of the trial was the same as that of Experiment 1.

The pretest and the posttest each consisted of 112 trials (8 test targets × 14 times) and lasted about 20 min. The second training session consisted of 10 blocks of 54 trials (9 training targets × 6 times) and lasted about 150 min. A 5-min rest period followed every two blocks. Three (M.O., T.S., and S.I.) of the 4 participants completed the three sessions as experimental subjects. One control subject (J.H.) did not participate in the training session and rested during it. The training direction was determined for each experimental subject (M.O. = T3, T.S. = T5, and S.I. = T7) as shown in Figure 9.

Data Analysis

The data analysis followed that used in Experiment 1, except that we used the distance between the target and the end point of the ballistic movement, instead of the PT, as an indicator of the *learning effect*. We detected the end point of the first ballistic movement using the curvature methods developed by Pollick and Ishimura (1995) and calculated the distance between the target and this point.

Figure 10 illustrates how the end position of the first ballistic movement was determined. We defined the end point of the ballistic movement as the point at which the movement was corrected for the first time. At that point, the tangential velocity is close to zero and the curvature drastically increases. Parts A and B of Figure 9 show typical trajectories in one trial. We calculated the curvature (C) at each point in the trajectory using the following equation:

$$C = \frac{\dot{x}\dddot{y} - \dddot{x}\dot{y}}{(\dot{x}^2 + \dot{y}^2)^{3/2}}, \tag{4}$$

where $\dot{x}$ and $\dot{y}$ are instantaneous velocities of the x- and y-coordinates of the hand in the plane and $\ddot{x}$ and $\ddot{y}$ are the corresponding accelerations at that point. In Part C of Figure 9, the curvature is plotted as a function of time from the onset of the first cue signaling the beginning of the trial. We defined the end point of the first ballistic movement as the point at which the curvature ex-

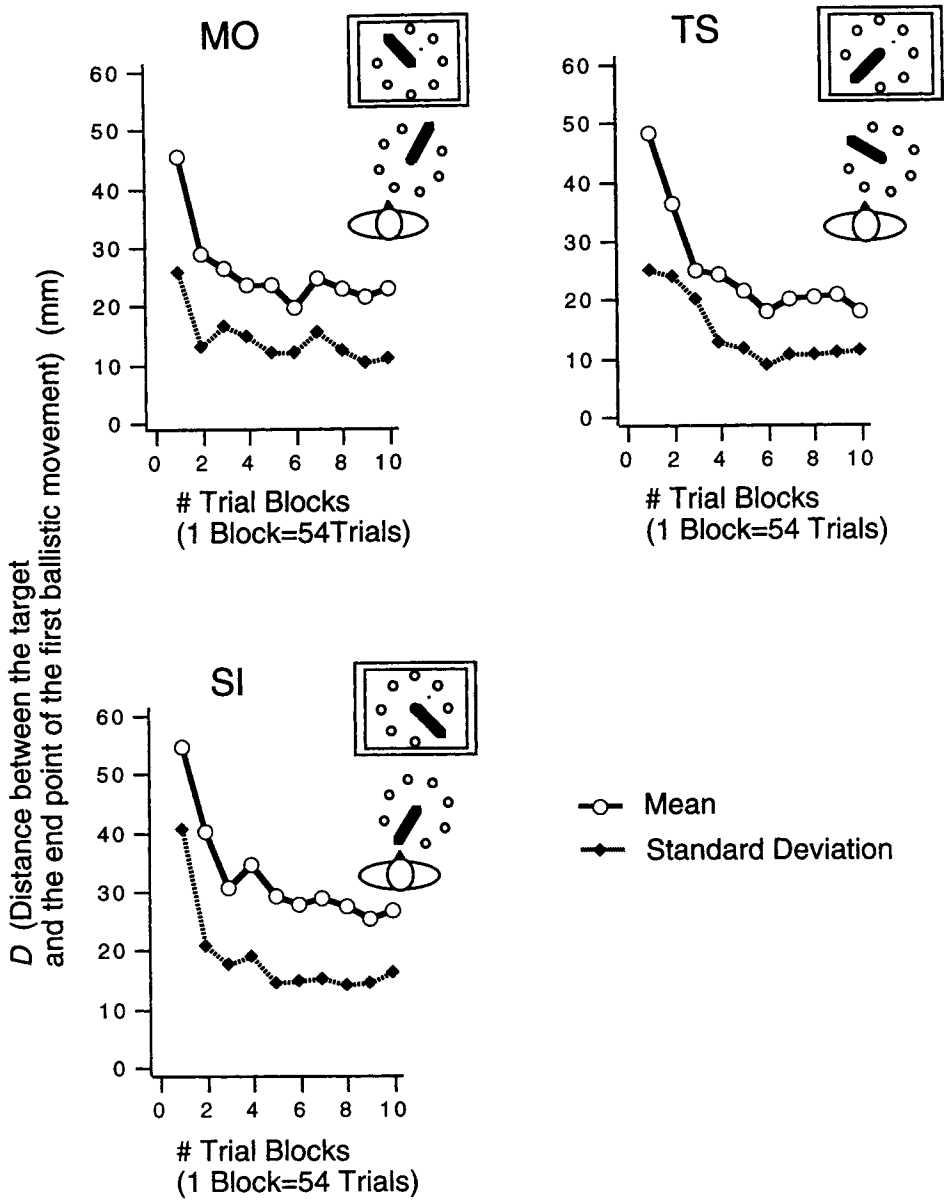

Figure 9. Learning profiles of 3 participants in the experimental group. The "clock" attached to each profile illustrates the training direction in both the hand space and on the CRT screen.

ceeded 0.5 mm^{-1} for the first time after the velocity was maximum (indicated by a square in Part C of Figure 9). The threshold value of 0.5 mm^{-1} was chosen during our pilot study as a compromise between two constraints: On the one hand, it should exceed the noise in the data; on the other hand, it should be small enough to detect the end of the ballistic movement reliably. At that point, the tangential velocity is almost at the local minimum, as shown in Part D of Figure 9. We calculated the distance between the center of the target and that point (*D*), as illustrated in Part A of Figure 9. We used the *D* values as an indicator of the effect of learning.

Another method for detecting the end point of the first ballistic movement is to find the first local minimum of the tangential velocity. However, we did not use this method because the curvature defined by Equation 4 changes more drastically than the tangential velocity at the point at which the corrective movement started (compare Parts C with D in Figure 9). If the feedforward trajectory is generated on the basis of the minimum torque-change principle (Uno et al., 1989), the curvature is close to zero during the movement, but it becomes infinitely large at the beginning and the end of the movement. The end point of the ballistic movement can be reliably estimated by detecting the point at which the curvature drastically increases.

The statistical analysis also followed that of Experiment 1, except for the treatment of outliers. The *D*s that exceeded mean ± 2 *SD*s (i.e., 95% confidence intervals of the *D* means) in the pretest and posttest were excluded from the statistical analysis as outliers. The maximum number of such *D*s for each target was 7 out of 112

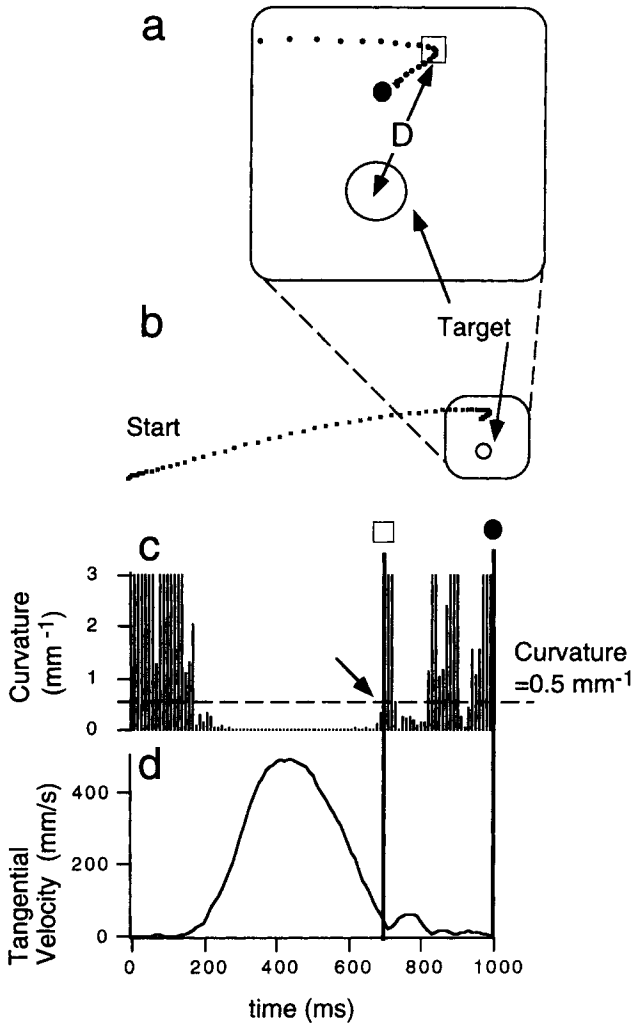

Figure 10. Typical trajectories in Experiment 2 showing the whole trajectory (b) and the part of it near the target (a). Curvature of the trajectory is calculated by Equation 4 (see text) and tangential velocity profiles plotted against time from the start signal (c and d). The end point of the trajectory is indicated by a solid circle, and the end point of the first ballistic movement detected by the method described in the text is indicated by a square.

in the pretest session and 6 out of 112 in the posttest session for all participants.

