

A computational theory of motor learning

WAYNE IBA AND PAT LANGLEY

Department of Information and Computer Science, University of California, Irvine, CA 92717, U.S.A.

Received August 20, 1987

Revision accepted December 18, 1987

In this paper we present a computational theory of human motor performance and learning. The theory is implemented as a running AI system called MAGGIE. Given a description of a desired movement as input, the system generates simulated motor behavior as output. The theory states that skills are encoded as *motor schemas*, which specify the positions and velocities of a limb at selected points in time. Moreover, there exist two natural representations for such knowledge; *viewer-centered* schemas describe visually perceived behavior, and *joint-centered* schemas are used to generate behavior. When the model acts upon these two representational formats, they exhibit quite different behavioral characteristics. MAGGIE performs the desired movement within a feedback control paradigm, monitoring for errors and correcting them when it detects them. Learning involves improving the joint-centered schema over many practice trials; this reduces the need for monitoring. The model accounts for a number of well-documented motor phenomena, including the speed-accuracy trade-off and the gradual improvement in performance with practice. It also makes several testable predictions. We close with a discussion of the theory's strengths and weaknesses, along with directions for future research.

Key words: machine learning, motor behavior, skill improvement, speed-accuracy trade-off, hill climbing, motor schemas, motor programs.

Dans cet article, nous présentons une théorie computationnelle de l'apprentissage et de la performance moteurs de l'homme. La théorie a été mise en oeuvre sous forme d'un système d'exécution IA appelé MAGGIE. Selon cette théorie, les aptitudes sont encodées sous forme de schémas moteurs, qui spécifient les positions et vitesses d'un membre à des points précis dans le temps. De plus, il existe deux représentations naturelles de ce type de connaissances : des schémas fondés sur le spectateur décrivant un comportement perçu visuellement, et des schémas fondés sur le groupage, utilisés pour produire un comportement. Lorsque le modèle réagit à ces deux formes de représentation, leurs caractéristiques de comportement sont très différentes. MAGGIE exécute le mouvement désiré à l'intérieur d'un paradigme de contrôle de réaction, détectant et corrigeant les erreurs le cas échéant. L'apprentissage nécessite l'amélioration du schéma fondé sur le groupage après de multiples essais; cela réduit le besoin de surveillance. Le modèle tient compte d'un certain nombre de phénomènes moteurs bien documentés, notamment le compromis vitesse-exactitude et l'amélioration graduelle de la performance avec la pratique. Il peut également faire plusieurs prévisions vérifiables. L'article se termine par une discussion sur les points forts et les points faibles de cette théorie, ainsi que sur les domaines dans lesquels des recherches futures pourraient être entreprises.

Mots clés : apprentissage automatique, comportement moteur, amélioration des aptitudes, compromis vitesse-exactitude, escalade, schémas moteurs, programmes moteurs.

[Traduit par la revue]

Comput. Intell. 3, 338-350 (1987)

1. Introduction

The ability to generate skilled movements is shared by experts in such diverse domains as professional basketball and concert music. Although skilled performance attracts one's attention, the acquisition of such expertise is even more intriguing. Many years of practice go into such improvement, and motor learning never completely halts; the old saying that "practice makes perfect" oversimplifies the process greatly. Our long-term goal is to develop a computational theory of motor skills and their acquisition. This theory should account for known aspects of human motor behavior, and ideally it should predict unobserved phenomena as well. In this paper, we describe the progress we have made by limiting our attention to one aspect of motor behavior — the task of refining an existing motor skill through practice.

Most research on motor behavior has focused on either very high level or very low level aspects of motor control. High-level work (e.g., Fikes *et al.* 1972; Segre 1987) has addressed issues of planning: generating a sequence of abstract operators that produce complex motor behavior. The low-level work in robotics has addressed the generation of appropriate torques and voltages (e.g., Swartz 1984) and in the case of neurophysiology has considered the excitation levels of various neurons (e.g., Arbib 1982; Davis 1976). Our main concern

lies with neither of these levels. Instead, our work focuses on an intermediate level at which torques and voltages need not be specified, but where the operators are still rather primitive.

There are two basic approaches to studying motor behavior that are orthogonal to these three levels of focus. The engineering approach, represented by robotics, has developed algorithms that control the movement of an arm, but these algorithms lack psychological plausibility. The "natural organism" approach, represented by neurophysiology, physiology, and psychology, has devised theories that account for observed phenomena, but often these theories are not computational in nature.¹ In our work on motor behavior, we have addressed both of these concerns: our theory is computational and also accounts for findings about human motor skills. In this respect, our results narrow the gap between robotics research and psychological work on motor control.

In this paper we introduce MAGGIE, a computational model of human motor behavior. We begin by reviewing some results from research on human motor skills. In Sect. 3 we present a detailed description of MAGGIE's performance system, along

¹Some psychologists have presented "computational" models that consist of mathematical equations, but that are not implemented as a coherent computer simulation.

with experimental results on the system's behavior. In Sect. 4 we describe the learning mechanism together with results from experiments on learning motor skills. We conclude by discussing the model's successes and failures in explaining various psychological phenomena, and by outlining our plans for further research in this area.

2. Background on human motor behavior

We do not have the space to present a complete review of the relevant research in motor behavior. We have selected material that will help motivate and justify claims made later in the paper. We first consider the basic structure of the human motor system at the level of the nervous system. After this, we review a number of motor phenomena that have been observed in both humans and animals. See Kelso (1982a) for a more comprehensive introduction.

2.1. Results from neurophysiology

The muscle structure in the higher mammals consists of two basic components. The *muscle fiber* constitutes the main part of the muscle; this includes the alpha neuron, which controls the degree to which the fiber contracts. This neuron is controlled by signals from higher centers of the nervous system via the spinal cord. Running through the center of the muscle fiber is the *muscle spindle*, a tiny sensory organ that sends afferent signals back to the spinal cord in proportion to its level of tension. If a load is placed on a limb, stretching the muscle fiber, the muscle spindle is stretched as well, sending signals to the spinal cord. The spinal cord sends return signals, causing the alpha neuron to fire more rapidly and increasing the contraction force in the main muscle so as to relieve the tension in the muscle spindle (Kelso 1982b).

This local configuration of the alpha neurons and the muscle spindles may explain the low-level stability observed in natural organisms. Pairs of muscles work in opposition to form a servolike mechanism, which maintains the necessary forces of contraction in the respective muscles. For example, when a horse is sleeping standing up, this servomechanism cycle keeps the horse from falling over in a gust of wind. Note that this low-level feedback cycle does not involve higher-level processing, but instead operates directly through the spinal cord. This means that organisms are able to set their limbs in a desired position relatively independent of loads and without higher-level sensory feedback. This observation will play an important role in our theory of motor behavior.

2.2. Results from experimental psychology

Experimental psychology provides an additional set of constraints on theories of human motor performance and learning. Below we summarize a number of well-established motor phenomena that have influenced the development of our theory.

One of the most robust findings is that performance improves with practice, and that this improvement occurs gradually. We also know that motor learning follows the *power law of practice*. This principle states that, if a given amount of practice p leads to an increment i in the quality of performance, then another increment of improvement i will require an order of magnitude more practice p^2 . This phenomenon rules out models that learn too quickly or too slowly.

The literature on motor behavior also distinguishes between *closed-loop* and *open-loop* modes of behavior. Movements performed in closed-loop mode can be altered while in progress as a result of feedback. This feedback indicates when errors have occurred and need to be corrected. Closed-loop

mode is commonly associated with relatively slow movements. Movements performed in open-loop mode can *not* be altered while in progress because feedback is either unavailable or unincorporated. Such movements are generally of a much faster nature (on the order of 200 ms) and, once initiated, are carried out to completion without changes resulting from the detection of errors. Since there is evidence for both modes of performance (Stelmach 1982; Schmidt 1982a), a complete model of human motor behavior should account well for both.

Another well-established motor phenomenon is the trade-off between *speed* and *accuracy*. In many cases, a subject can perform a skill reasonably well at a relatively slow rate, but performance deteriorates if he carries out the skill at a rapid rate. Fitts' law (Fitts and Peterson 1964) indicates a logarithmic relation between speed and accuracy, whereas Schmidt *et al.* (1979) and Meyer *et al.* (1982) report linear speed-accuracy trade-offs. Regardless of the quantitative results, any plausible model of human motor behavior must explain this deterioration of performance with increased speed.

Yet another well-documented result is that humans have a required delay before they are able to incorporate sensory feedback and initiate an alteration in their motor behavior. This delay has been found to be approximately 200 ms (Schmidt 1982a). For example, a subject might be told to perform a motor task but to watch for a signal indicating that he should perform a different task. On the average, it will take 200 ms from the onset of the signal to the initiation of the new motor task. This result places a clear restriction on the minimum cycle time for conscious modification of behavior in models of human motor behavior.