Results and Discussion

Learning Process During the Training Sessions

Distance between the target and the end point of the ballistic movement (D). In Figure 10, the mean and standard deviation of D in each block of trials are plotted separately for each of the 3 participants in the experimental group. Both the mean and standard deviation decreased with increasing number of trial blocks. The main effect of the trial block was significant, $F(6, 1{,}528) = 42.35$, $p < .0001$. We carried out post hoc comparisons (Tukey's HSD multiple comparisons) between trial blocks. They revealed significant differences for the first versus all the other blocks, the last versus the first block, and the last versus the second block, confirming that learning occurred during the training sessions.

Trajectories and velocity profiles. Parts A and B of Figure 11 show the typical trajectories of one participant (M.O.) in the first (Trials 1–9) and last (Trials 531–540) stages of the training session. Parts C and D of Figure 11 show the velocity profiles of these trials. The end points of these trajectories correspond to the end of the ballistic movements detected by the method described earlier. Therefore, no trajectories after the end of the ballistic movements are shown in these figures. The instantaneous tangential velocity at the end of the first ballistic movement is marked by a circle in Parts C and D of Figure 11. In the first and last stages, the trajectories were almost straight. However, the end points of the trajectories were distributed over a wide region in the first stage and were concentrated within the area where the training targets were located in the last stage. The velocity profiles in both the first stage and the last stage had a single-peak, bell-shaped appearance. However, they deviated from each other in the first stage but overlapped in the last stage. The deviation in the values of maximum velocity and the latency at the end of the ballistic movements decreased after training. These results suggest that the accuracy and stability of feedforward motor control increased at the end of the training session.

Comparisons Between the Pretest and the Posttest

Distribution of the end points of ballistic movements and velocity profiles. Parts A and B of Figure 12 show the typical distributions of end points of the first ballistic movements and velocity profiles of one participant (M.O.). The trajectories were omitted in Parts A and B of Figure 12 because they were nearly straight, as shown in Parts A and B of Figure 11, for all directions of the test targets. The end points were sparsely distributed over a wide region and often undershot the test targets in the pretest. However, they concentrated around each test target, and the tendency of the undershooting disappeared in the posttest. The velocity profiles in both the pretest and posttest had a single-peak, bell-shaped appearance. However, they deviated from each other in the pretest and overlapped in the posttest independent of the training direction. These results suggest that feedforward control became more accurate and stable in the posttest.

Distance between the target and the end point of the ballistic movement (D). In Figure 13, the mean D is plotted as a function of the test target direction in the same manner as in Figure 6. For participants in the experimental group, the mean and standard deviation of D varied with the test target direction in the pretest session, but they are constant in the posttest session. Furthermore, most of the mean D values in the posttest session (indicated by solid circles and solid lines) were shorter than those in the pretest session (open circles and dashed lines). The control subject, however, showed no consistent difference between the pre-

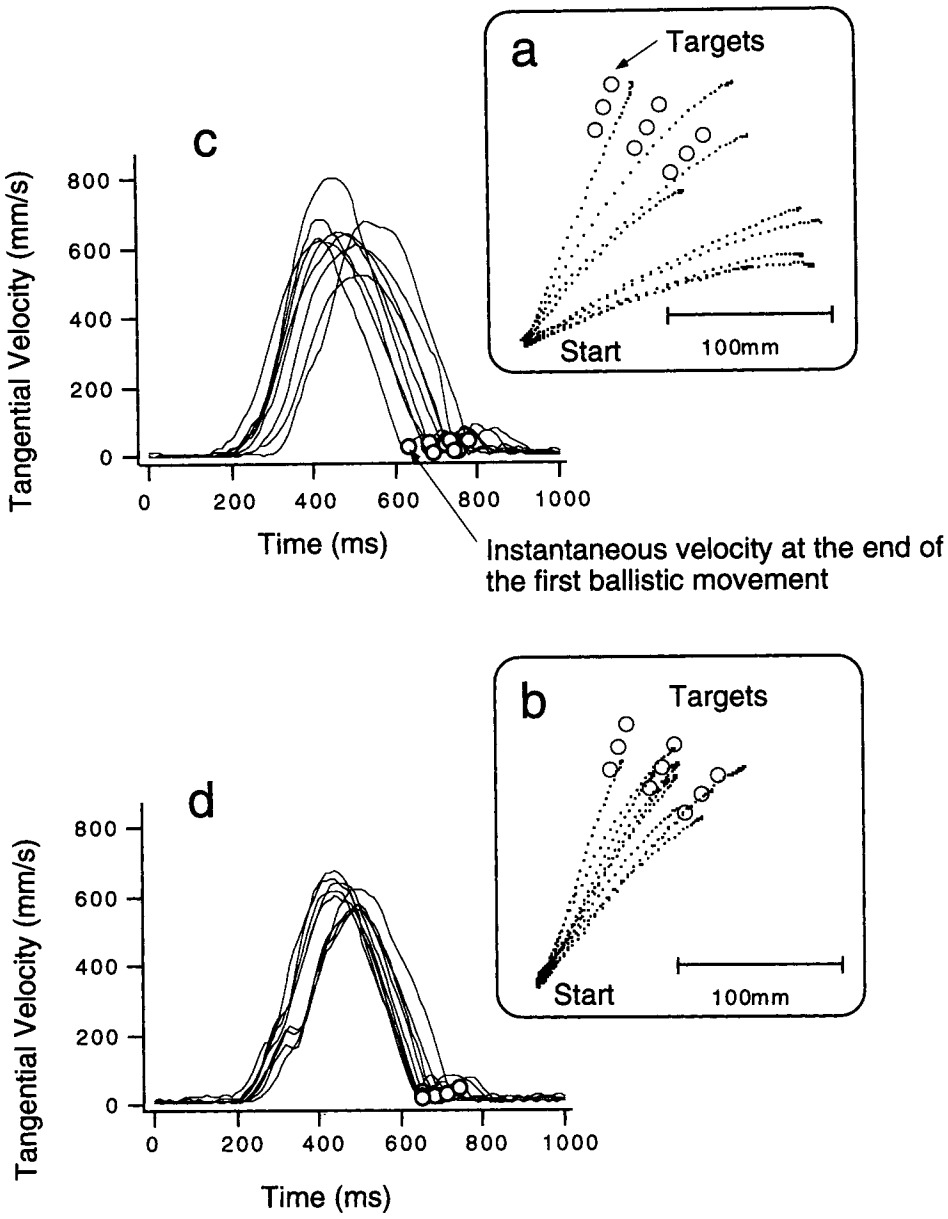

Figure 11. Typical trajectories of 1 participant (M.O.) in the first stage of the training session (Trials 1–9) superimposed on one graph (a). Those of the same participant in the last stage (Trials 531–540; b). Typical velocity profiles of the same participant in the first stage of the training session (c). The instantaneous velocity at the end of the first ballistic movement is marked by (circle a). Those of the same participant in the last stage (d).

test and the posttest. For the experimental subjects, the main effect of the test type (i.e., pretest or posttest) was significant, $F(1, 626) = 129.07$, $p < .0001$. However, for the control subject, the effect was not significant, $F(1, 206) = 0.17$, $p = .68$.

There was one exceptional case among the data mentioned earlier. M.O.'s mean D for T1 in the posttest was much larger than that in the pretest. The reason can be found in Part B of Figure 12. The end points of the ballistic movements aiming at T1 are indicated by an X in the figure. These end points shifted constantly toward T8 (i.e., in the counterclockwise direction). The mean directional error of these movements was 22.65°. The constant error must have caused the extra value of the mean D for T1 in the posttest. Such a phenomenon was not observed for the other participants.

The interaction between test type and angular discrepancy from the training direction was marginally significant for

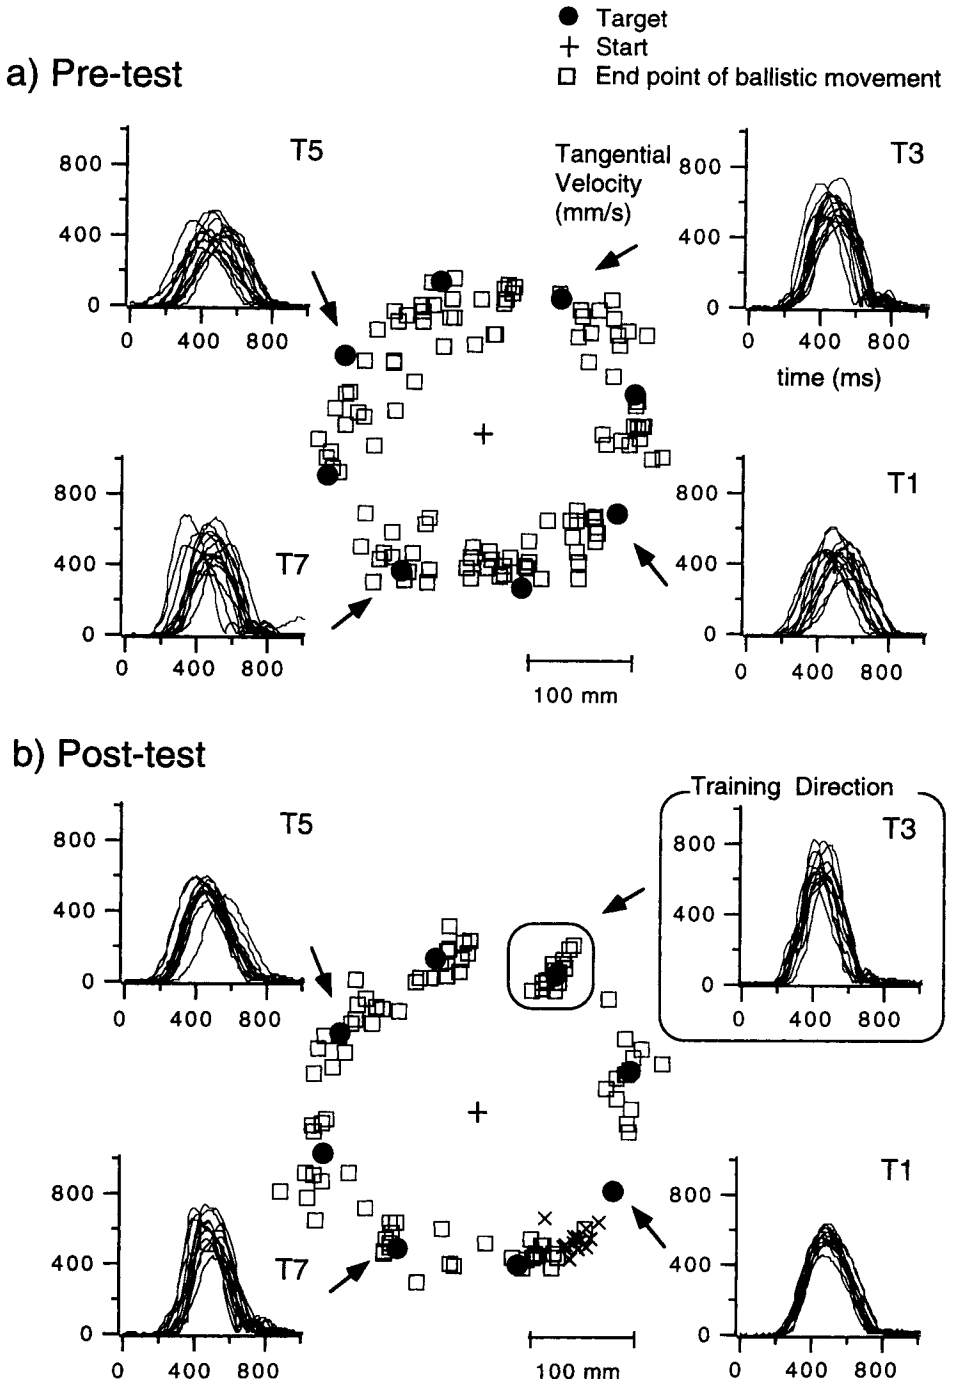

Figure 12. Typical positions of the end points of the first ballistic movements and velocity profiles of 1 participant (M.O.) in the pretest session superimposed on one graph (a). Those of the same participant in the posttest session (b). The end points of the ballistic movements in the posttest aimed at T1 are indicated by an *X*.

participants in the experimental group, $F(7, 645) = 1.99$, $p < .06$. We compared the mean *D*s of the pretest and posttest sessions for each angular discrepancy from the training direction using Tukey's HSD post hoc comparisons. Table 2 shows whether the difference between mean *D*s was significant in the pretest and posttest for each angular discrepancy from the training direction. The data in the table suggest that the effect of learning could be iden-

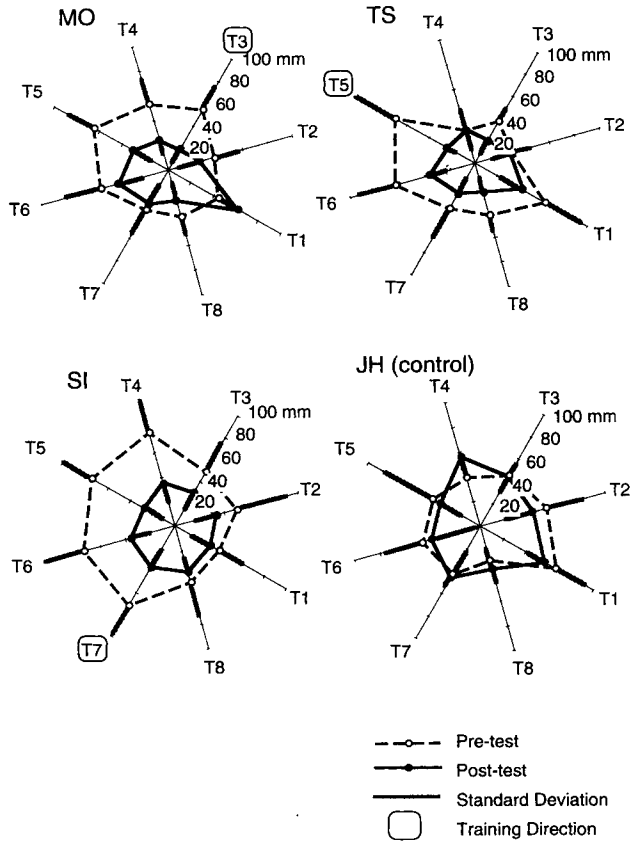

Figure 13. The mean distances of the pretest and posttest sessions plotted as a function of test target direction in the same manner as in Figure 6.

tified in the majority of directions (5 of 8 at the .05 level or below).

We removed those *D*s corresponding to end points, indicated by an *X* in Figure 12, when M.O. aimed at T1 (angular discrepancy from the training target = −90°) and carried out multiple comparisons on the residual *D*s. The results are shown in parentheses in Table 2. The difference between the mean *D*s in the pretest and the posttest for −90° was significant at the .001 level.

Consequently, the effect of learning was identified in various directions beyond the training direction. Furthermore, the effect was different among the angular discrepancies from the training direction.

General Discussion

We investigated generalization in visuomotor learning under kinematically transformed conditions to study the type of representation used by an internal kinematics model. In Experiment 1, there were no time constraints on the participant's performance except the instruction to acquire the target as rapidly as possible. Two phases were observed in the aiming performance: ballistic and corrective. In Experiment 2, the task and measurement were modified to investigate the accuracy of the ballistic movement. The accuracy and stability of feedforward control increased as the learning proceeded.

The results of Experiments 1 and 2 were similar in two respects: (a) The effect of learning could be identified in various directions other than the training direction and (b) the degree of the effect was different among the angular discrepancies from the training direction. The first point suggests that the generalization is global rather than local and that the representation the human motor system uses is not tabular. However, the second point suggests that the generalization is not as global as that predicted by structured representation. Thus, the generalization observed in our experiments was "intermediate," and it is impossible to classify the representation of the internal model as either of the two extreme types.

Extended Versions of Tabular and Structured Representations to Explain the Intermediate Generalization

Our findings suggest that the CNS is unlikely to use extreme types of the tabular and structured representations as an internal kinematics model. However, some extended types of the tabular or structured representation might have an intermediate generalization capacity. We now examine extended versions of these two extreme concepts.