A final phenomenon involves the transfer of skill between limbs or effectors. This is notably demonstrated by comparing handwriting generated with the dominant hand, the dominant hand on a blackboard, the opposite hand, a foot, and the mouth. The similarities between the resulting scripts are so strong that they can be easily recognized as being generated by the same person (see Hollerbach 1979). This indicates the presence of invariants in the human motor system, which any successful theory of motor behavior must explain.

3. A model of motor performance

Skill learning cannot occur in the absence of some performance task, and in this section we describe MAGGIE, a running system that implements our theory of motor performance.² Moreover, any performance task requires some environment in which to perform. Thus our model operates within a simulated environment that mimics the movement of jointed limbs. We begin by describing the specific performance task we are addressing — to carry out a well-specified motor skill with as little error as possible. We then introduce the inputs that are given to the model: the environment, the simulated arm, the sensorimotor interface, and the desired behavior. Next we consider the abstract *data types* and *operators* used by the performance system. Finally, we describe the details of MAGGIE's performance component, deferring discussion of learning issues until Sect. 4.

3.1. The task of feedback control

In a broad sense, our research goal is to develop a general

²A more detailed description of MAGGIE's performance mechanisms can be found in *A performance model of human motor behavior* by Iba and Langley (1987, unpublished manuscript).

computational model of human motor behavior. The performance component addresses only a small part of this task; specifically a simple form of *feedback control* (Wiener 1948). In this task, an agent is given some desired action sequence; it then manipulates its effectors to carry out this sequence as accurately as possible, using its sensors to detect differences between the desired and actual movement.³ The agent's sensors measure certain variables in the environment, which are then compared with the desired behavior. This comparison produces an error signal, which the agent uses to determine future effector applications. These in turn lead to changes in the environment.

For example, consider the task of painting the trim around a window. The desired behavior, or end result, is a coating of paint on the wood around the window but not on the glass itself. In this case, the brush can be viewed as the effector and the painter's eyes are the sensors. Here the significant variable measured in the environment is the gap between the edge of the paint and the glass window, measured at the point where the brush has most recently been applied. The error signal is obtained by observing this gap. If the paint is not reaching the edge of the window, commands are issued that move the brush closer to the glass as the stroke is in progress. Conversely, if paint begins to get on the glass, the brush is moved away from the window.

3.2. Inputs to the model

MAGGIE's performance component incorporates only very general assumptions about the nature of the agent and its environment. This generality requires one to provide additional constraints in the form of four inputs:

- a *simulated environment* in which to operate, along with a set of objects existing in this environment;⁴
- an *effector(s)* or *arm*, which can be manipulated by the agent and which has well-specified relations with other objects in the environment;
- a *sensorimotor interface*, which handles communication between the agent and the environment;
- a *desired behavior*, against which to compare actual effector movement.

The simulated environment

Rosenbaum (1985) has argued that motor behavior implies purposes and that purposes necessitate an agent. However, it makes no sense to refer to an agent in the absence of the environment in which it operates. One can conceive of alternative environments that obey different physical laws than operate in the real world, but since we are interested in human motor behavior we will consider a "standard" environment.

A complete specification of an environment entails listing all the objects and their associated attributes. Interactions between objects must be defined, such as collisions. For the purposes of developing and testing our model, we have developed a simple environment that contains objects with position, length, and velocity, but that ignores mass, friction, and force. In the

experiments reported below, the only objects in the world are the components of the agent's arm.

The arm

Since we are interested in human motor behavior, we will only consider jointed effectors, which we will call *arms*. Although the components or links of the arm are specified as ordinary objects in the environment, the arm merits special treatment here because of additional attributes inherently necessary for jointed movement. A joint is a relation that exists between two objects that are attached to each other.

In a more realistic world, a joint's attributes would include the type of joint, its friction coefficient, its maximum force and velocity, and its range of movement. However, we have restricted ourselves to a simplified type of ball-and-socket joint having only two attributes — the maximum velocity and rotational limits. Currently, we restrict each joint's range of motion to a hemisphere. In addition, rotation about the axis of a link is prohibited.

The sensorimotor interface

An agent will have difficulty interacting with its environment unless it can perceive that environment and control its effectors. In our simulation, both of these are accomplished through a sensorimotor interface. The "motor" component of the interface lets the agent control the motion of its arm. The "sensory" component relays sensory information to the agent about the location of objects, including the arm.

The transfer of sensory information can be viewed as a filtering operation. Essentially, the sensory filter takes a complete description of the world and passes a subset of this information to the agent. MAGGIE accepts two forms of sensory input: visual information giving the absolute positions and velocities of objects; and proprioceptive information giving the relative positions and velocities of the arm's joints (with respect to the previous joint⁵). Visual information is given in a *viewer-centered* representation, whereas proprioceptive information is provided in a *joint-centered* representation. We describe both of these coordinate systems in more detail later.

The motor interface can also be viewed as a filter, since not all possible motor commands are legal in the simulated world. For instance, if the agent specifies an arm movement that would take the arm outside the allowed ranges, the interface "clips" the movement so that the resulting movement is within the allowed limits. Except for such cases, motor control amounts to simply setting the relative positions of arm components to the values specified by the agent in the joint-centered representation. We believe this is a reasonable simplification in light of the motor behavior literature indicating that humans can "set" the positions of limbs without feedback (Kelso 1982b).

Desired behavior

Finally, an agent must be able to determine how well it is performing. The sensorimotor interface provides the position of the arm during a motion, but this must still be compared against some desired behavior. This comparison can be made easily if both the sensory information and the desired behavior are in the same representation. In MAGGIE, desired behavior is given as a sequence of sensory-level descriptions specifying

³Powers (1973) uses a somewhat different terminology. For instance, our "sensors" and "effectors" correspond to his "sensory functions" and "effector functions." Similarly, our "desired action sequence" is the same as his "reference signal."

⁴Some of these objects will correspond to the agent's *effectors*, which it can use to manipulate the environment.

⁵We define the *previous joint* of a particular joint to be the adjacent joint which is closer to the base of the arm in the kinematic chain.

the positions and velocities of the joints in viewer-centered coordinates. Associated with each description is a time value indicating when, with respect to the start of the movement, the sensory-level description should match the sensory input from the environment. This approach lets MAGGIE directly compare the desired behavior to the descriptions generated by the sensory filter.

3.3. Representations and data structures

Any computational model of motor skills requires some representation, and we will use the term *motor schema* for the memory structure that encodes the information necessary to carry out a particular movement. This is similar in spirit to Schmidt's (1982b) use of the term. Unlike Schmidt however, we use the term *motor program* not in reference to a stored memory structure, but for a sequence of motor commands generated dynamically from the motor schema. We will return to the distinction between the motor schema and the motor program shortly.

To be more specific, we will define a motor schema as a sequence of ordered pairs or data points (DP_1, DP_2, \dots, DP_n), where each data point, $DP_i = (t_i, \{ \langle J_k, p, v \rangle, \dots \})$, contains a time value t_i and a set of 3-tuples. The data points, DP_i , are ordered such that the time values, t_i , are in an increasing sequence: $t_i < t_j$ for $i < j$. Each 3-tuple is a set that contains: a joint name J_k , which identifies the joint described by the 3-tuple; a position p , which is the expected position of the joint in three-space at time t_i ; and a velocity vector v , which describes the desired velocity of the joint upon reaching the position p . Each data point contains a set of such 3-tuples; each describes a different one of the effector's joints, though not all joints need be specified. The exceptions are the first and last data point in the schema, which must specify a 3-tuple for every joint. We assume that for any pair of unique 3-tuples in this set, $J_k \neq J_l$. Because a motor schema specifies arm positions at only selected points in time over the course of a movement, we refer to a schema as a *sparse* representation over time.

Let us further define two different types of motor schemas. The first type, a *viewer-centered* schema, represents the position and velocity vectors using Cartesian three-space coordinates with the origin centered at the agent. In contrast, a *joint-centered* schema represents all positions and velocities in local joint-centered coordinates, where each local coordinate system is defined in terms of the previous joint.

In the viewer-centered representation, all joints are described in terms of a single Cartesian coordinate system. We assume this information is available as visual feedback during execution of a skill; it can also be used to describe another agent's actions. Thus, this framework can be used for both recognition and monitoring purposes. In this scheme, the common coordinate system is defined by an origin, set at the base (the first joint) of the effector, and the $x, y,$ and z axes. Given a more complete description of the agent, one can imagine other origins for a viewer-centered schema, such as the agent's eyes. We do not believe the choices of origin and axes will affect performance, assuming a linear translation from the chosen origin to the base of the effector.