First, participants might have used a sophisticated tabular representation rather than a simple one. The specific exem-

Table 2
Significance Levels of Difference Between the Mean D of the Pretest and Posttest as a Function of Angular Discrepancy From the Training Direction (Tukey's HSD Multiple Comparisons)

Angular discrepancy							
−135°	−90°	−45°	0°	45°	90°	135°	180°
$p < .05$	*ns* ($p < .001$)[a]	$p < .05$	$p < .001$	$p < .01$	$p < .05$	*ns*	*ns*

Note. HSD = honestly significant difference.
[a] This level of significance was obtained after removing *D*s corresponding to end points when participant M.O. aimed at T1; multiple comparisons were carried out on residual *D*s.

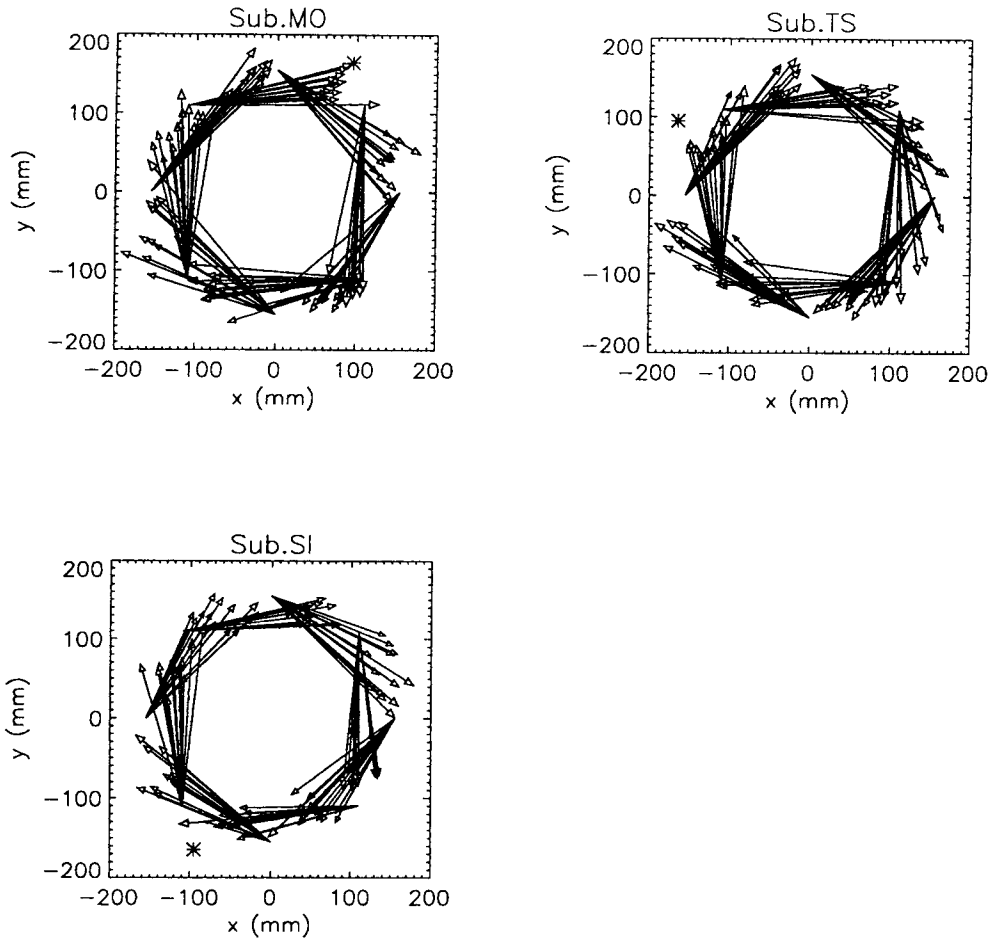

Figure 14. Vectors representing the displacement from the position of the target presented on the CRT screen to the end point of the first ballistic movement in the posttest of Experiment 2. The training direction for each of the participants is indicated by an asterisk. Sub. = Subject.

plar model, to which we referred in the introduction, is an example of such a representation. It stores specific exemplary pairs of input and output signals in memory and is a tabular representation in this respect. However, it has an intermediate generalization capacity because of the following operations: When it encounters a novel input signal, it calculates distances between the novel input signal and the stored exemplar signals in memory and finds some close exemplars. It then determines the output value by averaging the output values of these exemplars. It is important for the generalization properties of this representation to be characterized as follows: The effect of learning is most prominent when the input signal is close to that used during the training period, and the effect decreases monotonically as the input signal goes farther away from the training target (i.e., "graded response"). Computational models called *memory-based representations* (e.g., radial basis function [RBF; Poggio & Girosi, 1990] and K-means neural networks [Duda & Hart, 1973]) would produce results similar to this model's concerning the generalization properties.

Second, the respondents might have induced an improper structure or improper values of kinematic parameters in the structured representations. As we discussed in the *Method* section of Experiment 1, the participants were told about the 75° transformation of visual feedback around the center of the screen and about the 1:1 ratio of hand movement to cursor movement. However, there is no assurance that the participants represented the structure of the transformation (i.e., rotation around the center of the screen) properly. If the representation contained translation and magnification in addition to the rotation, then the participants' performance in the posttest would have been well suited for some particular target directions but not so for the other directions. Furthermore, they might have induced improper values of the parameters suited only for some particular target directions. It might have been difficult to acquire proper values through training because the training targets were restricted to the extremely small region as indicated in Figure 3. In these cases, intermediate generalization was observed.

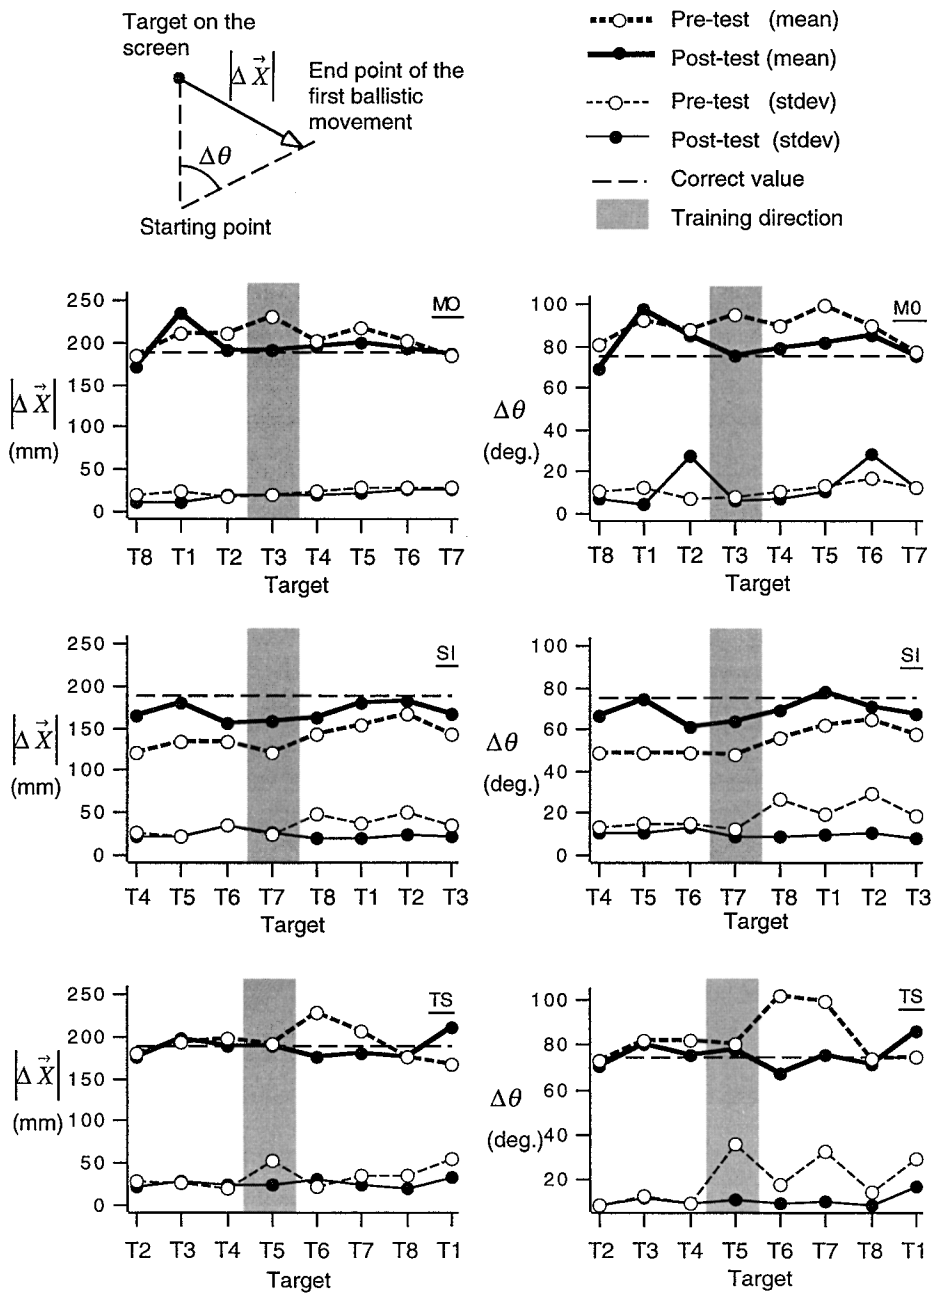

Figure 15. The top left panel illustrates the definitions of $|\vec{\Delta X}|$ and $\Delta\theta$. The other panels show the mean and standard deviation of magnitudes of the displacement vector ($|\vec{\Delta X}|$) and the angles of rotation ($\Delta\theta$) for each of the targets. Both of them are separately plotted for each of the participants and the pre- and posttests. MO, SI, and TS represent the principals.

Examinations of the Extended Versions of the Tabular and Structured Representations

We closely examined whether these extended versions could explain the results of our experiments. We first discuss the sophisticated tabular representation. We investigated whether a graded response would be evident in the participants' performance. Figure 14 shows vectors representing the displacement from the position of the target presented on the CRT screen to the end point of the first ballistic movement in the posttest of Experiment 2. These vectors directly indicate the transformation performed by the participants. We measured the magnitude of the vectors ($|\vec{\Delta X}|$) and the angles of rotation ($\Delta\theta$). The top of Figure 15

Table 3
Estimated Parameter Values by the Steepest Descent Method

	Model 1			Model 2				Model 3					
Participant	a: Rotation angle	b_1: Translation in x-axis (mm)	b_2: Translation in y-axis (mm)	a: Rotation angle	b_1: Translation in x-axis (mm)	b_2: Translation in y-axis (mm)	λ: Magnification	p_1	p_2	p_3	p_4	q_1	q_2
M.O.	−78.89°	−2.96	−0.51	−78.89°	−3.07	−0.49	0.95	0.20	0.86	−1.00	0.15	0.98	−3.01
S.I.	−68.65°	−0.87	−0.02	−68.65°	−0.86	−0.08	0.91	0.32	0.83	−0.86	0.34	−0.20	−0.96
T.S.	−75.34°	−2.22	−1.34	−75.34°	−2.28	−1.36	0.95	0.22	0.91	−0.93	0.25	1.16	−2.15

Note. p_1, p_2, p_3, and p_4 indicate parameters of linear transformation in Model 3. q_1 and q_2 indicate parameters of translation in Model 3.

illustrates the definitions of $|\overrightarrow{\Delta X}|$ and $\Delta\theta$, and the other parts show $|\overrightarrow{\Delta X}|$ and $\Delta\theta$ for each of the participants in the pre- and posttests. If the participant correctly performed the task and precisely acquired the target, $|\overrightarrow{\Delta X}|$ and $\Delta\theta$ became 188.7 mm and 75°, respectively (these values are indicated by broken horizontal lines in Figure 15). $|\overrightarrow{\Delta X}|$ and $\Delta\theta$ varied with the target direction in the pretest session, but they were consistently near the correct values in the posttest session for all participants in the experimental group.

We did not observe that the effect of learning was most prominent when the target was close to that used during the training period or that the effect monotonically decreased as it went farther away from the training target. Note that the difference between the value in the pretest and that in the posttest was almost constant for all targets in the case of 1 participant (S.I.). Such a uniform generalization property is apparently different from that of the sophisticated tabular representation. Thus, the memory-based representation, even in its sophisticated version, cannot explain the results.

Regarding the structured representation with an improper structure and with improper values of kinematic parameters, we tested three models corresponding to the kinematic transformation; they had different structures from each other. Two of them were based on the structure of the task used in Experiment 2, whereas the other was not. The first model had parameters representing the rotation angle (a) and the translation distances in the x- and y-axes (b_1, b_2). The mathematical formula of this model is represented by Equation 1. The second model had a parameter representing magnification (λ) in addition to those of the first model. Its mathematical formula is

$$\begin{pmatrix} x \\ y \end{pmatrix} = \lambda \begin{pmatrix} \cos a & \sin a \\ -\sin a & \cos a \end{pmatrix}^{-1} \begin{pmatrix} X - b_1 \\ Y - b_2 \end{pmatrix}. \tag{5}$$

The first model had a structure suited for the task in the experiment, whereas the second did not because the task required a rotational transformation without any magnification. We can call both models structured models because each of the parameters has a simple kinematic meaning (i.e., rotation, translation, or magnification). We could introduce another mathematical model when considering general linear transformation. If we had introduced a nonlinear model with many parameters, it might have represented the experimental results as a matter of course. Thus, we investigated how well a simple linear model given by

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} p_1 & p_2 \\ p_3 & p_4 \end{pmatrix} \begin{pmatrix} X \\ Y \end{pmatrix} + \begin{pmatrix} q_1 \\ q_2 \end{pmatrix} \tag{6}$$

could explain the results. However, this model is not a structured model according to the definition to which we referred in the introduction because we could not find simple kinematic meanings for the parameters (p_1, p_2, p_3, and p_4).

We investigated whether each of the models could explain the results of Experiment 2 in the following procedures. First, we estimated the parameter values of each

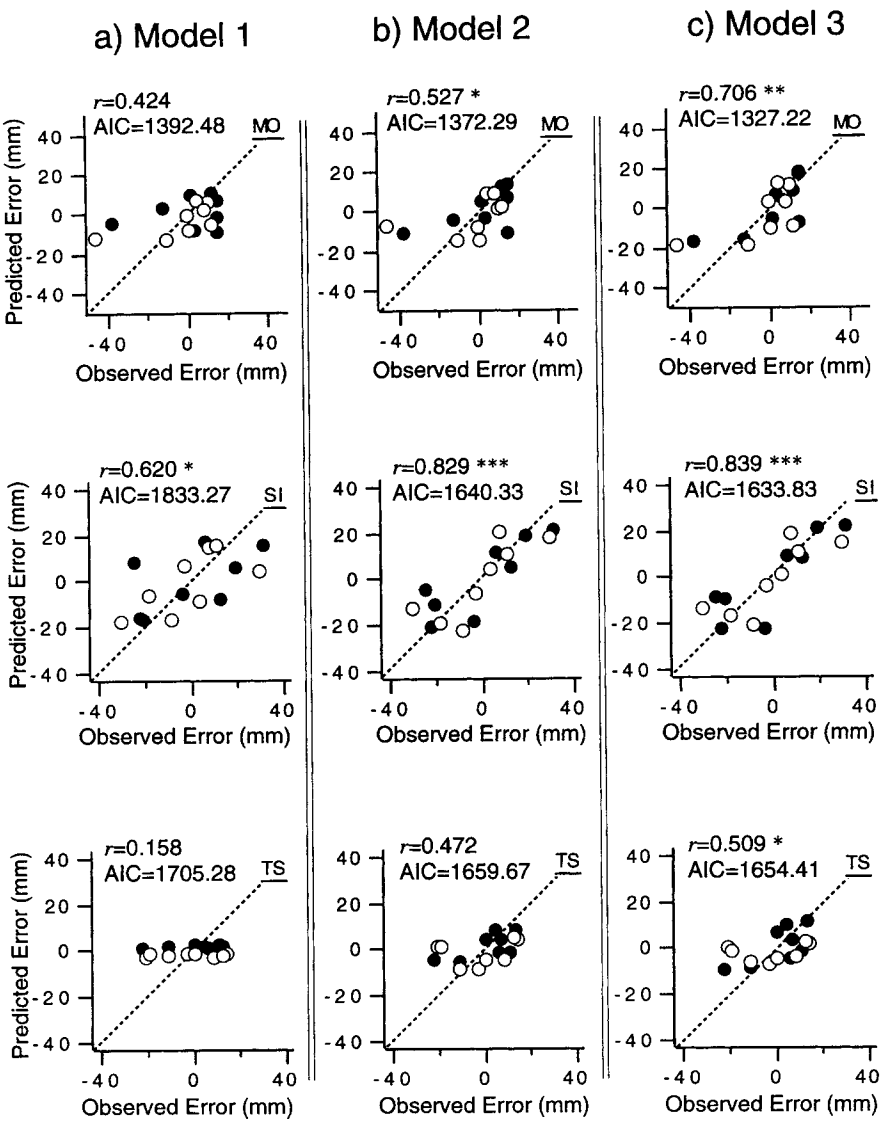

Figure 16. The correlation between the observed error distances in Experiment 2 and those predicted by each of the models. The distances were separately plotted for the x (solid circle) and y (open circles) directions. r = Pearson's correlation coefficient; AIC = Akaike's information criterion (Akaike, 1974). * $p < .05$. ** $p < .01$. *** $p < .001$.

model using the target positions on the screen (i.e., origins of the vectors shown in Figure 14) as input signals and the end points of the first ballistic movements (i.e., terminal points of the aforementioned vectors) in the posttest as output signals according to the steepest descent method (e.g., Bryson & Ho, 1975; see Appendix A). Second, we transformed the target position using each model and estimated parameter values and got the positions of the end points predicted by the models. Finally, we evaluated these models by investigating the correlation between the distance errors (i.e., distances between the target position and the end points) predicted by each model and those observed in the results of Experiment 2.