In a joint-centered representation, the rotation of each joint is described in terms of its own spherical coordinate system. We assume this form of information is available as proprioceptive feedback during execution; this representation can also be

used to actually generate motor behavior. MAGGIE uses a modified spherical coordinate scheme to represent joint rotations in each associated coordinate frame.⁶ That is, for each joint and its reference coordinate system, a triple $(\theta_x, \theta_y, \rho)$ consists of two angles of rotation, θ_x and θ_y , from the z axis about the x and y axis respectively, and a distance ρ from the origin.⁷

The coordinate system for a particular joint is defined in relation to the previous joint. For instance, the position and orientation of the coordinate system for an elbow would be described in the reference frame of its associated shoulder joint. Likewise, the wrist joint's coordinate system would be described in relation to the elbow joint. The convention we use to fix the coordinate system of joint J_i relative to that of joint J_{i-1} is as follows: the origin is placed at the end of the J_{i-1} th link; the z_i axis is made colinear with the J_{i-1} th link;⁸ the placement of the z_i axis determines the normal for the xy_i plane; the x_i axis is placed within this plane such that the xz_i plane is perpendicular to the xy_{i-1} plane; and the y_i axis is fully constrained by the placement of the other two. Initially, when all joint rotations are zero, the respective x and y axis for each joint are parallel while the respective z axis are all colinear. As any particular joint is rotated, an equivalent rotation (with respect to our convention) is applied to successive coordinate frames. Any set of rotations applied to the joints will always result in the J_{i-1} th link determining the normal to the xy_i plane and the direction of the z_i axis.

TABLE 1. Two representations of motor schemas for the straight-line task

Desired	
$\langle 1, \{ \langle J_1, \langle 100, 0, 0 \rangle, \langle 0, 0, -10 \rangle \}, \langle J_2, \langle 200, 0, 0 \rangle, \langle -2, 0, 2 \rangle \} \rangle$	
$\langle 6, \{ \langle J_1, \langle 95, 0, -31 \rangle, \langle -0.6, 0, -2 \rangle \}, \langle J_2, \langle 180, 0, 20 \rangle, \langle -2, 0, 2 \rangle \} \rangle$	
$\langle 20, \{ \langle J_1, \langle 100, 0, 0 \rangle, \langle 0, 0, 1 \rangle \}, \langle J_2, \langle 100, 0, 100 \rangle, \langle -2, 0, 2 \rangle \} \rangle$	
$\langle 39, \{ \langle J_1, \langle 0, 0, 100 \rangle, \langle -8.63, 0, 0 \rangle \}, \langle J_2, \langle 0, 0, 200 \rangle, \langle -2, 0, 2 \rangle \} \rangle$	
Joint-centered	
$\langle 1, \{ \langle J_0, \langle 0, 1.571, 100 \rangle, \langle 0, 0.1, 0 \rangle \}, \langle J_1, \langle 0, 0, 100 \rangle, \langle 0, -0.018, 0 \rangle \} \rangle$	
$\langle 6, \{ \langle J_0, \langle 0, 1.886, 100 \rangle, \langle 0, 0.021, 0 \rangle \}, \langle J_1, \langle 0, -0.856, 100 \rangle, \langle 0, 0.02, 0 \rangle \} \rangle$	
$\langle 20, \{ \langle J_0, \langle 0, 1.571, 100 \rangle, \langle 0, -0.01, 0 \rangle \}, \langle J_1, \langle 0, -1.571, 100 \rangle, \langle 0, 0.01, 0 \rangle \} \rangle$	
$\langle 39, \{ \langle J_0, \langle 0, 0, 100 \rangle, \langle 0, -0.086, 0 \rangle \}, \langle J_1, \langle 0, 0, 100 \rangle, \langle 0, 0.065, 0 \rangle \} \rangle$	

⁶Our convention for orienting a coordinate frame is noticeably different from the Denavit-Hartenberg notational convention commonly used in robotics work. We have developed our own convention to allow spherical joints, whereas the Denavit-Hartenberg system allows only revolute joints.

⁷This modified spherical coordinate scheme should not be confused with the standard spherical coordinates using triples of (ρ, θ, ϕ) . The details of this method would introduce unwanted complexity to the current discussion but can be found in *A performance model of human motor behavior* by Iba and Langley (1987, unpublished manuscript).

⁸In cases where one of the arm's link components is not straight, we use the endpoints of the link to determine the z axis.

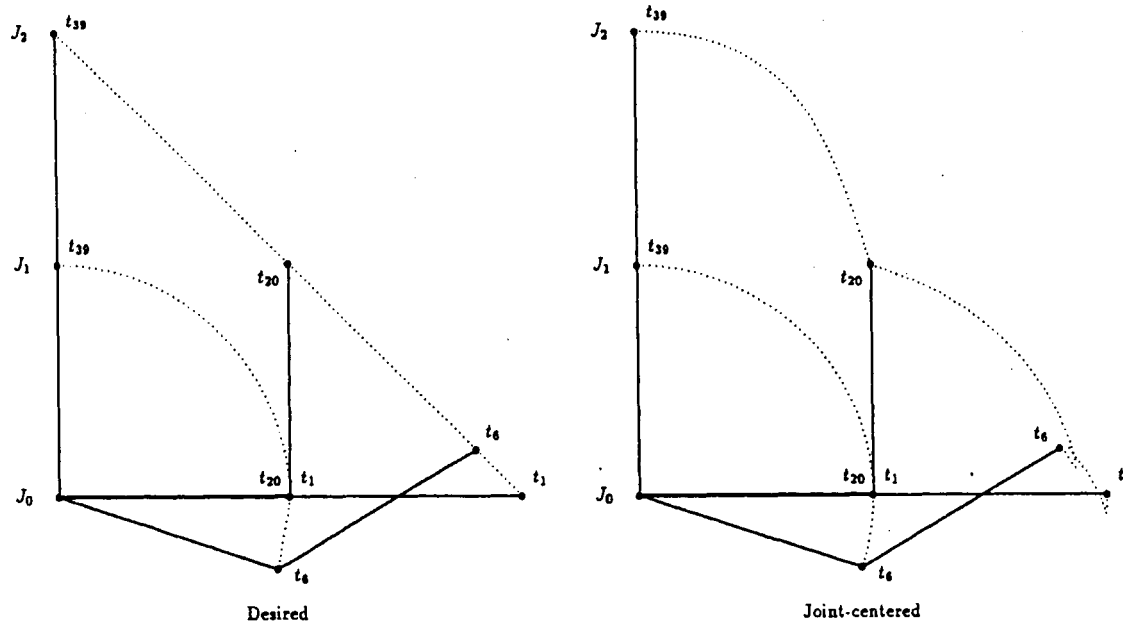


FIG. 1. Traces of the behavior defined by a viewer-centered and joint-centered motor-schema. The arm is shown at the appropriate positions for the times specified by the data points in the schemas (Table 1).

Table 1 shows an example of both types of schemas. The viewer-centered representation example is the specification for the desired behavior given to MAGGIE for the experiments reported later in the paper. The joint-centered example is a direct translation of the desired behavior into the alternate representation scheme. While these numbers may be incomprehensible, the underlying differences become evident when considering the movements they each respectively describe. Figure 1 presents the paths traced by MAGGIE's arm for the respective motor schemas shown in Table 1. Each picture shows the arm in the positions specified by the data points given in the schema. One can see that the joint-centered schema yields quite different results from the desired behavior, although it specifies the same data points as the other.

Although at first glance this dual representation might seem unnecessary, it lends considerably to the model's explanatory power. We propose that humans often acquire an initial motor skill in viewer-centered terms by observing another person performing that skill. The person then translates this description into a joint-centered schema when he attempts to execute the skill himself. It is important to note that these two representations have quite different behavioral characteristics. Each framework is able to specify any point in three-space, but when used by MAGGIE during the generation of behavior, they can produce quite different results. There are two factors acting in conjunction to cause this effect: the sparse representation used to represent motor schemas, and the method used by MAGGIE to generate the positions for the intervening times. As a result, some motions can be much more easily described in one scheme than the other. For instance, the viewer-centered framework can represent a straight line with only two points, whereas a joint-centered schema would require an infinite number of points. This differential power of the two representations predicts that some tasks will be more difficult to perform than others.

The sparse representation of a motor schema seems plausible for storing motor skills in long-term memory, but to actually generate motor behavior one must specify the missing points.

We will use the term *motor program* to refer to such a dense representation for a skill. A motor program can be viewed as a mathematical function that takes a time value as an argument. As output, it returns a set of triples defining, in local joint coordinates, the position for each joint at the given time. It is important to distinguish motor programs from joint-centered schemas. The latter specify the rotations and velocities of joints only at selected times; in contrast, motor programs specify joint rotations at every point in time. Such information can be generated dynamically from a joint-centered schema, as we discuss in the next subsection.

3.4. The performance component

Given a viewer-centered schema that describes some desired behavior, MAGGIE's performance system attempts to carry out this behavior using a specified limb. This involves a number of processes. First, the viewer-centered schema must be translated to a joint-centered representation. The resulting schema must then be "run" by generating an executable motor program and carrying out the specified actions. Simultaneously, the agent must monitor the resulting states, comparing actual positions with the intended positions as given in the viewer-centered schema. Execution and monitoring proceed in parallel until an error is detected. At this point, the system initiates an error correction process to return the limb to the desired path. Below we consider each of these processes in more detail.

From viewer-centered to joint-centered schemas

We assume that the agent begins with a viewer-centered description of a motor skill, presumably learned by observing another's actions or through problem solving. The first step in carrying out a motor skill involves applying an *inverse kinematic transform*⁹ to the viewer-centered schema resulting in a joint-centered representation that can be directly executed. This transformation must be done serially across the joints of a

⁹The details of this transformation are not important to this discussion but can be found in Wylie (1975).

limb, starting with the base joint and considering each successive joint in turn. This process can be time-consuming, and we believe it is one of the factors contributing to the slow and awkward nature of newly acquired skills. With practice, the joint-centered schema becomes fixed in long-term memory. At this point, one can execute the skill without invoking the transformation process; thus one could perform the skill more smoothly. MAGGIE does not currently model the acquisition stage for joint-centered schemas, but transforms a viewer-centered description into a joint-centered one initially and saves it.

Executing the joint-centered schema

Joint-centered schemas only specify the positions and velocities of the joints at selected points in time. Within our framework, the simulation of actual motor behavior requires the specification of either the relative locations or velocities for every joint at every simulated time step. Our *motor program*, as described above, satisfies this requirement since it generates the respective joint positions for every time value. MAGGIE does not store motor programs in memory; the system creates them in real time as it executes the skill. In our theory, this is accomplished by generating a spline for each joint between successive pairs of the points specified in the joint-centered schema.¹⁰ During a movement, when the limb reaches the end of the spline segment between two data points, DP_{i-1} and DP_i , the latter becomes the source and the next point in the sequence, DP_{i+1} , becomes the target for the next spline. This method yields a smooth, continuous curve throughout the execution of the schema.

Monitoring

As we have seen, there is no guarantee that behavior generated by the joint-centered schema will correspond to that specified in the viewer-centered description. Thus, MAGGIE must have some means of detecting divergences, and this is the role of the monitoring process. In order to make the necessary comparisons, the monitoring component uses the viewer-centered schema to generate a "pseudo" motor program. This program cannot be executed by effectors, but it specifies the desired position at each time during execution. When the difference obtained from this comparison becomes noticeable (i.e., exceeds a parameterized threshold), the monitor interrupts execution and invokes the error correction process. MAGGIE monitors for errors at a constant rate; this limits the speed at which it can execute a skill accurately, as we shall see shortly.

Error recovery

Once MAGGIE detects a significant divergence, it must still recover from that error. When invoked by the monitoring process, the error recovery mechanism applies a "burst of force" in a direction that will reduce the size of the error.¹¹ Error recovery involves generating a correction function¹² that modi-

¹⁰We assume that low-level neural circuitry can take relatively sparse inputs from a schema and generate such a motor program in real time.

¹¹This process models the type of corrections that result from error detection at the brain level of the nervous system, and not corrections resulting from servomechanisms at the spinal level.

¹²We use an inverted U type correction function (sin, parabolic, or absolute value) causing a gradual change in the limb's actual movement over the lifetime of the correction process. Note that this introduces another parameter — the type of correction function. Along with this, MAGGIE also allows the duration of the correction process to be adjusted.

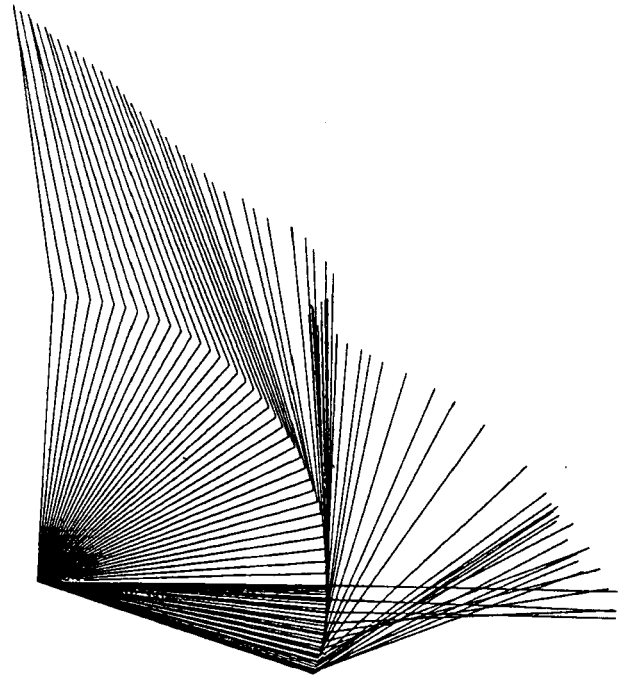


FIG. 2. Successive snapshots taken at regular intervals showing the movement of MAGGIE's arm during the execution, monitoring, and error correction of the straight-line schema.

fies the velocity of the arm for a short period of time. In the default condition, this function is generated such that the area under the curve is the same as the amount of error detected. This means that if the error remains constant, the path of the limb would return to the desired path after error correction has ended. The proportion of the area under the correction function to the size of the error can be adjusted by a compensation parameter.

Depending on the circumstances, the adjustment of this parameter can produce undercorrections or overcorrections. The former occurs in cases where the uncorrected behavior was about to begin reconverging with the idealized path, but had barely exceeded the error detection threshold before this occurred. Since the original motor program would have returned to the desired path on its own, an overcorrection will result. In contrast, undercorrections will occur if the uncorrected behavior is still diverging from the desired path. Such cases will require multiple calls to the error-recovery process.

Figure 2 shows successive snapshots of MAGGIE's arm during the execution of the straight-line schema. Each snapshot was taken at constant intervals so that one can perceive the velocity of the arm at different stages of the movement. Notice that with the monitoring and error-recovery processes, performance approximated the desired behavior more closely than the joint-centered specification alone (right-hand side of Fig. 1). Further improvements can be achieved by learning but we will return to this after quantitatively testing the performance mechanism.

3.5. Behavior of the performance system

We have implemented our model of motor behavior as a running FranzLisp program. Although the theory is independent of a particular arm instantiation, we have tested MAGGIE using a two-jointed arm with roughly human characteristics. Thus, the arm includes an upper arm and a forearm, the former

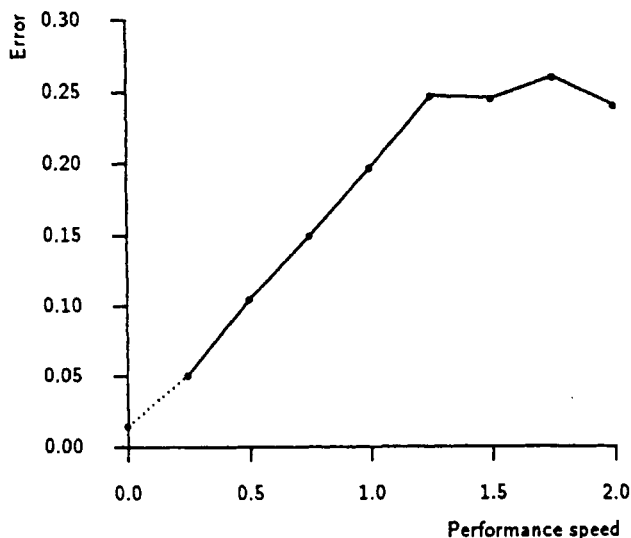


FIG. 3. Increasing error as a function of increasing performance speed.

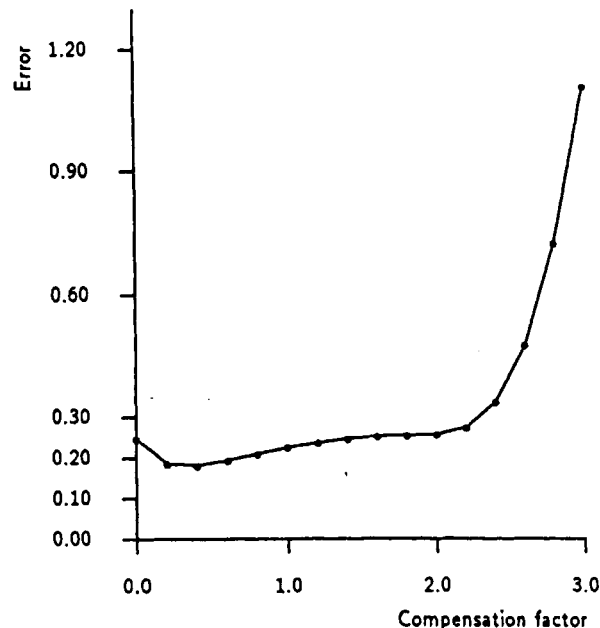


FIG. 4. Amount of error plotted for a range of compensation values.

rotating at a shoulder joint and the latter at an elbow joint. MAGGIE has been implemented to model motor behavior in three dimensions as described in the previous sections, but our tests to date have been run in two dimensions.