Table 3 shows the values of parameters estimated by the steepest descent method for each of the models. We calculated (a) the distances in the x and y directions between the target and the end points predicted by each model with the estimated parameter values and (b) those between the target and the observed end points that were averaged for each of the targets. Figure 16 shows the correlation between the distances for each model. Pearson's correlation coefficients exhibited the largest values and were significant in the third model for all respondents ($ps < .01$, .001, and .05 for participants M.O., S.I., and T.S., respectively). Thus, the third model seems to be the best for explaining the results of Experiment 2.

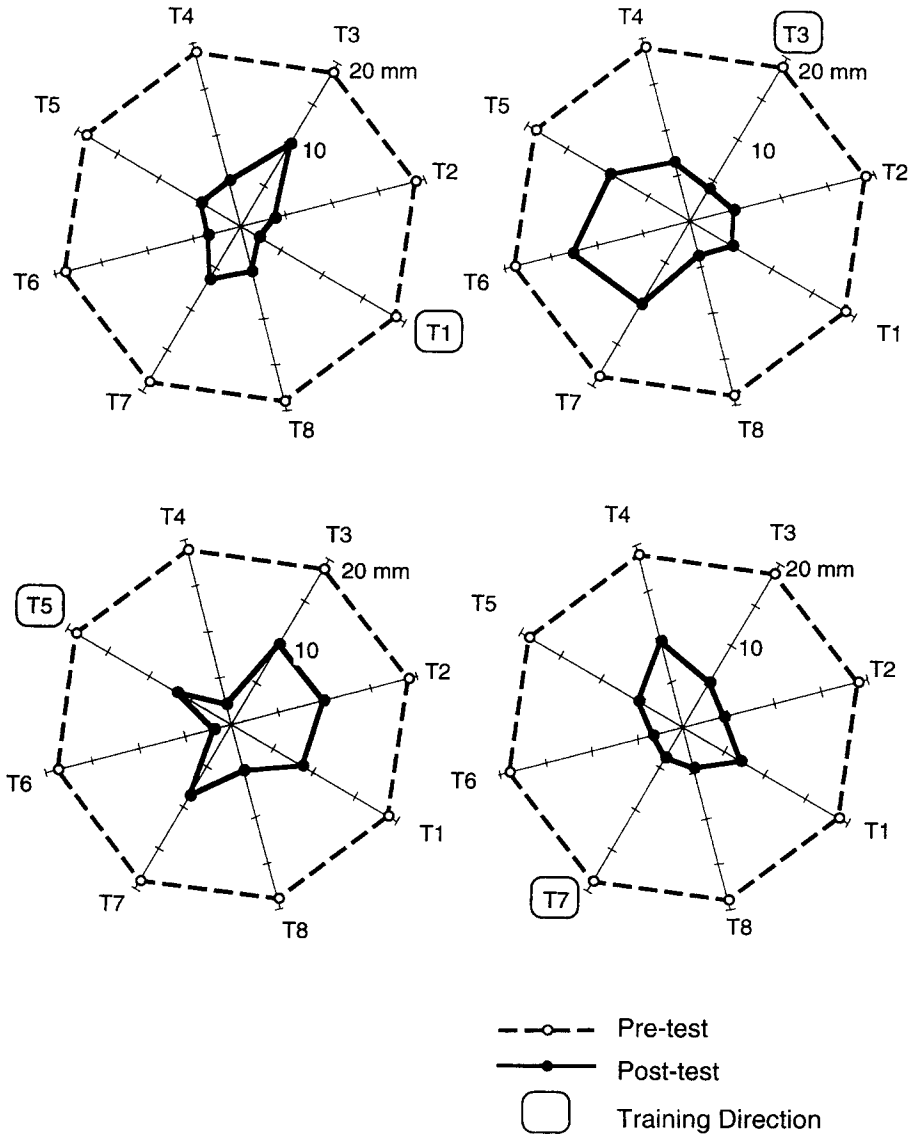

Figure 17. The distances of the pretest and posttest sessions in the stimulation data plotted as a function of training direction in the same manner as in Figure 6.

However, we had to take a number of parameters included in the three models into account before drawing a conclusion because a mathematical model can generally fit better given data as the number of parameters increases. The third model has six parameters, and the first and the second have three and four parameters, respectively. We calculated Akaike's information criterion (AIC; Akaike, 1974; see Appendix B) to evaluate the models taking these numbers of parameters into account. The results are also shown in Figure 16. The AIC value of the third model was the least for all of the participants (the model became better as the AIC value decreased). Thus, the third model was the best of the three. Furthermore, note that a matrix of the first four parameters (p_1, p_2, p_3, and p_4) in the third model did not correspond to any rotational matrices (see Table 3, Model 3). For example, in the case of M.O., p_1 is 0.20 and corresponds to cos(78.21°), whereas p_2 is 0.86 and corresponds to sin(59.03°). In this manner, the rotational angles given by the parameters in the same matrix are different from each other. We also could not find any simple physical meanings of these parameters for the other participants. These results suggest that structured representations, even in extended forms, cannot sufficiently explain the results of Experiment 2.

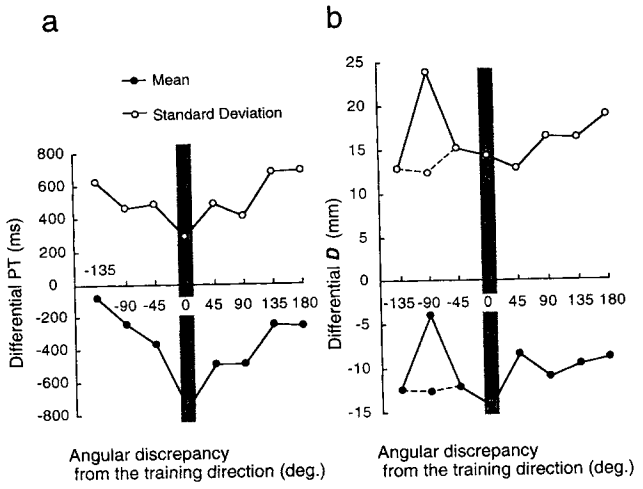

Figure 18. The mean and standard deviation of the differential performance times (PTs) in Experiment 1 (a) and distances (*D*s) in Experiment 2 (b) in the posttest sessions, plotted as a function of angular discrepancy from the training direction.

Generalization Properties of Connectionist Representation

A promising candidate to explain the intermediate generalization observed in our experiments is the learning in multiple-layer perceptrons. It can approximate any continuous function, including a linear one, by increasing the number of hidden units (Funahashi, 1989). Some theoretical approaches have discussed generalization properties of a neural network. Baum and Haussler (1989), for example, addressed the relation between the size of a neural network and its generalization ability. Poggio and Girosi (1990) pointed out that learning an input–output mapping from a set of examples in neural networks is closely related to classical approximation techniques; the networks are not only equivalent to generalized splines but are also closely related to the classical radial basis functions used for interpolation tasks.

As we mentioned before, the kinematic parameters are implicitly represented by a considerable number of synaptic weights in this approach. The number of parameters varies according to the number of hidden units in the network. Geman, Bienenstock, and Doursat (1992) pointed out that the mean-squared error of a network can be decomposed into a bias term and a variance term. If the number of parameters is too small, the performance of the network degrades because the bias increases; the bias can be reduced by increasing the number of parameters. A large number of parameters, however, causes a high variance, and large training samples are required to reduce the variance and to achieve an acceptable performance. They called this trade-off the "bias/variance dilemma."