Initial studies have focused on a skill that involves moving the hand through a straight line. We have already seen that such motions are easy to describe in viewer-centered coordinates, but that they are extremely difficult for a jointed arm to execute except in trivial cases (Hardy 1984). In a joint-centered representation, every joint must trace the path of an arc. Thus, MAGGIE can never completely achieve straight-line motion; it can only approximate such a path by stringing together a sequence of many small arcs, closely spaced in time. However, this requires learning, and in the current section we will limit our attention to performance phenomena.

In Sect. 2, we noted that one of the most robust findings about human motor behavior involved a trade-off between speed and accuracy. Since MAGGIE can run motor schemas at different speeds, we can test the model's ability to predict this trade-off. Figure 3 shows the results with the "straight-line" skill represented by the motor schemas of Table 1. Clearly, executing this schema at higher speeds leads to greater deviations from the desired motion, i.e., to lower accuracy.¹³ The relation is approximately linear for the range examined, replicating the results reported by Schmidt *et al.* (1979). This effect emerges naturally from the constant rate of monitoring. The more quickly the system runs a joint-centered schema, the fewer times it is able to check for errors and the larger they grow before correction.

We believe that this trade-off demonstrates the continuum between open- and closed-loop behavior. This continuum represents the amount of monitoring occurring during movements. When performing a skill slowly, one can make frequent adjustments, thus operating in a closed-loop mode. As the speed of the skill is increased, the performer can do fewer monitorings thereby moving the performance towards open-

loop mode along this continuum. We address a number of other issues related to this in Iba and Langley (unpublished manuscript, see footnote 2).

We have also noticed another intriguing regularity in MAGGIE's behavior. Recall that the implementation contains a parameter for scaling the amount of correction applied to a given error. Different settings of this parameter lead to different responses to error. Frequently the model detects an error as the deviation is becoming progressively greater, and radical corrective action is in order. However, such a remedy can also result in overcompensation, leading the model to "overshoot" the desired position or trajectory.

Figure 4 presents the effects on the model's behavior as one alters the value of this parameter. When the scheme is run quickly (making monitoring infrequent), increasing the amount of correction may lead to a reduction in the average deviation from the desired path. However, even higher settings can actually produce worse performance at a given speed. For instance, when attempting to follow a straight line, the hand may instead follow a jagged line that cuts back and forth across the desired path. Although we did not plan the model to behave in this fashion, we believe it makes sense. When monitoring occurs frequently, the system makes only minor errors and needs only minor corrective action. A high setting for the correction parameter will cause the system to overcompensate, and this can lead to wild oscillations.

MAGGIE also accounts for the transfer of motor skill between limbs. The model stores each joint-centered schema without reference to the particular limb involved. Thus, the system could take a schema designed for shoulder, elbow, and wrist joints and execute it on a different arm or even on a hip, knee, and ankle. However, to the extent that learning has fine tuned the schema for a given set of joints, performance will degrade drastically when it is run on limbs with different physical characteristics. We have not yet run tests to show the model predicts this behavior, but this is one of our priorities for future research.

In summary, MAGGIE explains a number of well-known phenomena relating to motor performance. However, our main

¹³We measure accuracy as the average deviation from the viewer-centered schema over the course of an execution. This score is computed by totaling the angular deviation at each joint and averaging over the total number of time steps simulated during that execution.

concern is with *learning*. In the following section we describe the model's learning components, along with its empirical behavior on this dimension and its relation to human motor learning.

4. Improving joint-centered motor schemas

Motor learning involves both the acquisition and the improvement of motor schemas. We envision a three-stage process: acquiring an initial viewer-centered schema through observation or problem solving; storing a joint-centered schema by repeatedly transforming the viewer-centered representation; and improving the joint-centered schema through repeated practice. Although we ultimately plan to model each of these processes, our main results involve the final stage of improvement and we will focus on that process in the remainder of the paper.

Let us reiterate the learning task we are attempting to model. Given an initial joint-centered schema that represents a motor skill, along with a viewer-centered schema for the same skill, modify the joint-centered schema so its behavior diverges from the viewer-centered description as little as possible. MAGGIE employs two interacting learning mechanisms to improve its joint-centered schemas. In this section we describe these mechanisms, along with their behavior on the line-drawing task discussed earlier.

4.1. Motivations for learning

Before describing the learning processes themselves, let us review the motivations for improving joint-centered schemas. Recall that MAGGIE's mechanisms for monitoring and error recovery let it execute a joint-centered schema with a moderate level of accuracy. Given this ability, why should the system bother to alter its schemas?

One obvious reason is that one prefers increased accuracy (with speed and attention held constant). That is, the overall error during the performance of a skill at a given speed should be smaller after learning has occurred. The error recovery process alone cannot accomplish this, but with sufficient learning, MAGGIE is able to mimic its viewer-centered schemas with arbitrary accuracy.

A second reason is the desirability of executing a skill either more quickly or with less attention (i.e., in open-loop mode). As stated above, our theory assumes that there is an upper limit on the rate at which monitoring can occur. Similarly, our theory assumes that monitoring is a conscious process requiring scarce attentional resources. However, improving a given joint-centered schema should lessen reliance on monitoring and error correction. This should have two beneficial effects. Learning should let one carry out a skill more rapidly without losing accuracy. It should also let one execute a skill with less attention, freeing resources to carry out other tasks in parallel.

Although monitoring and error correction give immediate aid in carrying out desired behaviors, learning provides a longer-term solution. We have said that viewer-centered and joint-centered representations lead to different interpolated behavior, but learning lets the latter approximate the former. For instance, one can simulate straight-line behavior with a joint-centered schema by adding a number of more densely spaced points to the schema, creating a sequence of very small arcs.

Thus, our learning model relies heavily on the distinction between viewer-centered and joint-centered schemas and the different representational powers of these two frameworks. It

also relies on the performance assumptions covered in the last section, specifically the mechanisms for monitoring and error recovery. This seems desirable; a learning system should not be independent of assumptions about representation and performance.

4.2. Assumptions and learning operators

The model makes a number of basic assumptions about the nature of the learning process. The main claims are:

- The model learns only when the errors are detected in the execution of a schema; MAGGIE is driven by failures rather than successes.
- Learning occurs after the trial during which errors occurred; this implies some memory for arm positions and velocities during the trial. We will call this the *motor buffer*.
- MAGGIE retains only one version of each joint-centered schema in its long-term memory; thus, it carries out a form of hill-climbing learning (Langley *et al.* 1987).

With these constraints in mind, let us consider the model's two operators for schema improvement.

Recall that MAGGIE specifies a motor schema as a sequence of points, each describing the locations and velocities of a set of joints. This suggests two natural approaches to modifying joint-centered schemas:

- modifying one of the fields in an existing point for a particular joint; or
- removing an existing point from the schema or adding an entirely new point.

The first of these seems the less drastic action, since it leaves the basic structure of the schema unaltered. However, there may be limits to what can be accomplished by modifying numeric values; in such cases, one may need to revise the schema structure by adding or removing points. To review, each data point consists of a time value and a set of 3-tuples. Each 3-tuple consists of a joint identifier, a position vector, and a velocity vector. In principle, any of the values in a data point, except the joint identifier, may be modified; however, our experiments to date have only considered adjusting the velocities. Nor have we examined the deletion of points from schemas; in its current form, MAGGIE only adds points.

TABLE 2. The learning algorithm

-
1. Select the failure point in motor-STM with the largest error.
 2. Find the best possible modification to the point values.
 3. Find the percentage improvement over the current form of the schema.
 4. If improvement > bias (should reflect density with respect to time),
 - (a) Then alter the schema with best (velocity) alteration found in step 2;
 - (b) Else add the selected failure point to the schema.
-

4.3. MAGGIE's learning algorithm

We have seen that error detection invokes the error recovery process, but it also triggers learning. Whenever the path of a joint diverges noticeably from the desired path, the monitor stores this "failure point" along with the currently desired point into the motor buffer. This lets MAGGIE delay learning until after the execution has been completed.

Table 2 presents the model's basic learning algorithm. Since a number of errors may occur in a given trial, the first step

involves selecting a failure point on which to base modification. In principle, one could use all errors noted in the trial to alter the schema. However, this would lead to much more rapid learning than observed in humans, so we limit the model to a single point. One explanation for this limit is that motor memory decays before additional points can be accessed. In any case, MAGGIE selects that failure point in the motor buffer with the largest error. We assume that larger errors require more processing than smaller ones and therefore are most easily available since they decay less rapidly. Thus, larger errors are reduced before smaller ones, giving a learning curve roughly similar to the power laws observed in human skill acquisition.¹⁴

Once MAGGIE has selected a failure point, it must decide between its two basic learning operators. One could simply add a new point wherever an error was detected. Since points specified in the schema are generally guaranteed to be reached at their respective times, performance would improve. Furthermore, the time between respective points would decrease, giving less occasion for deviating from the desired path. We have run experiments with this strategy and achieved good results (Langley *et al.* 1987). However, adding a point to a schema is a more radical operation than modifying the values of an existing data point. Therefore, MAGGIE incorporates a *bias* factor that discourages the addition of new data points in favor of modifications to existing points.

The current model only considers adjusting the values of velocity vectors. Furthermore, MAGGIE considers modifying only the two data points delimiting the segment of the schema containing the time of failure. That is, for the straight-line schema of Table 1, if the selected failure point was at time t_1 , then the second and third data points would be considered for modifications and would be said to "contain" the failure point. However, selecting among real-valued modifications still leads to an infinite branching factor, so we require some simplifying assumptions to help reduce the effective search space. We employ both an intelligent next state generator to propose a small number of possible alterations, and an evaluation function to select among the alternative modifications generated.

For two data points, DP_i and DP_j containing the failure point, the amount of adjustment A applied to each is inversely proportional to their respective distances (in time) from the failure point. That is, the closer the failure point is to DP_i , the larger the adjustment made to DP_i 's velocity. Although this does not guarantee an optimal modification, it is a reasonable alteration based upon the limited information available from the motor buffer. The amounts of adjustment that are considered are $A_i = Em_i$ to DP_i and $A_j = Em_j$ to DP_j , where m_i and m_j are computed by the following:

$$m_i = \frac{(t_F - t_i)}{(t_k - t_i)}$$

and

$$m_k = \frac{(t_k - t_F)}{(t_k - t_i)}$$

for failure point t_F , error vector E , and the associated time values for DP_i and DP_j , t_i and t_j .

¹⁴This is certainly not the only explanation of power laws; we direct readers to Rosenbloom and Newell (1987) for an alternative computational theory.

Based on this calculation, MAGGIE considers four possible ways of pairwise incrementing and decrementing the two data points discussed above by their respective amounts.¹⁵

$$\begin{array}{ll} (DP_i + Am_i, DP_j + Am_j) & (DP_i + Am_i, DP_j - Am_j) \\ (DP_i - Am_i, DP_j + Am_j) & (DP_i - Am_i, DP_j - Am_j) \end{array}$$

It may seem more straightforward to select the appropriate combination of adjustments by inspection of the error vector alone, but this is much more complicated than it appears. This results from the nature of the interpolation process used by the performance component and is considered elsewhere.¹⁶

Once the four combinations are generated, MAGGIE evaluates each alternative by generating a partial motor program for each case. The system examines the predicted performance of each program at the failure point, selecting the combination that minimizes error.¹⁷ If MAGGIE would proceed to compare this new partial motor program with the result of adding a completely new data point, it would always favor the creation of new points. This is because the comparison is made at the same point that the new data point would be added, therefore revealing no error. It is for this reason that we have included a bias against this response. As long as the best of the four possible modifications results in a percentage improvement greater than the bias factor, the modification is preferred. Only when none of the modifications considered can sufficiently improve the schema (at the time of failure), will a new data point be added to the schema. This bias factor has the effect of knocking the system out of local minima.

4.4. Behavior of the learning system

MAGGIE's learning methods are independent of a limb's dimensions and rotational constraints, but we have tested the system with the same arm described in Sect. 3. We have examined the system's learning behavior by running a number of experiments, again in two dimensions. We describe the results below, along with their relation to data on human motor learning.

Naturally, we would expect that as MAGGIE detects errors and modifies its joint-centered schema, its performance will improve on later executions. Figure 5 shows the model's average divergence from the desired path on eight successive trials with the "straight-line" schema. The figure indicates that the system's performance gradually improves with practice, modeling the basic phenomenon in human motor learning.

As we mentioned before, improvement over time is not sufficient for a psychologically plausible model of motor learning. The nature of MAGGIE's learning mechanism theoretically leads to a power law learning curve. This should arise from attending to the largest errors first, causing the most dramatic improvements in performance during early stages of practice. However, our preliminary results are inconclusive. A problem we face is that the reported human learning curves have measured performance either as the number of units produced per unit time, or as the average time to completion of

¹⁵Here we use + and - loosely for notational convenience. We assume that appropriate velocity vectors in the data point structure are accessed and updated according to the arithmetic operator indicated.

¹⁶Iba, W., and Langley, P. 1987. A performance model of human motor behavior. Unpublished manuscript.

¹⁷Another method would involve executing all four revised schemas in their entirety and comparing their resulting overall deviations. However, this would be very expensive computationally and we find it unlikely that humans carry out such computations unconsciously.

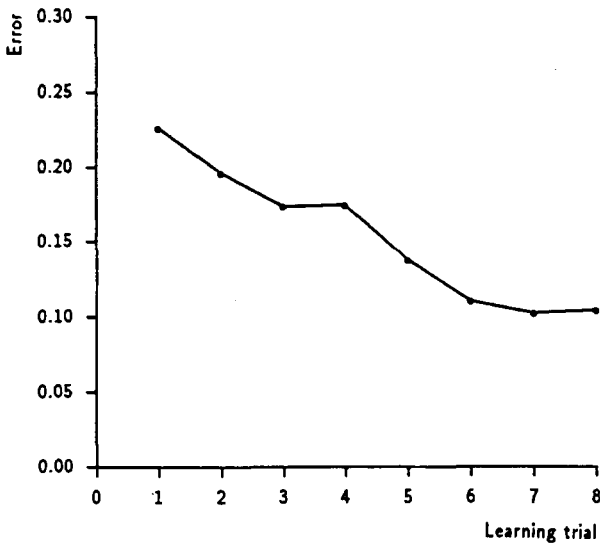


FIG. 5. Decreasing error as a result of learning after successive trials.

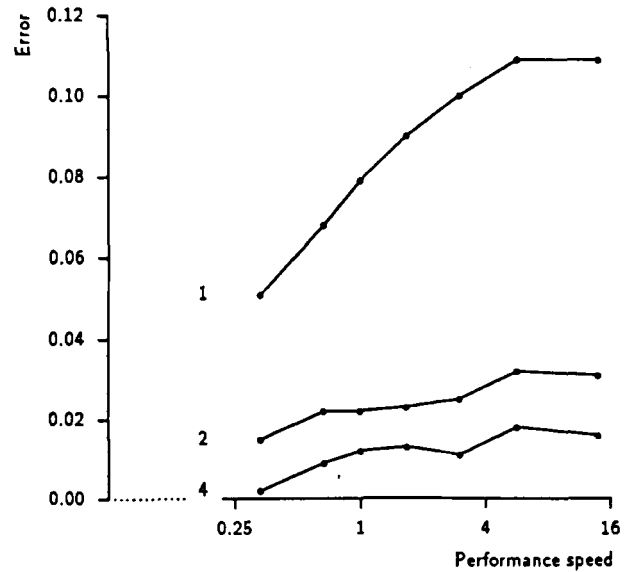


FIG. 6. Speed vs. accuracy after one, two, and four learning trials.

task. We must find new ways to test MAGGIE since our results are given as average error and therefore are not directly comparable. We also need to be able to run learning sessions over many more trials than we have to date. While we are not able to make strong claims at this time, the results displayed in the figure are encouraging.