Consequently, when the training data are restricted to particular sets of samples, as in our experiments, the bias for the training data can be reduced by increasing the number of parameters, but the variance cannot. Thus, in the case of a large number of parameters, there is no assurance that the network can correctly estimate a broad range of input–output mappings beyond the training region because the output values might be too sensitive to the location of the training points. In fact, our results show that the generalization in visuomotor learning is not perfect. For example, in Experiment 2, the learning effect was not present for some directions for 2 participants (M.O. and T.S.; see Figure 13).

Furthermore, it depends on the initial conditions of hidden units how a multilayer perceptron trained using small samples estimates input–output mappings beyond the training region. There is no assurance that the error of the network monotonically decreases as the input signal becomes similar to that used in the training. Note that the effect of learning did not monotonically decrease as the angular discrepancy from the training direction increased. Thus, the generalization properties observed in the results of Experiment 2 resemble those from "extrapolation"[3] by a neural network trained using small samples.

Recently, we examined the generalization properties in the learning of three-layer perceptrons by computer simulation (Imamizu, Uno, & Kawato, 1993). The purpose of this simulation was not to examine whether this representation could explain the results of Experiment 2 better than the representations we previously examined using the steepest descent method, but to investigate whether it could qualitatively reproduce the results after learning the same transformation that the human participants had learned.

We briefly describe the procedure of the simulation. The input signals for the network were the target positions on the screen, and the output signals corresponded to the end points of the first ballistic movements. First, the network learned the identity map from the screen to the hand coordinates before learning the rotation. Second, its performance was tested at the eight global targets used in Experiments 1 and 2. Third, it learned the function given by Equation 1 by using the nine local target positions in the hand coordinates as the teaching signals. These target positions were also used in Experiments 1 and 2. Finally, its performance was retested using the eight global target positions. In the third step and final step, random noises were added to the input and teaching signals to imitate the variability found in performance of human participants. Their minimum and maximum values were −5 and 5 cm, respectively.

Figure 17 shows the results of the simulation. They are similar to the results of Experiment 2 in two respects: (a)

[3] In this research, we focused on how participants "extrapolate" the relation between the stimuli and response of the practice pairs. However, the function forms of generalization might be understood much more clearly by studying the "interpolation" between the practice targets. For example, Bedford (1989) asked his respondents to learn to point at visual targets while wearing prism goggles that distorted the targets' perceived spatial positions. He studied how the respondents interpolated in pointing at new transfer targets and obtained some evidence for the primacy of linear functions during the induction of continuous stimulus–response relations.

The effect of learning could be identified in various directions other than the training direction and (b) the degree of the effect was different among the angular discrepancies from the training direction.

Differences Between Experiments 1 and 2

The results of Experiments 1 and 2 were similar in the two points discussed earlier, but they were different in that the learning effect in Experiment 1 was most prominent when the test target was close to the training direction, whereas this tendency was not as prominent in Experiment 2. Additionally, the range of generalization in Experiment 2 seemed to be wider than that in Experiment 1.

The difference in these results became more prominent when we studied the PTs of Experiment 1 and the *D*s of Experiment 2 in the posttest. We calculated the mean and standard deviation of differential PTs and *D*s for each participant in the experimental group in the posttest session (the differential method is described in Experiment 1). In Parts A and B of Figure 18, the differential PTs and *D*s in the posttest are plotted as a function of angular discrepancy from the training direction. All of the mean values are negative because the normalized zero corresponds to the mean PT or *D* of the pretest and posttest for each respondent. In Experiment 1 (see Part A of Figure 18), the mean and standard deviation of the PTs in the posttest decreased from −135° to the training direction (0°) and continuously increased from the training direction to 180°. In Experiment 2 (see Part B of Figure 18), the mean and standard deviation of *D* at −90° were extraordinarily large. However, if we remove the *D*s of participant M.O. for −90° for the reason mentioned earlier, the *D*s become relatively constant and small compared with the results of Experiment 1 (indicated by the dashed lines). Consequently, although the effect of learning was identified in various directions beyond the training direction, the generalization was relatively local in Experiment 1 and global in Experiment 2.

It might not be appropriate to directly compare the results of Experiment 1 with those of Experiment 2 because different measures were used. However, if there were a different effect of generalization between Experiments 1 and 2, there are two possible explanations.

First, there was a procedural difference between these two experiments, that is, whether corrective movement (feedback control) was used. In Experiment 1, the analysis of velocity profiles (see Figures 5 and 7) suggested that the early part of the movement was executed by feedforward control, whereas the latter part was executed by feedback control. On the other hand, in Experiment 2, the analysis (see Figures 11 and 12) suggested that almost the entire movement was executed by feedforward control. Thus, one can assume that the different posttest results in Experiments 1 and 2 were due to the different effect of learning in the feedforward control and feedback control. On the basis of these results, we guess that the range of generalization in the feedforward control was wider than that in the feedback control.

Second, the relatively local generalization observed in Experiment 1 might have arisen because the velocity of movement during the posttest was much higher than that in the pretest (see Figures 7 and 8). In this research, we restricted our attention to the internal models of kinematic transformations; however, there is no assurance that the models of the motor apparatus are built separately for kinematics and dynamics. If the CNS has one model capable of solving kinematics and dynamics simultaneously, the performance of the participants might depend on both the direction of the target and the velocity. In Experiment 1, the participants learned to aim at targets in a specific direction at high speed during the training session, but they might not have learned how to aim at targets in other directions at that speed. On the other hand, in Experiment 2, the PT was constant (600 ms), and the velocity in the posttest was almost compatible with that in the pretest.

In conclusion, the type of representation used in an internal kinematic model is nontabular and nonstructured. That is, it is unlikely that the CNS builds the internal model by either directly representing the kinematic parameters or representing the local association between sensory input and motor output. Instead, a promising candidate to explain the results of experiments is a neural network model with a medium number of neurons.

However, the computational characteristics of the connectionist representation that could reproduce the results of our experiments might resemble those of the structured representation. This is because, in this research, the structured model whose parameters were rotation angles, translation distances, and magnification could explain the results to some extent (but not sufficiently). We can find the reason for this in the task instruction of our experiments; the experimenter informed the participants of the nature of the transformation (i.e., 75° rotation) between the hand and the screen coordinates before the beginning of the experiments. One of our future goals is to investigate how such cognitive knowledge of the task affects the organization of connectionist representations.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, AC-19,* 716–723.

Albus, J. S. (1975). A new approach to manipulator control: The cerebellar model articulation controller (CMAC). *Journal of Dynamic Systems, Measurement, and Control, 97,* 220–227.

An, C. H., Atkeson, C. G., & Hollerbach, J. M. (1988). *Model-based control of a robot manipulator.* Cambridge, MA: MIT Press.

Atkeson, C. G. (1989). Learning arm kinematics and dynamics. *Annual Review of Neuroscience, 12,* 157–183.

Baum, E. B., & Haussler, D. (1989). What size net gives valid generalization? *Neural Computation, 1,* 151–160.

Bedford, F. L. (1989). Constraints of learning new mappings between perceptual dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 232–248.

Bryson, A. E., Jr., & Ho, Y. (1975). Parameter optimization prob-

lems. In A. E. Bryson & Y. Ho (Eds.), *Applied optimal control* (pp. 1–41). Washington, DC: Hemisphere.

Chamberlin, C. J., & Magill, R. A. (1992). A note on schema and exemplar approaches to motor skill representation in memory. *Journal of Motor Behavior, 24,* 221–224.

Cunningham, H. (1989). Aiming error under transformed spatial mappings suggests a structure for visual-motor mappings. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 439–506.

Duda, R. O., & Hart, P. E. (1973). Unsupervised learning and clustering. In R. O. Duda & P. E. Hart (Eds.), *Pattern classification and scene analysis* (pp. 189–260). New York: Wiley.

Flash, T. (1987). The control of hand equilibrium trajectories in multi-joint arm movements. *Biological Cybernetics, 57,* 257–274.

Flowers, K. (1975). Ballistic and corrective movements on an aiming task: Intention tremor and Parkinsonian disorders compared. *Neurology, 25,* 413–421.

Funahashi, K. (1989). On the approximate realization of continuous mapping by neural networks. *Neural Networks, 2,* 183–192.

Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and bias/variance dilemma. *Neural Computation, 4,* 1–58.

Georgopoulos, A. P., Kalaska, J. F., & Massey, J. T. (1981). Spatial trajectories and reaction times of aimed movements: Effects of practice, uncertainty and change in target location. *Journal of Neurophysiology, 46,* 725–743.

Hollerbach, J. M. (1982). Computers, brains and the control of movement. *Trends in Neurosciences, 56,* 189–192.