Our model of performance accounted for another robust finding: the trade-off between speed and accuracy. However, it seems natural to expect learning to affect this relation, and Fig. 6 shows how MAGGIE's speed-accuracy trade-off changes with practice on the straight-line schema. As the skill level improves, the trade-off curve becomes flatter and eventually disappears entirely. That is, modifications to the schema allow the system's behavior to approximate the desired behavior even without monitoring. This means that MAGGIE can execute the schema at a higher speed — even though there are fewer chances for monitoring — without seriously decreasing its accuracy.¹⁸

Another experiment considered the effect of MAGGIE's practice speed upon its learning rate; Fig. 7 shows the results. When learning at a high speed, performance improves for several trials and then stabilizes, but at a high error level. In contrast, slower practice leads to almost immediate asymptotic behavior but at a much lower error rate.¹⁹ The slope of the learning curves for all speeds tested are approximately the same; the difference between them are how quickly they stop improving and at what error level this occurs. This is a result of the practice speed on the number of possible monitorings. Since MAGGIE's learning is triggered by the monitoring process, there are fewer opportunities for improvement. This sug-

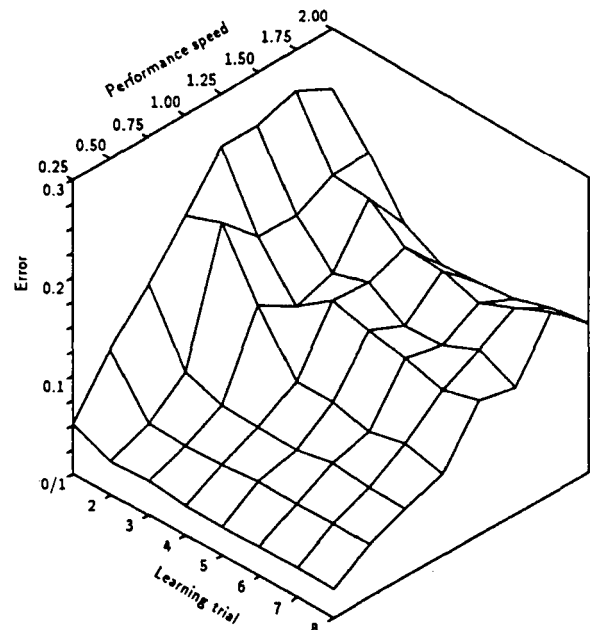


FIG. 7. Average error plotted as a function of practice speed and learning trial.

¹⁸This constitutes an untested second prediction of the learning theory: the speed-accuracy trade-off should disappear with practice. Actually, it implies a third prediction as well: learning should produce a transition in skills from closed-loop processing to open-loop mode, in which feedback is unnecessary and a motor skill can be carried out accurately with little attention. To our knowledge, neither of these behaviors have been reported in the psychological literature.

¹⁹This constitutes a fourth interesting and testable prediction about human motor learning: the speed at which a skill is practiced influences both the learning rate and the limit of possible improvements.

gests both an upper and lower limit on the effect, determined by the maximum and minimum number of possible monitorings. The lower limit occurs when the movement is so slow that the arm barely moves at all between monitorings; no learning would occur here because performance is already at the threshold of detectable errors. The upper limit arises when the speed of execution is so fast that the agent never gets a chance to monitor during movement; no learning would occur here either, since no failures would be detected.

We have already discussed MAGGIE's two learning mechanisms and the bias parameter that determines which one will be applied in a given situation. This parameter suggested a final experiment, in which we examined the model's learning behavior for different values of the parameter. Figure 8 shows

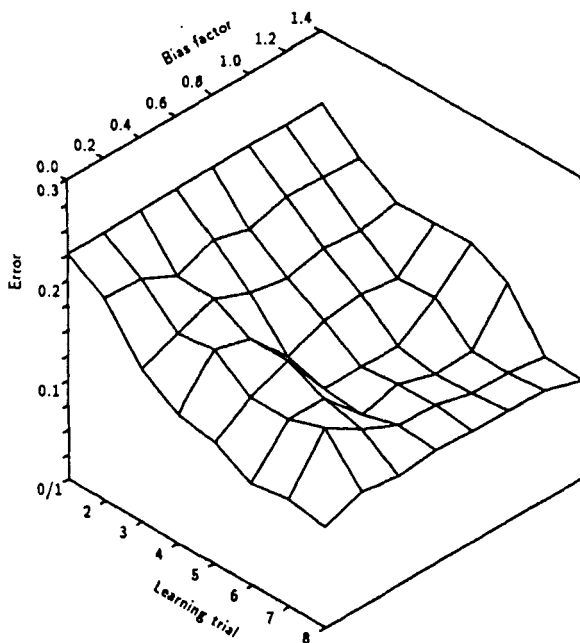


FIG. 8. Error plotted as a function of bias value and learning trial.

the set of learning curves that result on the straight-line task. At one extreme, we set the bias very small; this led MAGGIE to learn exclusively by altering velocities. At the other extreme, giving the bias a large value led the model to learn only by adding points. Intermediate biases led to mixed learning strategies.

Naturally, MAGGIE begins at the same level of error regardless of the bias factor. The figure also reveals that the system arrives at essentially the same performance level after eight learning trials, but that the learning rates vary according to the bias. However, the relation is definitely not monotonic. Note that a "medium" bias yields a trough or canyon in which the learning rates are greater than for either high or low biases.²⁰

This behavior can be explained by supposing that overly conservative and overly rash learning strategies each have drawbacks. When the system is too reluctant to add new points, velocity changes give only minor improvement; when the system adds points too eagerly, it has little chance to fine tune the altered structure. The medium bias level led MAGGIE to add a new point every two or three trials, causing change in structure when needed but also giving the system time to fine tune the restructured schema.

In Sect. 2.2, we discussed a number of phenomena in the psychological literature that constrain plausible models of human motor behavior. Here we have presented the results from a number of experiments used to test MAGGIE. These results support the psychological plausibility of our theory although not all are conclusive. We have also presented a number of predictions made by the theory. We are continuing to look for results in the literature that would confirm or falsify these predictions. We are encouraged to further develop and test our system in light of these results and predictions.

²⁰We do not see how this aspect of MAGGIE's learning behavior leads to any testable predictions about human motor learning. However, it does help us understand the workings of the model itself.

5. Discussion

Now that we have described our theory of motor behavior and its implementation in MAGGIE, let us turn to its evaluation. We begin on a positive note, considering the theory's successes in terms of both explained and predicted phenomena. We then examine its limitations, along with directions for future work that these suggest.

5.1. Successes of the theory

MAGGIE and the theory it implements make up a coherent computational theory of human motor skills. To our knowledge, it constitutes the only such theory in existence that accounts for observed phenomena. Roboticians have proposed a variety of computational methods for motor control, but these were never intended as models of human behavior. Similarly, psychologists have developed theories of motor behavior, but for the most part, these have not been instantiated in computational terms. The theory provides the first computational explanation of motor performance and learning; this is its most basic contribution to cognitive science and artificial intelligence.

Let us briefly review the phenomena that MAGGIE successfully models. These are:

- the trade-off between speed and accuracy;
- the distinction between closed-loop and open-loop behavior;
- the transfer of motor skills across limbs; and
- the gradual improvement of motor performance with practice.

In the model, these behaviors emerge from representational differences between the two types of motor schemas combined with the limited rate at which monitoring can occur. Learning improves the joint-centered descriptions and thus reduces reliance on monitoring and error correction.

The same mechanisms lead to several predictions about human motor behavior. These include:

- a reduction in the speed-accuracy trade-off with practice;
- a gradual transition from closed-loop behavior to open-loop behavior;
- an effect of practice speed on learning rate and asymptotic performance.

Each of these predictions should be simple to test, and we look forward to feedback from experimentalists along these lines. If the predictions are accurate, this will be convincing evidence in favor of the theory. If not, then the manner in which they are disconfirmed will suggest directions in which to modify the model.

We evaluate the relative worth of a theory based upon both the phenomena that it explains as well as the predictions that it makes. The predictions should not be a part of the phenomena to be explained. That is, one should take a set of phenomena and develop a theory to account for these. Then the designers should step back and ask "what other predictions of phenomena does the theory make?" Additionally, the complexity of a theory in terms of constraining parameters helps determine its value or promise. While our theory will not be the last word on motor behavior, it rates highly with respect to all three of these criteria.

5.2. Limitations of the theory

Whether or not the above predictions are borne out, the

existing theory has a number of limitations that require extensions. For instance, there is mounting evidence for a preparation stage prior to the onset of movement (Kelso 1982*b*). This suggests that motor programs are generated before motion is initiated, whereas our current theory assumes it is generated dynamically. Another problem involves the componential transfer of motor skills, which suggests that such skills are organized hierarchically. The current theory only handles skills at a single level and makes no proposal for their integration into larger structures.

We have also focused on motor skills that involve no objects other than the agent. This includes a wide class of skills, but much of human motor behavior involves interactions with other objects. In some cases, the agent has direct influence over the object during only part of the schema. For instance, in the first stages of throwing a ball one has immediate control over the ball's location. However, once the ball is released, its trajectory is almost entirely a function of the arm's earlier motion. In such cases, improvement requires taking into account knowledge of results (e.g., the quality of the ball's flight), and the current theory makes no statement about this aspect of learning.

Nor must all motor learning involve modification of joint-centered schemas; there are undoubtedly cases in which the initial viewer-centered schema can be improved as well. For example, suppose the agent acquires a viewer-centered description by observing another agent perform some task. There are a number of ways such learning by imitation can lead to inaccurate schemas: attentional limitations may cause important details to be omitted from the learned schema; the imitated agent's limbs may differ in important ways from the learner's limbs; or the imitated agent may simply execute the skill poorly itself. In each of these cases, the learner would need to improve its viewer-centered description, either by observing the other agent many times or by reasoning from knowledge of results.