Imamizu, H., Uno, Y., & Kawato, M. (1993). Generalization of visuo-motor learning. *Society for Neuroscience Abstracts, 19,* 1594.

Jordan, M. I. (1990). Motor learning and the degree of freedom problem. In M. Jannerod (Ed.), *Attention and performance XIII* (pp. 796–836). Hillsdale, NJ: Erlbaum.

Kawato, M. (1990). Computational schemes and neural network models for formation and control of multijoint arm trajectory. In T. Miller, R. Sutton, & P. Werbos (Eds.), *Neural networks for control* (pp. 197–228). Cambridge, MA: MIT Press.

Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical network model for motor control and learning of voluntary movement. *Biological Cybernetics, 57,* 169–185.

Kawato, M., & Gomi, H. (1992). The cerebellum and VOR/OKR learning models. *Trends in Neurosciences, 15,* 445–453.

Keele, S. W. (1981). Behavioral analysis of movement. In J. M. Brookhart, V. B. Mountcastle, V. B. Brooks, & S. R. Geiger (Eds.), *Handbook of physiology: Section I. The nervous system* (pp. 1391–1414). Bethesda, MD: American Physiological Society.

Koh, K., & Meyer, D. E. (1991). Function learning: Induction of continuous stimulus–response relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 811–836.

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

Poggio, T., & Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science, 247,* 978–982.

Pollick, F. E., & Ishimura, G. (1995). *The three-dimensional curvature of straight-ahead movement.* Manuscript in preparation.

Rumelhart, D. E. (1986). *Parallel distributed processing.* Cambridge, MA: MIT Press.

Saltzman, E. (1979). Levels of sensorimotor representation. *Journal of Mathematical Psychology, 20,* 91–163.

Saltzman, E. (1987). Skilled actions: A task-dynamic approach. *Psychological Review, 94,* 84–106.

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

Schmidt, R. A. (1988). The process of learning. In R. A. Schmidt (Ed.), *Motor control and learning* (pp. 457–491). Champaign, IL: Human Kinetics Publishers.

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

Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectory in human arm movement: Minimum torque-change model. *Biological Cybernetics, 61,* 89–101.

Appendix A

Parameter Estimation Using the Steepest Descent Method

We explain how we estimated parameter values from the experimental data using the steepest descent method (e.g., Bryson & Ho, 1975). Take the case of the first model. The number of trials was 112 in Experiment 2, and there were 112 pairs of input (the target positions on the screen: X and Y) and output (the end points of the first ballistic movements: x and y) as the observed data. According to the first model (i.e., Equation 1), the relation between the ith input and output signals is represented by

$$x_i = (X_i - b_1)\cos a + (Y_i - b_2)\sin a$$

$$y_i = -(X_i - b_1)\sin a + (Y_i - b_2)\cos a.$$

Our purpose is to find a set of parameter values (a, b_1, and b_2) that best fit the observed pairs of input and output signals. This task is equivalent to minimizing the error function E given by

$$E = \frac{1}{2}\sum_{i=1}^{112}[(x_i^M - x_i)^2 + (y_i^M - y_i)^2].$$

Here x_i^M and y_i^M are the outputs of the model for the ith input. The best set of parameters can be obtained in an iterative manner (described next).

We set some adequate values of parameters and calculated E for the first time. We changed each of the parameter values so as to decrease the value of E according to the following iteration rule:

$$a^{k+1} = a^k - \varepsilon \cdot \frac{\delta E}{\delta a^k},\ b_1^{k+1} = b_1^k - \varepsilon \cdot \frac{\delta E}{\delta b_1^k},\ \text{and}\ b_2^{k+1} = b_2^k - \varepsilon \cdot \frac{\delta E}{\delta b_2^k}.$$

(*Appendix A continues on next page*)

Here, a^k, b_1^k and b_2^k are the parameter values in the kth iteration, and ε is a positive value that is small enough (from 10^{-5} to 10^{-8}) to avoid divergence. The parameter values each converged to the best value after 100,000 iterations.

We also estimated the parameter values in the same manner for the other models, although the numbers of their parameters are different from the first.

Appendix B

Akaike's (1974) Information Criterion

Akaike's (1974) information criterion is a method used to evaluate models taking a number of parameters into account. According to Akaike, this criterion is defined as

$$\mathrm{AIC} = -2\log(\hat{\theta}) + 2(k+1)\,. \tag{A1}$$

Here, $L(\hat{\theta})$ is the maximum likelihood of a model F_θ, and k is the number of parameters. Thus, the first term corresponds to the goodness of the fit, and the second is a penalty for increasing the number of parameters. If the input signals of the model are w_1, w_2, . . ., w_n, the output signals are $F_\theta(w_1)$, $F_\theta(w_2)$, . . ., $F_\theta(w_n)$, and the observed experimental data are z_1, z_2, . . ., z_n, then the error of the model for the ith input signal is given by

$$\varepsilon_i = z_i - F_\theta(w_i)\,.$$

If the errors distribute in the normal distribution whose mean and variance are 0 and σ, respectively, the model's probability density function p is described as

$$p(\varepsilon) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left\{\frac{[z_i - F_\theta(w_i)]^2}{2\sigma^2}\right\}.$$

Its likelihood function $L(\theta)$ is represented by

$$L(\theta) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi\sigma}} \exp\left\{\frac{[z_i - F_\theta(w_i)]^2}{2\sigma^2}\right\}.$$

This function becomes maximum when σ equals the variance of the observed data ($\hat{\sigma}$). Thus,

$$L(\hat{\theta}) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi\hat{\sigma}}} \exp\left\{\frac{[z_i - F_\theta(w_i)]^2}{2\hat{\sigma}^2}\right\}.$$

By substituting this relation into Equation A1, AIC can be written as

$$\mathrm{AIC} = -2n\log\frac{1}{\sqrt{2\pi\hat{\sigma}}} + \sum_{i=1}^{n}\left\{\frac{[z_i - F_\theta(w_i)]^2}{\hat{\sigma}^2}\right\} + 2(k+1)\,.$$

We used this equation to get AIC values. First, we calculated the squared sum of the errors of the model (i.e., the numerator of the second term). If the position of the end point of the first ballistic movement is (x_i, y_i) and those predicted by the model are (x_i^M, y_i^M), then the calculation becomes

$$\sum_{i=1}^{n} [(X_i - X_i^M)^2 + (y_i - y_i^M)^2]\,.$$

Here, we assume that errors in the x and y directions are independent of each other. Second, we calculated the variance of the observed data ($\hat{\sigma}$). Let the distance between the end point of the first ballistic movement and the target position in the hand coordinates (x_i^T, y_i^T) be $\bar{\rho}$, and so forth,

$$\bar{\rho} = \sqrt{(x_i - x_i^T)^2 + (y_i - y_i^T)^2}\,.$$

Then, the mean and variance are represented by $\bar{\rho}$ and σ_ρ, respectively, and one can get both values directly from the observed data. We now consider the variance of ρ in the x direction because we have assumed that errors in the x and y directions are independent of each other. The distance in the x direction is given by

$$\varepsilon_x = \rho\sin\phi\,.$$

Here, ϕ is defined as arctan $([y_i - y_i^T]/[x_i - x_i^T])$. Let the mean and variance of ε_x be $E(\varepsilon_x)$ and $V(\varepsilon_x)$, respectively. Also, suppose $P(\rho)$ and $Q(\phi)$ represent distribution functions of ρ and ϕ, respectively. Then

$$E(\varepsilon_x) = \int P(\rho)\rho d\rho \int Q(\phi)\sin\phi d\phi = \bar{\rho}\int_0^{2\pi} \frac{1}{2\pi}\sin\phi d\phi\,.$$

Because $\int_0^{2\pi} \frac{1}{2\pi}\sin\phi d\phi$ is constantly zero, $E(\varepsilon_x)$ becomes zero. Thus,

$$V(\varepsilon_x) = E\{[\varepsilon_x - E(\varepsilon_x)]^2\} = E(\varepsilon_x^2)$$

and

$$E(\varepsilon_x^2) = \int P(\rho)\rho^2 d\rho \int Q(\phi)\sin^2\phi d\phi$$

$$= \sigma_\rho \int_0^{2\pi} \frac{1}{2\pi}\cdot\frac{1-\cos 2\phi}{2} d\phi$$

$$= \frac{\sigma_\rho}{2}\,.$$

We used $V(\varepsilon_x)$ as the variance of the observed data ($\hat{\sigma}$).

Received November 10, 1993
Revision received September 15, 1994
Accepted October 14, 1994 ■