Even in terms of joint-centered learning, we have limited our treatment to the transfer of skill from viewer-centered descriptions, but other approaches are possible. One might also create joint-centered schemas from proprioceptive sense data. For instance, one might use problem-solving methods to generate a sequence of motor schemas for achieving some goal. Information in the motor buffer (Sect. 4.2) resulting from the execution of this sequence could then be used as the basis for a new joint-centered description. Initially, this schema might not even have an associated viewer-centered schema, so the improvement techniques currently implemented in MAGGIE would not find much use.

We mentioned earlier the process of translation between viewer-centered and joint-centered schemas, and that MAGGIE did not model this process in a satisfactory fashion. Nor have we explained the manner in which "mental practice" can improve performance without explicit practice. Finally, conscious experimentation may also play a role in human motor learning. We have seen that adding new points can knock a schema out of a local maximum, and extreme perturbations on locations or velocities might have a similar effect. These all seem fruitful directions for further research, and we plan to address them in future versions of MAGGIE.

Most important, we must test the model's behavior on additional motor skills so we can evaluate its generality. The computational results described earlier in the paper were obtained with runs on the straight-line schema. We anticipate analogous

results on other schemas, but we must test this prediction and carefully examine any differences that arise. Such differences need not invalidate the theory, since they may also arise in human behavior. But we must clearly run MAGGIE on a wide variety of motor skills and attempt to understand the full range of its behavior.

5.3. Summary

In this paper we have presented a computational theory of human motor behavior and its implementation in MAGGIE. The model assumes that two distinct representations underly motor skills, one based on viewer-centered coordinates and the other using joint-centered descriptions. Each type of schema consists of a sequence of "points" that describe the locations and velocities of relevant joints at successive points in time. Motor behavior involves translating from the viewer-centered scheme to the joint-centered scheme, and then interpolating intermediate points to produce actual behavior.

We found that these two frameworks have different representational capabilities, each describing some motions better than the other. For this reason, the translation process is inherently imperfect and MAGGIE must continually monitor its behavior for deviations from the desired path. When errors become noticeable, the system invokes an error-recovery process that attempts to put the movement back on track. The model assumes a lower limit on the frequency of monitoring, and this limitation led naturally to the speed-accuracy trade-off and the distinction between closed-loop and open-loop behavior.

MAGGIE learns only in response to a detected error. In some cases, the system alters the velocity of a point in the schema; in others, it actually adds a new point. Both learning methods ultimately lead to improvements in performance, letting the joint-centered schema more closely approximate the viewer-centered description. This learning process accounted for a number of observed behaviors and predicted additional phenomena that have not been reported in the literature.

Our initial tests of the model have been encouraging, but more work remains to be done. We need to study MAGGIE's behavior on a variety of motor skills, and we need to extend the system along a number of dimensions. We feel that MAGGIE is a good initial model, but we have far to go before achieving a truly general and robust theory of human motor performance and learning.

Acknowledgements

We would like to thank John Gennari, Glenn Iba, and Oliver Selfridge for helpful comments on an earlier draft of this paper. We would also like to thank the members of the UCI World Modelers project — David Benjamin, John Gennari, Kevin Thompson, and Patrick Young — for numerous discussions leading to the ideas reported herein. Maggie and Kirsten Iba provided many insights into early childhood motor control. This work was supported by Contract MDA 903-85-C-0324 from the Army Research Institute and by a gift from Hughes Aircraft Company.

ARBIB, M. A. 1982. Modeling neural mechanisms of visuomotor coordination in frog and toad. *In* Lecture notes in biomathematics, Vol. 45. Edited by S. Amari and M. A. Arbib. Springer-Verlag, Berlin, West Germany.

DAVIS, W. J. 1976. Organizational concepts in the central motor networks of invertebrates. *In* Neural control of locomotion. Edited by R. M. Herman, S. Grillner, P. S. G. Stein, and D. G. Stuart.

- Plenum Press, New York, NY.
- FIKES, R. E., HART, P. E., and NILSSON, N. J. 1972. Learning and executing generalized robot plans. *Artificial Intelligence*, 3: 251–288.
- FITTS, P. M., and PETERSON, J. R. 1964. Information capacity of discrete motor responses. *Journal of Experimental Psychology*, 67: 103–112.
- HARDY, S. 1984. Robot control systems. In *Artificial intelligence: tools, techniques, and applications*. Edited by T. O'Shea and M. Eisenstadt. Harper and Row, New York, NY.
- HOLLERBACH, J. 1979. Understanding manipulator control by synthesizing human handwriting. In *Artificial intelligence: an MIT perspective*. Edited by P. H. Winston and R. H. Brown. MIT Press, Cambridge, MA.
- KELSO, J. A. S. *Editor*. 1982a. Human motor behavior: an introduction. Lawrence Erlbaum Associates, Hillsdale, NJ.
- . 1982b. Concepts and issues in human motor behavior: coming to grips with the jargon. In *Human motor behavior: an introduction*. Edited by J. A. S. Kelso. Lawrence Erlbaum Associates, Hillsdale, NJ.
- LANGLEY, P., GENNARI, J. H., and IBA, W. 1987. Hill-climbing theories of learning. *Proceeding of the Fourth International Workshop on Machine Learning*, Morgan Kaufmann, Irvine, CA.
- MEYER, D. E., SMITH, J. E. K., and WRIGHT, C. E. 1982. Models for the speed and accuracy of aimed movements. *Psychological Review*, 89(5): 449–482.
- POWERS, W. T. 1973. Feedback: beyond behaviorism. *Science*, 179: 351–356.
- ROSENBAUM, D. A. 1985. Motor programming: a review and scheduling theory. In *Motor behavior: programming, control, and acquisition*. Edited by H. Heuer, U. Kleinbeck, and K.-H. Schmidt. Springer-Verlag, Berlin, West Germany.
- ROSENBLOOM, P., and NEWELL, A. 1987. Learning by chunking: a production system model of practice. In *Production system models of learning and development*. Edited by D. Klahr, P. Langley, and R. Neches. MIT Press, Cambridge, MA.
- SCHMIDT, R. A. 1982a. More on motor programs. In *Human motor behavior: an introduction*. Edited by J. A. S. Kelso. Lawrence Erlbaum Associates, Hillsdale, NJ.
- . 1982b. The schema concept. In *Human motor behavior: an introduction*. Edited by J. A. S. Kelso. Lawrence Erlbaum Associates, Hillsdale, NJ.
- SCHMIDT, R. A., ZELAZNIK, H., HAWKINS, B., FRANK, J. S., and QUINN, J. T. JR. 1979. Motor-output variability: a theory for the accuracy of rapid motor acts. *Physiological Review*, 86(5): 415–451.
- SEGRE, A. M. 1987. Explanation-based learning of generalized robot assembly plans. Technical Report UILU-ENG-2208, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL.
- STELMACH, G. E. 1982. Motor control and motor learning: the closed-loop perspective. In *Human motor behavior: an introduction*. Edited by J. A. S. Kelso. Lawrence Erlbaum Associates, Hillsdale, NJ.
- SWARTZ, N. M. 1984. Arm dynamics simulation. *Journal of Robotic Systems*, 1(1): 83–100.
- WIENER, N. 1948. *Cybernetics: control and communication in the animal and the machine*. Wiley, New York, NY.
- WYLIE, C. R. 1975. *Advanced engineering mathematics*. 4th ed. McGraw-Hill, New York, NY.

Wayne Iba received his B.A. in computer and information sciences at the University of California, Santa Cruz, and he is now a Ph.D. student in the Department of Information and Computer Science at the University of California, Irvine. His research has focused on machine learning, including concept acquisition using previously learned concepts, learning relational concepts from noisy examples, and models of human motor learning. He has also worked in the area of game playing and is interested in all aspects of human cognition and perception.

Pat Langley is an associate professor in the Department of Information and Computer Science at the University of California, Irvine. Before coming to UCI, Dr. Langley was a research scientist at Carnegie-Mellon University, where he received his Ph.D. in cognitive psychology in 1979. He has taught numerous courses in artificial intelligence, including graduate seminars on the topic of machine learning. Dr. Langley has published over 60 papers, the majority describing research on AI approaches to learning; this work has spanned a variety of learning problems, including language acquisition, heuristics learning, scientific discovery, concept formation, and motor learning. He is coauthor of the book *Scientific Discovery: Computational Explorations of the Creative Processes* and is coeditor of the collection *Production System Models of Learning and Development*. He is also the executive editor of the journal *Machine Learning*.